

DISSERTATION

submitted to the
Combined Faculties of the Natural Sciences
and Mathematics

of the
Ruperto-Carola-University of
Heidelberg, Germany
for the degree of
Doctor of Natural Sciences

Put forward by
ANN-KATHRIN PERREVOORT
born in Gronau (Westfalen)
Oral examination: 22nd of May, 2018

Sensitivity Studies on New Physics
in the Mu3e Experiment and
Development of Firmware for the
Front-End of the Mu3e Pixel Detector

Referees:

Prof. Dr. Niklaus Berger

Prof. Dr. Stephanie Hansmann-Menzemer

In loving memory of Heinrich Perrevoort
from whom I inherited a good portion of optimism
and the confidence to meet this challenge.

ABSTRACT

The upcoming *Mu3e* experiment aims to search for the lepton flavour violating decay $\mu^+ \rightarrow e^+e^-e^+$ with a final sensitivity of one signal decay in 10^{16} muon decays. As this decay is forbidden in the Standard Model of particle physics, an observation would be a clear indication for New Physics beyond the Standard Model.

The ambitious sensitivity goal imposes several challenges on the experiment, among which a high track and data rate. In the context of this thesis, an on-line algorithm for hit sorting by time stamp is developed as part of the firmware of the pixel detector front-end. System integration tests in test beam measurements with pixel sensor prototypes yield data transfer efficiencies in excess of 99.95 % and demonstrate the suitability of the concept for the data acquisition system of *Mu3e*.

The feasibility of background-free $\mu \rightarrow eee$ searches in the first phase of the experiment is confirmed in detailed simulation studies. A single event sensitivity to the signal decay of $2.3 \cdot 10^{-15}$ is estimated—an improvement by more than two orders of magnitude compared to existing limits. In addition, a model-independent effective field theory approach is developed and used to estimate the sensitivity for different New Physics interactions.

Furthermore, the physics programme of the experiment is expanded by exploring two additional processes involving New Physics: lepton flavour violating two-body decays of the muon and muon decays with promptly decaying dark photons. For both channels, it is shown that a highly competitive sensitivity can be achieved.

This work thus highlights the unique possibilities of New Physics searches with rare muon decays in the *Mu3e* experiment.

ZUSAMMENFASSUNG

Das *Mu3e*-Experiment plant eine Suche nach dem Leptonfamilienzahl-verletzenden Zerfall $\mu^+ \rightarrow e^+ e^- e^+$ mit einer finalen Sensitivität von einem Signalzerfall in 10^{16} Myonzerfällen. Da dieser Zerfall im Standardmodell der Teilchenphysik verboten ist, wäre eine Beobachtung ein eindeutiges Zeichen für Neue Physik jenseits des Standardmodells.

Das ambitionierte Sensitivitätsziel birgt eine Reihe von Herausforderungen für das Experiment, unter anderem hohe Spur- und Datenraten. Im Rahmen dieser Arbeit wird ein Algorithmus zum Zeit-Sortieren der Treffer des Pixeldetektors in Echtzeit für die Front-End-Firmware entwickelt. In Systemintegrationstests in Teststrahlungsmessungen mit Pixelsensor-Prototypen werden Datentransfer-Effizienzen oberhalb von 99.95 % erzielt, was die Eignung des Konzepts für das Datennahmesystem von *Mu3e* nachweist.

Die Machbarkeit von Untergrund-freien $\mu \rightarrow eee$ -Suchen in der ersten Phase des Experiments wird in detaillierten Simulationsstudien bestätigt. Eine Einzelereignis-Sensitivität von $2.3 \cdot 10^{-15}$ für den Signalzerfall wird erwartet—eine Verbesserung um mehr als zwei Größenordnungen gegenüber existierenden Ausschlussgrenzen. Außerdem wird eine Modell-unabhängige Methode basierend auf effektiven Feldtheorien entwickelt, mit dem die Sensitivität für verschiedene Wechselwirkungen Neuer Physik bestimmt wird.

Zusätzlich wird das Physikprogramm des Experiments erweitert indem zwei weitere Prozesse Neuer Physik untersucht werden: Leptonfamilienzahl-verletzende Zweikörperzerfälle des Myons und Myonzerfälle mit prompt zerfallenden dunklen Photonen. Für beide Zerfallskanäle wird gezeigt, dass eine konkurrenzfähige Sensitivität für das Experiment erreicht werden kann.

Diese Arbeit betont somit die einzigartigen Möglichkeiten für Suchen nach Neuer Physik in seltenen Myonzerfällen mit dem *Mu3e*-Experiment.

CONTENTS

I	INTRODUCTION	1
1	THE STANDARD MODEL AND BEYOND	3
1.1	The Standard Model of Particle Physics	3
1.2	Lepton Flavour Violation and Physics Beyond the Standard Model	9
1.2.1	Lepton Flavour Violating Muon Decays in Effective Theories	13
1.2.2	Previous and Current Experiments Searching for Lepton Flavour Violation	16
1.3	Family Symmetries and the Familon	22
1.3.1	Experimental Limits on $\mu \rightarrow eX$	24
1.4	Dark Photons	27
1.4.1	Experiments Searching for Dark Photons	30
2	THE MU3E EXPERIMENT	35
2.1	Signal and Background	35
2.2	Experimental Concept	39
2.2.1	Muon Beam	43
2.2.2	Target	44
2.2.3	Magnet	44
2.2.4	Pixel Tracker	45
2.2.5	Timing Detector	51
2.2.6	Data Acquisition	52
2.2.7	Track and Vertex Reconstruction	54
2.3	Simulation of the Mu3e Experiment	59
2.3.1	Generators of Physics Processes	60
2.4	Upgrades to the Phase I Experiment	61
2.4.1	Phase II of the Mu3e Experiment	61
2.4.2	The Mu3e-Gamma Experiment	62
II	DATA ACQUISITION	63
3	THE FRONT-END OF THE MU3E PIXEL DETECTOR	65
3.1	Readout on the Pixel Sensor	65
3.2	Data Acquisition at the Front-End of the Mu3e Pixel Detector	68
3.2.1	Data Rates at the Front-End	70
3.3	The MuPix Telescope	74
3.3.1	Data Path	75
3.4	The Hit Sorter Entity for MuPix7	77

3.4.1	Writing Process	78
3.4.2	Reading Process	82
3.4.3	Resource Usage	86
3.4.4	Tests of the MuPix7 Hit Sorter	86
3.5	The Hit Sorter Entity for MuPix8	87
3.5.1	Writing Process	90
3.5.2	Reading Process	92
3.5.3	Resource Usage	97
3.5.4	Tests of the MuPix8 Hit Sorter	97
3.6	Summary on the Pixel Detector Front-End	100
III	PHYSICS STUDIES	107
4	THE SIGNAL DECAY MEDIATED VIA EFFECTIVE OPERATORS	109
4.1	Searches for $\mu \rightarrow eee$ Events with Mu3e	110
4.1.1	Simulation of Signal and Background Events	112
4.1.2	Reconstruction	113
4.2	Background Study	116
4.2.1	Pre-Selection	116
4.2.2	Selection and Cuts	116
4.3	Kinematics of the Signal Decay	131
4.4	Sensitivity to Effective Operators	138
4.4.1	Effective Operators	138
4.4.2	Dipole Operator	141
4.4.3	4-Fermion Operators	142
4.4.4	Selection Cuts	145
4.4.5	Sensitivity	150
4.5	Summary on $\mu \rightarrow eee$ Searches	153
5	LEPTON FLAVOUR VIOLATING TWO-BODY DECAYS OF THE MUON	161
5.1	Searches for $\mu \rightarrow eX$ with Mu3e	161
5.2	Estimating the Sensitivity	163
5.2.1	Toy Monte Carlo Approach	165
5.3	Simulation of $\mu \rightarrow eX$ Events	169
5.3.1	Event Generator	169
5.3.2	Reconstruction	170
5.3.3	Selection and Cuts	170
5.4	Expected Branching Fraction Limits	182
5.5	Analysis in Momentum and Polar Angle	190
5.5.1	Event Distributions	190
5.5.2	Expected Branching Fraction Limits	191
5.6	Momentum Calibration	197
5.6.1	Validation of the Method	199

5.7	Momentum Bin Size	201
5.8	Comparison to Existing Experimental Limits	204
5.9	Summary on Lepton Flavour Violating Two-Body Decays	209
6	SEARCHES FOR ELECTRON-POSITRON RESONANCES	213
6.1	Searches for Dark Photons	213
6.2	Estimating the Sensitivity	216
6.2.1	Toy Monte Carlo Approach	217
6.3	Simulation of Muon Decays with Prompt Dark Photon Decay	221
6.3.1	Event Generator	221
6.3.2	Reconstruction	223
6.3.3	Selection and Cuts	224
6.4	Expected Sensitivity to the Branching Fraction	231
6.5	Expected Sensitivity to the Kinetic Mixing Parameter	244
6.6	Systematic Uncertainties	245
6.6.1	Fraction of Bhabha Events	247
6.6.2	Bin Size	250
6.6.3	Restrictions on the Three-Particle Mass	250
6.7	Projections of Dark Photon Searches with the Mu3e Experiment	251
6.8	Summary on Electron-Positron Resonance Searches .	260
7	SUMMARY, CONCLUSION, AND OUTLOOK	265
7.1	Firmware on the Pixel Detector Front-End	266
7.2	Sensitivity Studies on New Physics	267
IV	APPENDIX	273
A	DATA ACQUISITION	275
A.1	MuPix Data Protocol	275
B	THE SIGNAL DECAY MEDIATED VIA DIFFERENT EFFECTIVE OPERATORS	279
B.1	Effective Operators	279
B.2	Dalitz Analysis	282
B.3	Dalitz Plots	283
C	LEPTON FLAVOUR VIOLATING TWO-BODY DECAYS OF THE MUON	287
C.1	Data Selection	287
C.2	Expected Branching Fraction Limits	292
C.3	Momentum Calibration	303
D	PUBLICATIONS	305
E	BIBLIOGRAPHY	307
F	ACKNOWLEDGMENTS	327

OUTLINE OF THE THESIS

Searches for lepton flavour violation are sensitive probes for New Physics beyond the Standard Model. Particularly stringent constraints can be achieved with lepton flavour violating muon decays as muons can be produced in large quantities.

The upcoming *Mu3e* experiment searches for the lepton flavour violating decay $\mu^+ \rightarrow e^+e^-e^+$ with an ultimate sensitivity to one signal decay in 10^{16} muon decays. This poses several experimental challenges which are met with a novel experimental concept and innovative detector technologies. In essence, a high number of muon decays needs to be observed with excellent momentum, time, and vertex resolution. A key component of the experiment is thus a tracking detector built from thin silicon pixel sensors.

In the context of this thesis, parts of the data acquisition firmware for the pixel detector front-end are developed. The pixel sensors feature an on-chip readout system but the implemented readout sequence does not strictly conserve the time order of the hits. This time order needs to be restored on the front-end boards in order to enable event filtering at a later stage of the data acquisition chain. The implementation and evaluation of the firmware-based time sorting algorithm are outlined in chapter 3.

The sensitivity of the *Mu3e* experiment to the signal decay $\mu \rightarrow eee$ greatly depends on the performance of the detector. In chapter 4, the background suppression capability and resulting signal sensitivity are evaluated by means of a detailed detector simulation. The signal decays are further modelled in an effective field theory approach in order to account for variations in the kinematics induced by different New Physics models. This enables model-independent estimates of the sensitivity and comparisons to limits of other lepton flavour violation searches.

The large momentum acceptance and precise tracking of electrons and positrons in the *Mu3e* detector as well as the large number of expected muon decays enables competitive searches for other New Physics processes besides $\mu \rightarrow eee$ decays. Two examples are studied in this work. A search strategy and sensitivity study for lepton flavour violating two-body decays of the muon $\mu \rightarrow eX$ are presented in chapter 5. Muon decays involving dark photons that immediately decay into an electron-positron pair $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ are investigated in chapter 6.

Part I

INTRODUCTION

The Standard Model of particle physics describes the elementary particles and their interactions remarkably well, and yet it is known to be incomplete. Ideas for extensions are abundant but observational hints for New Physics are scarce. Motivated by neutrino oscillations and the discrepancy between measurements and theoretical calculations of the anomalous magnetic moment of the muon, the lepton sector seems to be a promising candidate for New Physics searches.

THE STANDARD MODEL AND BEYOND

Today's understanding of particle physics is condensed in the Standard Model (SM). Numerous experimental observations are explained with great precision by the SM. Despite its success, the SM is known to be incomplete because it does not incorporate gravity, lacks an explanation of neutrino masses and provides no candidate for dark matter. Experiments are thus searching in a large range of observables for hints of New Physics beyond the SM.

In this chapter, the SM is reviewed and a few shortcomings are discussed. As the violation of lepton flavour is a promising signature of New Physics, selected models beyond the SM are presented with focus on lepton flavour violating muon to electron transitions and dark photons. In addition, the current and future experimental situation of such searches is summarized.

1.1 THE STANDARD MODEL OF PARTICLE PHYSICS

The Standard Model of particle physics is a mathematical description of the fundamental constituents of matter and their interactions—with the exception of gravity—in the form of a quantum field theory.

The particle content and force carriers of the interactions in the SM are shown in figure 1.1. The matter fields consist of twelve fermions with spin $\frac{1}{2}$, six leptons and six quarks, as well as their antiparticles. Leptons and quarks are arranged in three generations. The first generation contains the elementary particles that make up everyday matter: the electrically charged electron lepton and its neutral counterpart, the electron-neutrino ν_e , as well as the up-quark with charge $+\frac{2}{3}$ and the down-quark with charge $-\frac{1}{3}$. The latter two are bound to protons and neutrons and thus form atomic nuclei. Moving on to the second and third generation, the particles have higher masses but otherwise share the characteristics of their first generation equivalent. Neutrinos are in general massless in the SM. The leptons in the second and third generation are the muon μ and muon-neutrino ν_μ and the tau τ and tau-neutrino ν_τ , and the quarks are the charm- and strange-quark and the top- and bottom-quark, respectively.

Interactions are mediated through the exchange of gauge bosons

with spin 1. Quantum electrodynamics (QED) describe the electromagnetic interaction acting through the exchange of massless virtual photons. All electrically charged particles have electromagnetic interactions, i. e. the charged leptons e , μ and τ , as well as all quarks.

The strong force acts on the colour charge of the particles and thus only on quarks. It is mediated by a set of eight massless gluons. Due to the strong interaction, quarks are always bound to hadrons in nature: three quarks of different colour form colour-neutral baryons, and a quark-antiquark pair with the same colour and anti-colour forms a colour-neutral meson. Combinations of larger numbers of quarks and anti-quarks are allowed as long as they remain colour-neutral. Evidence for tetraquarks and pentaquarks exists [1, 2]. The strong interaction is described by quantum chromodynamics (QCD).

The weak force acts on the weak isospin under which all fermions are charged. It is thus the only way neutrinos interact in the SM. Interactions are mediated by the massive W^\pm ($m_W = 80.4 \text{ GeV}$) and Z bosons ($m_Z = 91.2 \text{ GeV}$)¹. The weak interaction has a few peculiarities. It acts only on the left-handed component of the particle fields and is thus maximally parity-violating. And in interactions mediated by W bosons, the flavour of quarks (u , d , c , s , t , b) can be changed whereas the quark flavour is conserved for Z exchange. The transition probability is described with the Cabibbo–Kobayashi–Maskawa (CKM) matrix.

The electromagnetic and weak interaction can be unified above the electroweak symmetry breaking scale of about 246 GeV . In the Higgs mechanism a complex scalar field is introduced which causes the spontaneous breaking of the electroweak symmetry due to a non-zero vacuum expectation value. As a result, a massive scalar particle arises, the Higgs boson. It is the only fundamental scalar in the SM. The W and Z bosons acquire mass in the Higgs mechanism. Fermion masses are generated in Yukawa interactions with the Higgs field.

One of the goals of particle physics is to further unify the electroweak and strong force in a *Grand Unified Theory* (GUT), and ultimately also with the gravitational force in a *Theory of Everything*.

The SM can be condensed to the following Lagrangian [4]

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{QCD}} + \mathcal{L}_{\text{EW}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}} , \quad (1.1)$$

¹ Throughout the thesis, natural units are used, i. e. $\hbar = c = 1$. Energies, momenta and masses are therefore given in units of eV .




Figure 1.1: The Standard Model of particle physics. Figure adapted from [3].

which is invariant under transformation of the group $SU(3)_c \times SU(2)_L \times U(1)_Y$. $SU(3)_c$ denotes the gauge group of the strong interaction, and $SU(2)_L$ and $U(1)_Y$ the gauge groups of the weak isospin and hypercharge Y . The SM Lagrangian has contributions from the strong (QCD) and electroweak (EW) interactions as well as from the Higgs field and Yukawa interactions

$$\mathcal{L}_{\text{QCD}} = \bar{q} (i\gamma^\mu \partial_\mu - m_q) q - g_s (\bar{q} \gamma^\mu T_a q) G_\mu^a - \frac{1}{4} G_{\mu\nu}^a G_a^{\mu\nu}, \quad (1.2)$$

$$\begin{aligned} \mathcal{L}_{\text{EW}} = & \bar{\psi} \gamma^\mu \left(i\partial_\mu - \frac{g}{2} \boldsymbol{\tau}_L \cdot \mathbf{W}_\mu - g' \frac{Y}{2} B_\mu \right) \psi \\ & - \frac{1}{4} \mathbf{W}_{\mu\nu} \cdot \mathbf{W}^{\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu}, \end{aligned} \quad (1.3)$$

$$\mathcal{L}_{\text{Higgs}} = \left| \left(i\partial_\mu - \frac{g}{2} \boldsymbol{\tau} \cdot \mathbf{W}_\mu - g' \frac{Y}{2} B_\mu \right) \phi \right|^2 - V(\phi), \quad (1.4)$$

$$\mathcal{L}_{\text{Yukawa}} = -G_1 \bar{\psi}_L \phi \psi_R + iG_2 \bar{\psi}_L \boldsymbol{\tau}_2 \phi^* \psi_R + \text{h. c.} . \quad (1.5)$$

The Lagrangian of equation 1.2 describes the interactions of quarks (Dirac spinor q and mass m_q) and gluons (gauge fields G_μ^a and field strength tensor $G_{\mu\nu}^a$). The coupling strength of the strong interaction is g_s . γ^μ are the Dirac matrices and T_a the Gell-Mann matrices. In the electroweak Lagrangian of equation 1.3, interactions of fermions




Figure 1.2: Decays of the muon in the Standard Model.

(Dirac spinor ψ) with the gauge fields of $SU(2)_L$ (\mathbf{W}_μ and coupling strength g) and of $U(1)_Y$ (B_μ and coupling strength g') and thus the photon, W^\pm and Z boson are summarized. $\boldsymbol{\tau}_L$ are the Pauli matrices, in this case only acting on left-handed fermions. The Lagrangian of equation 1.4 describes interactions of the W^\pm , Z and photon with the complex scalar Higgs field ϕ . $V(\phi)$ denotes the Higgs potential. Yukawa interactions of fermions and the Higgs are given in equation 1.4, in which G_1 and G_2 are matrices of Yukawa couplings.

MUONS IN THE STANDARD MODEL The muon is the charged lepton of the second generation. It is about 200 times heavier than the electron with a mass of 105.66 MeV [1] and decays with a lifetime of 2.20 μs via the weak interaction exclusively into electrons, neutrinos and photons. The dominant decay mode with a branching fraction of almost 100% is the so-called Michel decay $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$. With a branching fraction of about 1.4% an additional photon is emitted $\mu^+ \rightarrow e^+ \gamma \bar{\nu}_\mu \nu_e$, and with a branching fraction of $3.4 \cdot 10^{-5}$ this photon converts internally to an $e^+ e^-$ pair $\mu^+ \rightarrow e^+ e^- e^+ \bar{\nu}_\mu \nu_e$. The branching fractions are summarized in table 1.1, and the corresponding Feynman diagrams are shown in figure 1.2.

Measurements of the Michel decay parameters of the decay $\mu \rightarrow e \nu \bar{\nu}$ are precision tests of the electroweak theory. The latest measurements have been performed by the *TWIST* collaboration [5]. No deviation from the Standard Model has been found.

LIMITATIONS OF THE STANDARD MODEL With the discovery of the Higgs boson [6, 7], all particles predicted by the SM have been observed. Although the SM successfully describes a wide range of experimental observations with great precision, it is known to have some shortcomings and is thus regarded to be incomplete. An obvious deficiency is that the SM provides no description of gravity. Although gravitational effects are negligible at typical en-

Decay mode	Branching fraction	Comment
$\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$	$\approx 100\%$	
$\mu^+ \rightarrow e^+ \gamma \bar{\nu}_\mu \nu_e$	$(1.4 \pm 0.4)\%$	subset of $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$ with $E_\gamma > 10\text{MeV}$
$\mu^+ \rightarrow e^+ e^- e^+ \bar{\nu}_\mu \nu_e$	$(3.4 \pm 0.4) \cdot 10^{-5}$	experiment with trans- verse momentum cut $p_e^T > 17\text{MeV}$

Table 1.1: Branching fractions of SM muon decay modes (values taken from [1]).

ergy scales in high energy physics, they dominate the interactions at cosmological scales and thus need to be addressed.

Many astrophysical observations cannot be explained by the amount and distribution of ordinary *luminous* matter in the universe and thus hint at the existence of another form of matter, the *dark matter*, which interacts gravitationally and has only very weak interactions with photons, if at all. The observations include the distributions of rotational velocities of stars orbiting the galactic centre [8–11], gravitational lensing [12], anisotropies in the cosmic microwave background (CMB) [13, 14], and the formation of large scale structures [15]. Although these phenomena can be addressed in a variety of theories, only dark matter explains them all at once. Furthermore, the observed accelerated expansion of the universe [16] is attributed to the so-called dark energy. Recent measurements by the *Planck* space observatory estimate the contribution to the energy-matter density in the universe of ordinary baryonic matter to be 5%, of dark matter to be 26%, and of dark energy to be 69% in the Λ CDM (Lambda cold dark matter) model [17]. The SM provides no explanation for dark matter or dark energy and thus falls short to describe 95% of the matter and energy content of the universe².

A further open question is the imbalance of matter and antimatter in the observable universe known as baryon asymmetry. Matter and antimatter are assumed to be produced in equal amounts during the Big Bang. For the asymmetry to arise, a few conditions must hold—the so-called Sakharov conditions [18]—among which the violation of the combined charge and parity symmetry (CPV). The

² Although only weakly interacting, neutrinos are no viable dark matter candidates. Due to the small mass, neutrinos remain relativistic during the expansion of the universe after the Big Bang. The observed large scale structures however require dark matter particles to be non-relativistic.

weak interaction is known to violate CP though the amount of CPV is insufficient to explain the baryon asymmetry. In fact, it is a further puzzle of the SM that no CPV is observed in strong interactions although QCD does not require CP to be conserved (strong CP problem).

The observation of neutrino oscillations gives even more compelling reasons for the need to extend the SM [19–24]. The transition from a neutrino of one flavour to another violates the conservation of the individual lepton flavour³. Although not explicitly conserved by the interactions, lepton flavour is an accidental symmetry of the SM due to the particle content. The neutrinos are further required to be massive in order to perform mixing. The SM provides however no mechanism for neutrinos to acquire mass. So far, the extremely small mass of the neutrinos is not yet directly measured. Only the differences of the squares of the masses are known from the observed oscillations as well as upper limits for example from measurements of the end point of β decay spectra. The mixing is parametrized by the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix. In fact, the neutrino is a mysterious particle. Up to now, it is not proven whether neutrinos are Dirac (neutrino and anti-neutrino exist) or Majorana (the neutrino is its own anti-particle) particles.

A few more open questions concern the hierarchy problem (fine-tuning is needed to keep quantum corrections to the Higgs mass small enough so that it remains consistent with the measured value of 125 GeV) and the number of generations (exactly three generations of neutrinos are observed in Z decays). Moreover, the SM has 26 free parameters⁴ for which it lacks explanation.

Apart from neutrino oscillations, only a few more observations give hints for physics beyond the SM. These include discrepancies in the measurements of the proton charge radius. Results obtained with muonic hydrogen spectroscopy deviate by more than 7 standard deviations from results obtained with spectroscopy of normal hydrogen and with electron-proton scattering [25–28]. It is however not yet excluded that the discrepancy is related to an experimental artefact. A further discrepancy is observed for the measured and calculated value of the anomalous magnetic moment of the muon. The current deviation amounts to 3.6 standard deviations [29]. Huge efforts are ongoing in theory to reduce the uncertainties of the cal-

³ The leptons in each generation are assigned an individual flavour quantum number: electron flavour L_e for e and ν_e , L_μ for μ and ν_μ , and L_τ for τ and ν_τ .

⁴ The 26 parameters comprise the Yukawa couplings of all twelve fermions, the coupling strengths of the gauge interactions, the vacuum expectation value and mass of the Higgs, the mixing angles of the CKM and PMNS matrix, and the strong CP phase.

ulation. In addition, a new experiment—*E989* at Fermilab—to measure the muon magnetic moment with increased precision is currently in commissioning [30, 31]. Thus, the discrepancy is expected to be confirmed or disproved in the coming years. In addition, several deviations from the SM prediction in B decays involving muons or taus are reported [32–35]. All these discrepancies seem to be related to flavour and muons in particular.

Nonsurprisingly, numerous models exist trying to extend the SM and to solve these puzzles. Predictions of these models need experimental confirmation or refutation. Experiments can be roughly classified in two categories. At the high-energy frontier, experiments aim at the direct production of new particles at colliders. As the new particles are expected to be very heavy, a high centre of mass energy is required. Experiments of the high-precision or high-intensity frontier search for small deviations from the SM that might be induced by new particles in quantum loops or by new interactions with small coupling strengths. This requires usually the observation of a large number of events, hence the name high-intensity frontier. Searches for rare decays such as the $\mu \rightarrow eee$ decay that is going to be investigated in the upcoming *Mu3e* experiment belong to the high-intensity frontier.

1.2 LEPTON FLAVOUR VIOLATION AND PHYSICS BEYOND THE STANDARD MODEL

In the SM, the flavour of the individual leptons is conserved. With the observation of neutrino oscillations however it became evident that the lepton flavour of neutral leptons is violated in nature. As a consequence, also the flavour of charged leptons will be violated in some order in perturbation theory.

The contribution to the branching fraction of $\mu \rightarrow e\gamma$ and similar lepton flavour violating muon decays that stem from the neutrino masses and mixing alone is extremely small [36–38]

$$\mathcal{B}(\mu \rightarrow e\gamma) = \frac{3\alpha}{32\pi} \left| \sum_{i=2,3} U_{\mu i}^* U_{ei} \frac{\Delta m_{i1}^2}{M_W^2} \right|^2 < 10^{-54}, \quad (1.6)$$

due to a GIM (Glashow-Iliopoulos-Maiani) mechanism [39] like suppression. Herein, α denotes the fine structure constant, $U_{\alpha i}$ the elements of the neutrino mixing matrix, Δm_{ij}^2 the differences of the squared neutrino masses, and M_W the mass of the W boson. The large suppression is caused by the smallness of Δm_{ij}^2 ($\mathcal{O}(10^{-3} \text{ eV}^2)$)




Figure 1.3: Lepton flavour violating muon decays in the SM extended by neutrino mixing and masses.





Figure 1.4: Past and future experiments searching for charged lepton flavour violating decays. Adapted from [40].


and the comparably large W boson mass ($M_W = 80.4 \text{ GeV}$). The corresponding Feynman diagrams with neutrino mixing in a loop are shown in figure 1.3. The branching fraction for $\mu \rightarrow eee$ is even smaller because of the additional vertex of the photon conversion.

As SM contributions are absent and the contribution from neutrino mixing is negligible, the violation of the flavour of charged leptons becomes an ideal probe to search for New Physics. Any observation would be an unambiguous signal for physics beyond the SM. This is why over the past decades many experiments have been performed and also for the future more experiments are planned pushing the sensitivity level further down (see figure 1.4).



(a) SUSY: loop level diagram with sleptons and neutralinos.

(b) SUSY: tree level diagram with sneutrinos.



(c) Doubly charged scalar.

(d) Extended electroweak sector.

Figure 1.5: The $\mu \rightarrow eee$ signal decay mediated in various models beyond the SM.

The violation of the charged lepton flavour is closely linked to the generation of neutrino masses and thus naturally arises in many New Physics models that aim at explaining the observed neutrino oscillations [36]. A few models are briefly discussed to give a glimpse on the range of physics associated with lepton flavour violation in the charged lepton sector.

In supersymmetric theories (SUSY), each particle of the SM is paired with a corresponding *superpartner*. Fermions get a bosonic supersymmetric counterpart (*sfermions*) and bosons a fermionic one (*gauginos* and *Higgsinos*). The popularity of SUSY models stems from the fact that they naturally solve the hierarchy problem. Each quantum loop contribution of SM particles to the Higgs mass is compensated by a corresponding loop of SUSY particles. As the cancellation cannot be perfect, SUSY must be a broken symmetry. In addition, in many supersymmetric models the lightest neutralino is a viable dark matter candidate. Up to now, SUSY has evaded observation [41] which also supports the statement that if existent SUSY must be a broken symmetry.

Whereas leptons acquire mass in Yukawa interactions, slepton masses arise from the breaking of SUSY. A misalignment of the mass matrices of leptons and sleptons can induce flavour violation of the SM particles and thus enhance the decay rates to observable levels. An example for $\mu \rightarrow eee$ is shown in figure 1.5a [42]. If R-parity⁵ is conserved, the decay is mediated at lowest order via gamma or Z penguin diagrams—in this case with smuons ($\tilde{\mu}$), selectrons (\tilde{e}) and neutralinos ($\tilde{\chi}$)—or box diagrams. In R-parity violating SUSY, $\mu \rightarrow eee$ can also be mediated in tree level diagrams [43, 44] (see figure 1.5b).

In combination with the seesaw mechanism, also the smallness of the masses of the left-handed neutrinos can be explained [36, 45]. SUSY models can be further assumed to be part of a larger symmetry group at a high energy scale and become thus a potential realization of grand unified theories [46].

In left-right symmetric models of the weak interaction, parity is conserved above the symmetry breaking scale which leads amongst others to right-handed neutrinos and neutrino masses. Such models can be realized with a Higgs triplet consisting of a neutral (Δ^0), singly (Δ^+) and doubly charged scalar (Δ^{++}) [47–56]. Processes like $\mu \rightarrow eee$ and muonium-antimuonium oscillation $\mu^- e^+ \leftrightarrow \mu^+ e^-$ can occur at tree level (see figure 1.5c).

Models with an extended gauge sector or with extra dimensions give rise to new heavy vector bosons such as Z' [57–59]. With flavour

⁵ R-parity is a quantum number which differentiates SM and SUSY particles.

off-diagonal couplings, a Z' can mediate $\mu \rightarrow eee$ decays (see figure 1.5d).

The abundance of models with lepton flavour violation turns searches for decays like $\mu \rightarrow e\gamma$ and $\mu \rightarrow eee$ into powerful probes of New Physics. It is however difficult to pick a representative model. Therefore, in the following a model-independent approach in effective field theory is applied.

1.2.1 Lepton Flavour Violating Muon Decays in Effective Theories

An effective theory is a physics model that describes all physics effects below a certain energy scale. It can be regarded as a low-energy approximation of a usually more complex underlying theory that is derived by integrating out heavy degrees of freedom not relevant at the considered scale. In fact, the SM can be interpreted as an effective field theory itself.

New Physics can be parametrized in effective theories in a model-independent way. With the underlying theory of New Physics being unknown, a bottom-up approach is adopted. The SM Lagrangian is extended with effective operators of increasing dimension. The operators are constructed respecting the involved particle content and symmetries. Above a chosen dimension, operators are omitted as the contribution at low energies decreases with the dimension of the operator. The coupling of the operators is typically expressed in terms of a dimensionless Wilson coefficient C and some powers of energy scale which are required to preserve an overall dimension of four for every term in the Lagrangian.

In the case of lepton flavour violating transitions of muons to electrons, the following Lagrangian can be formulated [60, 61] (see also [36, 46, 62])

$$\begin{aligned}
\mathcal{L}_{\text{eff}} = & \mathcal{L}_{\text{QED}} + \mathcal{L}_{\text{QCD}} \\
& + \frac{1}{\Lambda^2} \left[C_L^{\text{D}} \mathcal{O}_L^{\text{D}} + \sum_{f=q,\ell} \left(C_{ff}^{\text{VLL}} \mathcal{O}_{ff}^{\text{VLL}} + C_{ff}^{\text{VLR}} \mathcal{O}_{ff}^{\text{VLR}} + C_{ff}^{\text{SLL}} \mathcal{O}_{ff}^{\text{SLL}} \right) \right. \\
& \quad \left. + \sum_{h=q,\tau} \left(C_{hh}^{\text{TLL}} \mathcal{O}_{hh}^{\text{TLL}} + C_{hh}^{\text{SLR}} \mathcal{O}_{hh}^{\text{SLR}} \right) + C_{gg}^{\text{L}} \mathcal{O}_{gg}^{\text{L}} \right. \\
& \quad \left. + \text{L} \leftrightarrow \text{R} \right] + \text{h.c.} .
\end{aligned} \tag{1.7}$$

The Lagrangian is valid below a scale Λ within the range $m_W \geq \Lambda \gg m_b$ with m_W being the mass of the W boson and m_b the mass of the heaviest involved quark, the b quark. \mathcal{L}_{QED} and \mathcal{L}_{QCD} are the QED

and QCD Lagrangians of the SM. The weak interaction is omitted as it is not relevant for the considered lepton flavour violating processes. The effective operators \mathcal{O} are given by

$$\mathcal{O}_L^D = e m_\mu (\bar{e} \sigma^{\mu\nu} P_L \mu) F_{\mu\nu}, \quad (1.8)$$

$$\mathcal{O}_{ff}^{VLL} = (\bar{e} \gamma^\mu P_L \mu) (\bar{f} \gamma_\mu P_L f), \quad (1.9)$$

$$\mathcal{O}_{ff}^{VLR} = (\bar{e} \gamma^\mu P_L \mu) (\bar{f} \gamma_\mu P_R f), \quad (1.10)$$

$$\mathcal{O}_{ff}^{SLL} = (\bar{e} P_L \mu) (\bar{f} P_L f), \quad (1.11)$$

$$\mathcal{O}_{hh}^{SLR} = (\bar{e} P_L \mu) (\bar{h} P_R h), \quad (1.12)$$

$$\mathcal{O}_{hh}^{TLL} = (\bar{e} \sigma_{\mu\nu} P_L \mu) (\bar{h} \sigma^{\mu\nu} P_L h), \quad (1.13)$$

$$\mathcal{O}_{gg}^L = \alpha_s m_\mu G_F (\bar{e} P_L \mu) G_{\mu\nu}^a G_a^{\mu\nu}. \quad (1.14)$$

Herein, $P_{R(L)}$ is the operator of the right(left)-handed projection, $F_{\mu\nu}$ the photon and $G_{\mu\nu}^a$ the gluon field strength tensor. e denotes the electron charge, m_μ the muon mass, α_s the strong coupling constant and G_F the Fermi coupling constant. As usual, γ_μ are the gamma-matrices and $\sigma_{\mu\nu} = \frac{i}{2} [\gamma_\mu, \gamma_\nu]$. The fermion fields $f \in \{u, d, c, s, b, e, \mu, \tau\}$ include all leptons and quarks except for the top quark. In the fields $h \in \{u, d, c, s, b, \tau\}$, all of the quark fields of f are contained and the tau lepton.

Relevant for $\mu \rightarrow eee$ decays are dipole operators (superscript D, see also figure 1.6a), and scalar (S) and vector (V) operators that mediate the decay as contact interaction of four fermions (see figure 1.6b). For muon to electron conversion on nuclei, additional scalar and tensor (T) operators as well as gluon (index gg) operators are included in the Lagrangian.

The effective theory approach allows to compare constraints of various lepton flavour violating observables. It has to be taken into account, that the energy scale μ of $\mu \rightarrow e\gamma$ and $\mu \rightarrow eee$ decays is at m_μ whereas the scale of $\mu N \rightarrow eN$ conversion is typically 1 GeV. For a comparison, they have to be propagated to a common scale, for instance to $\mu = m_W$. In this transition, the evolution of the renormalization group causes the Wilson coefficients to run and mix. In this way, also contributions from operators including quark or tau fields arise for example for $\mu \rightarrow eee$.

In figure 1.6, resulting constraints from current and projected future limits from $\mu \rightarrow e\gamma$, $\mu \rightarrow eee$ and muon conversion searches are shown for selected operators. At the high scale m_W , two operators are assumed to be non-zero. The Wilson coefficients are evolved to the respective scale of the observable, e. g. m_μ for $\mu \rightarrow eee$, and the

predicted rates are compared to the experimental or projected limits. The still allowed regions in parameter space for the investigated coefficients at m_W are shown.

In figure 1.7a, dipole and vector operators are compared. Typically, $\mu \rightarrow e\gamma$ puts strong constraints on the dipole operator, but the precision of future conversion experiments will yield comparable constraints. The vector operator is better constrained by $\mu \rightarrow eee$ than in the other two observables as it directly mediates $\mu \rightarrow eee$ but enters at loop level in the other two observables. An interesting feature to notice is the cancellation of vector and dipole operators in the case of $\mu \rightarrow e\gamma$ and $\mu N \rightarrow eN$ that results in blind regions in the parameter space. As the vector operator enters these processes at loop level, the cancellations occur. Therefore, this is not observed for $\mu \rightarrow eee$.

A dipole and a scalar operator are compared in figure 1.7b. Again, cancellation effects occur in $\mu \rightarrow e\gamma$ and $\mu N \rightarrow eN$ as the operators mediate these decays at loop level. $\mu \rightarrow eee$ is mediated at tree level for both operators and thus yields the strongest constraints.

Figure 1.7c shows the parameter space of a dipole operator and a scalar operator involving b quarks. As expected, $\mu N \rightarrow eN$ conversion is a powerful tool to test operators of quark fields. Through mixing, also $\mu \rightarrow e\gamma$ and $\mu \rightarrow eee$ become sensitive to quark operators. In the blind region of $\mu N \rightarrow eN$ in this operator combination, complementary information of the other two processes can be obtained. From this simple scenario of only two contributing operators, it can be deduced that in charged lepton flavour violation it is not a single experiment but the interplay of the searches that yields the best sensitivity. Because of the complementarity of these observables, the best knowledge about the potential New Physics is obtained from the combination of the results of all three searches.

In fact, further observables—from neutrino physics to magnetic and electric dipole moments of leptons, rare decays of taus, mesons, W , Z and Higgs bosons, and direct searches for new heavy particles at colliders—complete the picture. Effective theories can provide a toolkit to compare the results. The approach presented here is however only valid up to the electroweak symmetry breaking scale and does for instance not respect the $SU(2)$ symmetry of the weak interaction. Going beyond the electroweak symmetry breaking scale requires a SM effective field theory.

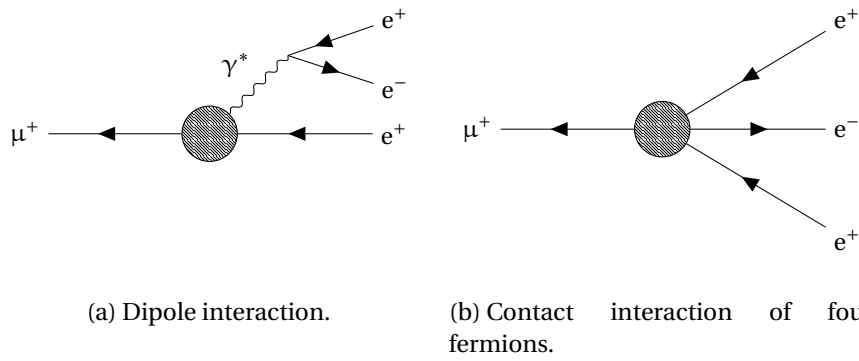


Figure 1.6: The $\mu \rightarrow eee$ signal decay mediated via effective operators of dipole and 4-fermion interactions.

1.2.2 Previous and Current Experiments Searching for Lepton Flavour Violation

As the range of models beyond the SM is huge that predict enhanced rates for lepton flavour violation, searches for lepton flavour violation in the charged lepton sector are a powerful test of New Physics. Lepton flavour violation can manifest in a range of observables: rare muon and tau decays, muon to electron conversion on nuclei, muonium-antimuonium oscillations, and decays of Z, Higgs and heavy neutral particles. A few current experimental limits are given in table 1.2.


Muon to electron transitions yield particularly stringent limits as muons can be produced in large quantities. Experiments searching for $\mu \rightarrow e\gamma$, $\mu \rightarrow eee$ and $\mu N \rightarrow eN$ are briefly discussed in the following. A detailed description of the *Mu3e* experiment is given in chapter 2. The projected sensitivities of future experiments are summarized in table 1.3.

1.2.2.1 The MEG Experiment


The latest results on lepton flavour violation searches of muons stem from the $\mu^+ \rightarrow e^+\gamma$ decay.

As it is a two-body decay, it has a distinct signature. Usually, muon decays at rest are observed. One searches for a positron and a photon with an energy of half the muon rest mass (about 52.83 MeV). Both particles are emitted in a back-to-back topology, have a common vertex, and appear coincidentally.

One source of background is the Standard Model decay $\mu^+ \rightarrow e^+\gamma\bar{\nu}_\mu\nu_e$. As the neutrinos escape detection, a positron and a photon from a common vertex are observed just as in the signal decay.



(a) Dipole and 4-fermion vector interactions.



(b) 4-fermion scalar and vector interactions.


(c) Dipole interactions and scalar interactions involving b fields.

Figure 1.7: Current (solid lines) and future (dashed lines) constraints by muon to electron transition observables on Wilson coefficients from the effective Lagrangian of equation 1.7 at the scale $\mu = m_\mu$. The enclosed area is the allowed parameter space. $\mu \rightarrow e\gamma$ decays are indicated in green, $\mu \rightarrow eee$ in red and $\mu N \rightarrow eN$ conversion in blue. The figures are taken from [61].

Process	Limit	Confidence Level	Experiment
$\mu \rightarrow e\gamma$	$4.2 \cdot 10^{-13}$	90 % CL	MEG [63]
$\mu \rightarrow eee$	$1.0 \cdot 10^{-12}$	90 % CL	SINDRUM [64]
$\mu N \rightarrow eN$ (on Au)	$7.0 \cdot 10^{-13}$	90 % CL	SINDRUM II [65]
$\tau \rightarrow e\gamma$	$3.3 \cdot 10^{-8}$	90 % CL	BaBar [66]
$\tau \rightarrow \mu\gamma$	$4.4 \cdot 10^{-8}$	90 % CL	BaBar [66]
$\tau \rightarrow \mu\mu\mu$	$2.1 \cdot 10^{-8}$	90 % CL	Belle [67]
$Z \rightarrow e\mu$	$7.3 \cdot 10^{-7}$	95 % CL	CMS [68]
$Z \rightarrow e\tau$	$9.8 \cdot 10^{-6}$	95 % CL	OPAL [69]
$Z \rightarrow \mu\tau$	$1.2 \cdot 10^{-5}$	95 % CL	DELPHI [70]
$H \rightarrow e\mu$	$3.6 \cdot 10^{-2}$	95 % CL	CMS [71]
$H \rightarrow e\tau$	$6.1 \cdot 10^{-1}$	95 % CL	CMS [72]
$H \rightarrow \mu\tau$	$2.5 \cdot 10^{-1}$	95 % CL	CMS [72]

Table 1.2: Experimental upper limits on lepton flavour violating observables in the charged lepton sector.

Process	SES	Experiment
$\mu \rightarrow e\gamma$	$6 \cdot 10^{-14}$	MEG II [73]
$\mu \rightarrow eee$	$\leq 1 \cdot 10^{-16}$	Mu3e (phase II) [74]
$\mu N \rightarrow eN$ (on Al)	$2 \cdot 10^{-14}$	DeeMee [75]
$\mu N \rightarrow eN$ (on Al)	$3 \cdot 10^{-17}$	COMET (phase II) [76]
$\mu N \rightarrow eN$ (on Al)	$3 \cdot 10^{-17}$	Mu2e [77]
$\mu N \rightarrow eN$ (on Al)	$\leq 1 \cdot 10^{-18}$	PRISM/PRIME [78]

Table 1.3: Projected single event sensitivity (SES) of future experiments searching for lepton flavour violation with muons.




Figure 1.8: Longitudinal and transverse view of the *MEG* detector [63].

Since the two particles do not fully match the criteria on energy and the back-to-back topology, this background can be suppressed by a good energy and spatial resolution. In addition, there is accidental background, i. e. combinations of positrons, for example from the dominant $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$ decay, with photons from $\mu^+ \rightarrow e^+ \gamma \bar{\nu}_\mu \nu_e$, Bremsstrahlung or positron annihilation. Hence, in addition timing resolution becomes important to suppress accidental background. The most recent result on $\mu^+ \rightarrow e^+ \gamma$ is set by the *MEG* experiment [63, 79] which was operated at the Paul Scherrer Institute (PSI) until 2013. The detector is shown in figure 1.8. A muon beam of about $3 \cdot 10^7 \mu/s$ is stopped on a target in the centre of the experiment. The positrons are tracked in a drift chamber system so that the momentum can be measured by the bending radius in the magnetic field. The time is measured with a timing counter system made of scintillating bars. The photons are measured in a liquid Xenon calorimeter read out by photo-multiplier tubes.

A total of $7.5 \cdot 10^{14}$ muons has been stopped in the *MEG* experiment. The data set is analyzed in a combined blind and maximum likelihood analysis and no significant excess is found. The final result excludes the decay $\mu^+ \rightarrow e^+ \gamma$ down to branching fractions of $\mathcal{B} < 4.2 \cdot 10^{-13}$ at 90% confidence level [63]. It is currently the most stringent bound in charged lepton flavour violating decays.

At the moment, the *MEG* experiment undergoes an upgrade, called *MEG II* [73]. The positron tracker and timing detector are replaced and the liquid Xenon calorimeter is upgraded. The thus improved energy, angular and timing resolution allow for running at higher muon rates. *MEG II* has a projected single event sensitivity of $5 \cdot 10^{-14}$ in 3 years of data taking, about one order in magnitude better than *MEG*.




Figure 1.9: Sketch of the *SINDRUM* experiment [80]. The incoming muons are stopped on the target (T). The multiwire proportional chambers (C) measure the momentum and the hodoscope (M) the timing of the decay products.

1.2.2.2 *The SINDRUM Experiment*

The latest measurement on $\mu^+ \rightarrow e^+e^-e^+$ decays was performed by the *SINDRUM* experiment at the Swiss Institute for Nuclear Research (now PSI).

A drawing of the experiment is shown in figure 1.9. The incoming muon beam is stopped on a hollow double cone target made of a thermo plastic foam. The experiment is placed in a solenoidal magnetic field so that the momentum of the decay electrons and positrons can be measured with an arrangement of multiwire proportional chambers. The timing is measured with a plastic scintillator hodoscope read out with photo-multipliers.

During the run time of *SINDRUM*, no signal decay was observed. An upper limit of $\mathcal{B} < 1.0 \cdot 10^{-12}$ is set at 90 % CL.

The *Mu3e* experiment plans to measure $\mu^+ \rightarrow e^+e^-e^+$ with novel detector technologies and ultimately reach a single event sensitivity of $1 \cdot 10^{-16}$.

1.2.2.3 *Muon to Electron Conversion Experiments*

This type of experiments searches for muon to electron conversion in muonic atoms, e. g. of gold or aluminum. Therefore, a beam of μ^- is used in contrast to μ^+ beams in $\mu^+ \rightarrow e^+\gamma$ and $\mu^+ \rightarrow e^+e^-e^+$ searches. The presence of the nucleus is required to take the recoil in order not to violate energy and momentum conservation. Slow muons are stopped in the target material, captured into the atomic

orbits of the target nuclei, and decay emitting electrons. In the case of conversion, mono-energetic electrons with the energy of the muon mass minus the recoil and binding energy are expected which is about 105 MeV. For all Standard Model decays, the energy of these electrons is below this energy.

Backgrounds are any electrons with an energy of about 105 MeV that reach the detection area. These can stem for example from the muon decay in orbit $\mu^-N \rightarrow e^- \nu_\mu \bar{\nu}_e N$, which has an energy distribution with a tail towards 105 MeV. Another possible background is beam electrons from pion decays in the upstream beam. Prompt background is usually suppressed by using a pulsed beam and waiting for a certain while after a spill before starting the measurement. The latest measurement of $\mu^-N \rightarrow e^-N$ has been performed by the *SINDRUM II* experiment using a gold target [65]. No signal event is observed and thus a limit on the muon conversion rate over the muon capture rate is set

$$\mathcal{R}_{\mu e} = \frac{\Gamma(\mu\text{Au} \rightarrow e\text{Au})}{\Gamma(\mu\text{Au})} < 7 \cdot 10^{-13} \text{ at } 90\% \text{ CL.} \quad (1.15)$$

Currently, various muon conversion experiments are in planning. The *Mu2e* experiment at Fermilab and the *COMET* experiment at J-Parc are briefly mentioned.

Mu2e [77] plans to study up to 10^{18} muons stopped on a aluminum target. The experiment is shown in figure 1.10. A pulsed proton beam with spills every 1700 ns hits a tungsten target producing amongst others a lot of pions. The production solenoid is designed to capture most of these pions. In the transport solenoid, the pions and muons from pion decay are carried to the detector region. Here, charge and momentum of the particles are selected. By the end of the transport solenoid, the beam is a nearly pure muon beam. In the detector region, these muons hit a target of aluminum foils where they stop, form muonic atoms, and decay. The muon stopping rate is about $1.5 \cdot 10^{10} \mu/s$. The decay electrons are measured with a tracker made of straw tubes and a CsI calorimeter. The tracker is dimensioned such that electrons from the dominant Michel decay mode are outside of the acceptance. In order to suppress beam-related background, data taking starts only about 700 ns after a spill. In addition, there are vetoes for cosmic radiation.

Commissioning of *Mu2e* is planned for 2020 with first data in 2021. The single event sensitivity is $\mathcal{R}_{\mu e} < 2.9 \cdot 10^{-17}$.

The *COMET* (COherent Muon to Electron Transition) experiment at J-Parc foresees a staged approach [76] (see figure 1.11). As in the *Mu2e* experiment, muons are generated by hitting a target—




Figure 1.10: Sketch of the *Mu2e* experiment [81].

graphite in this case—with a proton beam. The production solenoid enhances the yield for pion capture, and with the transport solenoid, charge and momentum of the particles are selected. In phase I, the transport section does a 90 degrees turn. Then, the muons hit a aluminum target where they decay. The target is located inside a cylindrical drift chamber measuring the momenta of the decay electrons. In addition, trigger hodoscopes are instrumented.

In phase II, a longer, C-shaped muon transport section will be used. The special shape allows for a better momentum selection. The muons then hit the target and decay. The target and detector section are separated. The additional electron transport system aids in eliminating charged particles with different momenta and charge. Furthermore, there is no direct line of sight between target and detector, hence, the detector is shielded from neutral particles. Also, the detector is upgraded with straw tubes and an electron calorimeter.

The phase I of *COMET* will be commissioned in 2018 with first data taking expected in 2019. A single event sensitivity of $3 \cdot 10^{-15}$ is envisaged. In addition, background measurements for the phase II experiment are performed. The single event sensitivity of phase II is $2.6 \cdot 10^{-17}$.

1.3 FAMILY SYMMETRIES AND THE FAMILON

The flavour structure of the fermions is one of the puzzles of the SM. As the conservation of the individual lepton flavour numbers $L_{e,\mu,\tau}$ is implicit in the SM, it is a tempting possibility to introduce an additional flavour symmetry (see e. g. [83–94]). If this symmetry is spontaneously broken, as indicated by the observation of neutrino oscillations, Nambu-Goldstone bosons arise [88, 95, 96]. These are massless if a continuous global symmetry is spontaneously broken, and acquire mass if the symmetry is also explicitly broken (pseudo-Nambu-Goldstone bosons). In the case of a broken flavour symmetry,

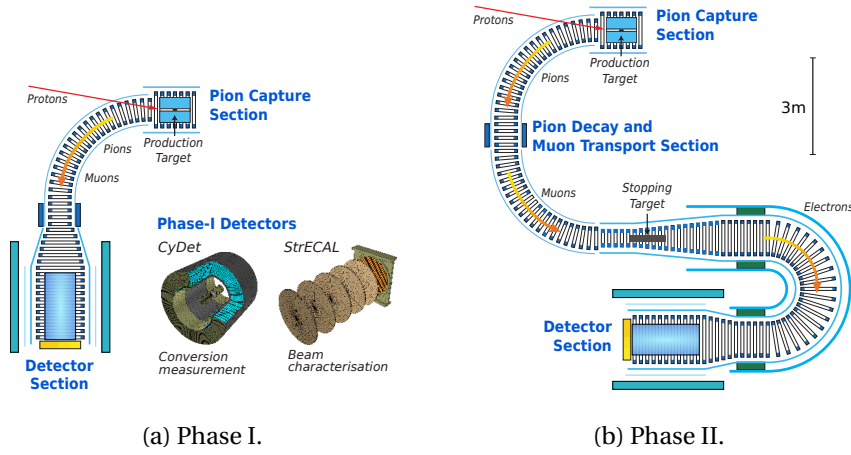


Figure 1.11: Sketch of the *COMET* experiment in the two phases [82].

the Nambu-Goldstone bosons are called *familons*⁶. The mechanism is similar to the spontaneous breaking of a Peccei-Quinn symmetry that gives rise to axions [97, 98]. In many configurations, the familon is also a suitable candidate for cold dark matter [90, 99–101].

Familons are emitted in flavour changing processes like $\mu^+ \rightarrow e^+ X$ decays with the neutral boson X being the familon. These decays are studied in this thesis.

The concept of familons is rather generic. There exist also more complete models that rely on the idea of a spontaneously broken flavour symmetry, for instance models including flavon fields by which the structure of the mass matrices can be explained [102, 103]. Similarly to familons, flavons give rise to $\mu \rightarrow eX$ decays and other flavour changing processes. Also a dark Higgs can be associated with $\mu \rightarrow eX$ decays [104].

The following considerations are exemplary and closely follow [83, 90]. For simplicity, only electrons and muons are addressed although familons can interact with all fermions including quarks. Furthermore, a purely off-diagonal coupling is assumed i. e. muon-electron exchange.

The Standard Model Lagrangian is extended with a complex scalar field ϕ corresponding to the familon. Interactions with fermions are via derivative couplings.

$$\mathcal{L} \ni \frac{1}{2} \partial_\mu \phi \partial^\mu \phi - \frac{1}{2} m_X^2 \phi^2 + c_L \frac{\partial_\mu \phi}{\Lambda} \overline{\mu}_L \gamma_\mu e_L + c_R \frac{\partial_\mu \phi}{\Lambda} \overline{\mu}_R \gamma_\mu e_R \quad (1.16)$$

⁶ A similar concept exists for the difference of baryon and lepton number ($B - L$). The corresponding Nambu-Goldstone boson for a spontaneous breaking of this symmetry is the *majoron*. Majorons can have off-diagonal couplings and mediate $\mu \rightarrow eX$ decays among others.

Herein, μ and e denote the muon and electron fields and $c_{L,R}$ is a coefficient for the interaction with left- and right-handed fermions, respectively. Λ is the mass scale of symmetry breaking and m_X the mass of the familon.

The interaction term can be further simplified

$$c_L \frac{\partial_\mu \phi}{\Lambda} \bar{\mu}_L \gamma_\mu e_L + c_R \frac{\partial_\mu \phi}{\Lambda} \bar{\mu}_R \gamma_\mu e_R \quad (1.17)$$

$$= a \frac{\partial_\mu \phi}{\Lambda} \bar{\mu} \gamma_\mu e + b \frac{\partial_\mu \phi}{\Lambda} \bar{\mu} \gamma_\mu \gamma_5 e \quad (1.18)$$

$$= a \frac{\phi}{\Lambda} (m_\mu - m_e) \bar{\mu} e + b \frac{\phi}{\Lambda} (m_\mu + m_e) \bar{\mu} \gamma_5 e \quad (1.19)$$

using the coefficients

$$a = \frac{c_R + c_L}{2}, \quad b = \frac{c_R - c_L}{2}. \quad (1.20)$$

It becomes apparent that scalar and pseudo-scalar interactions can exist. This can lead to an angular dependence in the differential decay rate of $\mu \rightarrow eX$ decays depending on the size and relative sign of a and b [105]

$$\frac{d\Gamma}{d\cos\theta} \propto (1 + ab \cos\theta). \quad (1.21)$$

The angle θ is spanned by the muon spin direction and the momentum of the electron.

1.3.1 Experimental Limits on $\mu \rightarrow eX$

Muons can be produced in large quantities and the SM decay modes are well understood. Therefore, muon decays are an ideal probe to search for familons and alike. In table 1.4, the current experimental bounds on such muon decays are summarized. For the *Mu3e* experiment, the most relevant process is $\mu^+ \rightarrow e^+ X$. The experiments with the current best limits on $\mu^+ \rightarrow e^+ X$ are presented in the following⁷.

⁷ For searches of $\mu^+ \rightarrow e^+ \gamma X$, it is important to detect the photon as well in order to identify the familon by missing energy. Such a search cannot be performed with *Mu3e*. The limit on $\mu^+ \rightarrow e^+ X$, $X \rightarrow e^- e^+$ stems from $\mu^+ \rightarrow e^+ e^- e^+$ searches at the *SINDRUM* experiment which is already presented in section 1.2.2.2.

Decay	Comment	Limit
$\mu^+ \rightarrow e^+ X$	$m_X = 0 \text{ MeV}$	$2.6 \cdot 10^{-6}$ [106]
$\mu^+ \rightarrow e^+ X$	$13 \text{ MeV} < m_X < 80 \text{ MeV}$	$9 \cdot 10^{-6}$ [5]
$\mu^+ \rightarrow e^+ \gamma X$	$m_X < 2m_e$	$1.3 \cdot 10^{-9}$ [105]
$\mu^+ \rightarrow e^+ X, X \rightarrow e^- e^+$	$\tau_X < 1 \text{ ns}$	$\mathcal{O}(10^{-10})$ [107]

Table 1.4: Current strongest upper limits at 90 % CL on the branching fraction of muon decays involving familons. In the case of $\mu^+ \rightarrow e^+ X, X \rightarrow e^- e^+$, the acceptance and thus upper limit depends strongly on the assumed life time τ_X of the familon.

1.3.1.1 Massless Bosons

In the experiment by Jodidio et al. at *TRIUMF* [106], the endpoint of the positron momentum spectrum from μ^+ was precisely measured in a search for right-handed currents. A sketch of the experimental setup is shown in figure 1.12. A highly polarized μ^+ beam is stopped in the muon polarimeter on a target either consisting of a high purity metal foil (Ag, Al, Au or Cu) or of liquid helium. In these materials, depolarization due to muonium production is strongly suppressed. Within the muon polarimeter, a magnetic field can be applied. With a magnetic field transverse to the beam direction, the muon spins precess leading to a vanishing average polarization. With the spin-holding longitudinal magnetic field, the spins of the muons are not affected. The spin-precessed data is used for calibration, the spin-held data for analysis. The momentum of the decay positrons is measured with a spectrometer based on a dipole magnet and drift chambers.

The dominant Michel decay mode of the muon $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$ has a differential decay rate that strongly depends on the positron energy E_e and the angle θ between the positron momentum and the muon spin direction. At leading order in the SM it reads [46]

$$\frac{d^2\Gamma(\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e)}{dx d\cos\theta} \propto x^2 [(3-2x) + P \cos\theta (2x-1)] \quad (1.22)$$

$$\text{with } x = \frac{2m_\mu E_e}{m_\mu^2 + m_e^2}. \quad (1.23)$$

In this experiment, the muon polarization points upstream, and thus the spin-holding operation mode corresponds to measuring at $\cos\theta = -1$. From SM muon decays, no events are expected at $\cos\theta = -1$ and $x = 1$ (see figure 1.13). In $\mu \rightarrow eX$ decays however, the positron has $x = 1$ because of the massless X , and is emitted isotrop-




Figure 1.12: Sketch of the experiment by Jodidio et al. [106].

ically. Thus, events are also expected at $\cos\theta = -1$ and $x = 1$. A limit on $\mu \rightarrow eX$ decays with $m_X = 0 \text{ MeV}$ is derived from a total of $1.8 \cdot 10^7$ muon decay events in the spin-holding mode. No excess was observed at the endpoint of the positron spectrum. The upper limit at 90 % CL is found to be $2.6 \cdot 10^{-6}$.

1.3.1.2 Massive Bosons

The *TRIUMF Weak Interaction Symmetry Test* experiment (*TWIST*) has collected data from a total of $5.8 \cdot 10^{10}$ μ^+ decays [5]. Highly polarized muons are stopped on an aluminum or gold foil in the centre of the detector (see figure 1.14). The decay positrons are tracked in a solenoidal magnetic field with an arrangement of drift chambers and proportional chambers. Potential $\mu \rightarrow eX$ decays with massive X would show up as an excess on the positron momentum spectrum at a momentum that is characteristic for the mass m_X . In addition, *TWIST* has investigated the angular distribution.

The upper branching fraction limits on $\mu \rightarrow eX$ decays are shown in figure 1.15 [108]. For isotropic decays, the 90 % CL limit is $9 \cdot 10^{-6}$ on average for m_X between 13 MeV and 80 MeV. If the signal positron is preferentially emitted in the direction of the muon spin and thus in the dominant direction of the positron from $\mu \rightarrow e\nu\nu$, the




Figure 1.13: Differential decay rate of Michel decays $\mu^+ \rightarrow e^+\bar{\nu}_\mu\nu_e$ at different angles θ between the positron momentum and the anti-muon spin. The case of $\cos\theta = -1$ resembles the spin-holding operation mode, and $\cos\theta = 0$ the spin-precessing mode.

average limit is slightly higher at $10 \cdot 10^{-6}$. The sensitivity improves if the signal positron is emitted opposite to the muon spin direction. Here, the limit is at $6 \cdot 10^{-6}$ on average. *TWIST* has also set limits on massless X. No improvement over existing results was achieved except for the case in which the positron is emitted in muon spin direction. This limit is $58 \cdot 10^{-6}$ at 90 % CL.

A $\mu \rightarrow eX$ search like in the case of Jodidio et al. cannot be conducted with the *Mu3e* experiment. The very forward region is not instrumented and the muon beam is not as highly polarized.

However, *Mu3e* is capable to perform similar searches as performed by *TWIST*. With a significantly higher expected number of muon decays, *Mu3e* can probe way smaller branching fractions. A sensitivity study for $\mu \rightarrow eX$ decays is presented in chapter 5.

1.4 DARK PHOTONS

The SM of particle physics is remarkably successful in explaining the interactions of the known matter (with the exception of the neutrino oscillations). It fails however to provide a description of dark matter which constitutes about 26 % of the energy-matter density in the uni-




Figure 1.14: Sketch of the TWIST experiment [109].




Figure 1.15: Limits on $\mu \rightarrow eX$ decays (90% CL) as derived by the *TWIST* experiment [108]. Isotropic decays have $A = 0$. For $A = -1$, the positron is preferentially emitted in muon spin direction, and for $A = +1$ in the opposite direction.

verse [17]. Only little is known about dark matter. A dark sector can be defined consisting of particles which interact gravitationally and which are not charged under any of the forces of the SM [110, 111]. The dark sector can include mediator particles that couple to the particles of the SM through *portals*. The portal interactions which are allowed by the symmetries of the SM include the vector, Higgs, neutrino, and axion portals. The following considerations are restricted to the vector portal as it induces processes that can be studied with *Mu3e*.

The mediator of the vector portal is a light, massive vector boson A' from a new, broken $U(1)_D$ gauge group [110–114]. The interaction with SM particles arises through kinetic mixing with the hypercharge gauge boson of the SM $U(1)_Y$ symmetry [115–117]. At masses of A' below GeV, the dominating mixing is with the SM photon. This is why A' is dubbed *dark photon*.

The Lagrangian of this interaction is [118]

$$\mathcal{L}_{A'} = -\frac{\epsilon}{2} F'_{\mu\nu} F^{\mu\nu} - \frac{1}{4} F'_{\mu\nu} F'^{\mu\nu} + \frac{1}{2} m_{A'} A'_\mu A'^\mu. \quad (1.24)$$

Herein, $F^{\mu\nu}$ and $F'^{\mu\nu}$ denote the field strength tensor of the SM photon and of the dark photon, respectively. $m_{A'}$ is the dark photon mass, and ϵ is the kinetic mixing parameter. The first term in equation 1.24 describes the kinetic mixing while the second and third term are the dark photon self-interaction and mass term.

Through mixing, a coupling of the dark photon to the electromagnetic current J_{EM}^μ arises

$$\mathcal{L}_{A'} \supset \epsilon e A'_\mu J_{EM}^\mu, \quad (1.25)$$

which determines the production of dark photons from and decay into SM particles.

In *Mu3e*, the dark photon can be emitted from the muon, positron or intermediate W boson in $\mu \rightarrow e\nu\nu$. If the dark photon mass exceeds $2m_e$, decays into electron-positron pairs are possible. Feynman diagrams of $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays are shown in figure 1.16. In addition to visible decays into SM particles⁸, the dark photon can also have invisible decay modes into particles χ of the dark sector if these are kinematically allowed. This reduces the decay rate for visible dark photon decays, accordingly.

⁸ If the dark photon mass is below $2m_e$, the dark photon can still decay into three photons. As this process is loop-induced, it becomes irrelevant for dark photon masses above the two electron threshold.




Figure 1.16: Dark photons emitted in muon decays. The dark photons decay visibly into an electron-positron pair $\mu \rightarrow e\nu\gamma(A' \rightarrow ee)$.

The dark photon is a possible explanation for the long-standing discrepancy between the theoretical and experimental value of the anomalous magnetic moment of the muon [29, 119–121], although the corresponding parameter space is by now basically completely disfavoured by various experimental constraints. Dark photons or similar particles can also be linked to astrophysical anomalies such as the observed positron excess in cosmic rays [122–128], as well as to the recently reported anomalies in the nuclear transitions of ^8Be [129]. In the latter case, a *protophobic* gauge boson is required with small or zero couplings to the proton in order to be consistent with the dark photon constraints from neutral pion decay [130].

Apart from the universal dark photon that couples to all charged matter fields of the SM through kinetic mixing, a few more light gauge bosons are possible, i. e. for the $B - L$ and $L_i - L_j$ symmetries with $i, j \in \{e, \mu, \tau\}$ [131–134]. The phenomenology is similar to the one of the universal dark photon but of course only couplings to the involved matter fields arise.

1.4.1 Experiments Searching for Dark Photons

Experiments searching for dark photons vary by the production and detection mode and are thus sensitive to different areas in the dark photon parameter space.

Dark photons can be produced via bremsstrahlung of a beam of electrons or protons impinging on a fixed target with atomic number Z $e^-/pZ \rightarrow e^-/pZA'$, or via annihilation $e^+e^- \rightarrow \gamma A'$ in e^+e^- colliders or with an e^+ beam on a fixed target. In hadron colliders, dark photons can be produced in Drell-Yan reactions $q\bar{q} \rightarrow A'$. Furthermore, dark photons can be produced in meson decays, e. g. $\pi^0 \rightarrow \gamma A'$ and $K \rightarrow \pi A'$.

Visible dark photons can be detected in bump hunts in the invariant mass spectrum of the final state and thus usually with high back-

ground. The decay modes are in general $A' \rightarrow l^+ l^-$ and $A' \rightarrow h^+ h^-$, with l denoting leptons and h hadrons. For $e^+ e^-$ colliders and meson decays, the initial state is known, and thus an invisible dark photon can be identified in missing mass distributions.

Decays of visible dark photons can also be searched for at displaced vertices. Again, a bump hunt in the invariant mass distribution is performed. Depending on the dark photon mass and the kinetic mixing parameter, the width and decay length of the dark photon are given by [118]

$$\Gamma_{A' \rightarrow l^+ l^-} = \frac{\alpha \epsilon^2}{3} m_{A'} \sqrt{1 - \frac{4m_l^2}{m_{A'}^2}} \left(1 + \frac{2m_l^2}{m_{A'}^2}\right) \quad (1.26)$$


$$c\tau_{A' \rightarrow l^+ l^-} = \frac{\hbar c}{\Gamma_{A' \rightarrow l^+ l^-}} \approx 0.8 \text{ mm} \left(\frac{10^{-4}}{\epsilon}\right)^2 \frac{10 \text{ MeV}}{m_{A'}}. \quad (1.27)$$

Displaced vertex searches can thus probe very small kinetic mixing parameters.


A dark photon can also be emitted in muon decays $\mu \rightarrow e \nu \nu A'$. The sensitivity of the *Mu3e* experiment to muon decays with visible dark photon decay is evaluated in chapter 6. Due to the similar masses of muons and pions, a comparable dark photon mass region is probed as in measurements of neutral pion decays for example by the *NA48/2* experiment [130].

In figure 1.17, experimental bounds on visible dark photon decays are outlined for existing and future experiments (see [110] and references therein). Constraints can be derived for instance from measurements of the magnetic moment of the electron and muon (a_e, a_μ). In beam dump experiments, typically dark photons with long decay lengths and thus very small mixing parameters and low masses are tested (*E774, E141, E137*). Experiments detecting dark photons from annihilation, bremsstrahlung or meson decay in bump searches or via missing mass span a wider range in dark photon mass but the reach in ϵ is limited by luminosity (*KLOE, NA48/2, BaBar, HADES, PHENIX, A1, APEX, PADME, VEPP-3, DarkLight, HPS, MMAPS, LHCb, Belle-II, BES III*). With searches for displaced vertices additional regions in parameter space at lower ϵ can be explored. These regions are usually separated from the accessible parameter space of bump hunts due to the resolution of the measurement of the decay length.

The entire parameter space in which a dark photon can explain the observed discrepancy between the measured and calculated value



(a) Sketch of the reach of dark photon searches. In *A* are bump hunts and missing mass searches. *B* are displaced vertex searches, and *C* beam dump experiments.



(b) Excluded dark photon parameter space from existing (grey solid areas) and future experiments (open areas). In the green band, the dark photon could explain the anomalous magnetic moment of the muon.

Figure 1.17: Reach in the parameter space of different dark photon searches. Figures taken from [110].

of the anomalous magnetic moment of the muon is disfavoured by experimental observations.

THE MU3E EXPERIMENT

The upcoming *Mu3e* experiment [74, 135, 136] is going to search for the lepton flavour violating decay $\mu^+ \rightarrow e^+ e^- e^+$ with the aim of a final sensitivity of one signal decay in 10^{16} observed muon decays¹, improving on the results of the preceding *SINDRUM* experiment by four orders of magnitude [64].

This poses challenges to the experimental design. In the order of several 10^{16} muon decays need to be observed which requires high muon beam rates and a detector and data acquisition system that are capable of coping with these rates. In addition, an excellent background suppression is indispensable to guarantee background-free operation. This requires a novel detector concept and new detector technologies.

The *Mu3e* experiment is conducted in two stages. In the first phase, the experiment is operated at an existing muon beam line at the Paul Scherrer Institute (PSI). A sensitivity of the order of 10^{-15} is envisaged. The ultimate sensitivity of 10^{-16} requires higher muon beam rates than currently available, and thus a new beam line. This is under investigation at PSI.

In the following, the experimental concept is presented with focus on the phase I detector design.

2.1 SIGNAL AND BACKGROUND

The signature of a $\mu^+ \rightarrow e^+ e^- e^+$ signal decay is defined by two positrons and one electron that emerge coincidentally in time from a common decay vertex. The topology of the signal decay and background processes is sketched in figure 2.1. The invariant mass of the three decay particles equals the muon rest mass

$$m_{eee} \equiv \left(\sum_{i=1}^3 (\mathbf{P}_i)^2 \right)^{\frac{1}{2}} = m_{\mu} = 105.66 \text{ MeV}, \quad (2.1)$$

¹ If not stated differently, the terms *muon* and *electron* denote in the following both, muon and antimuon and electron and positron, respectively.

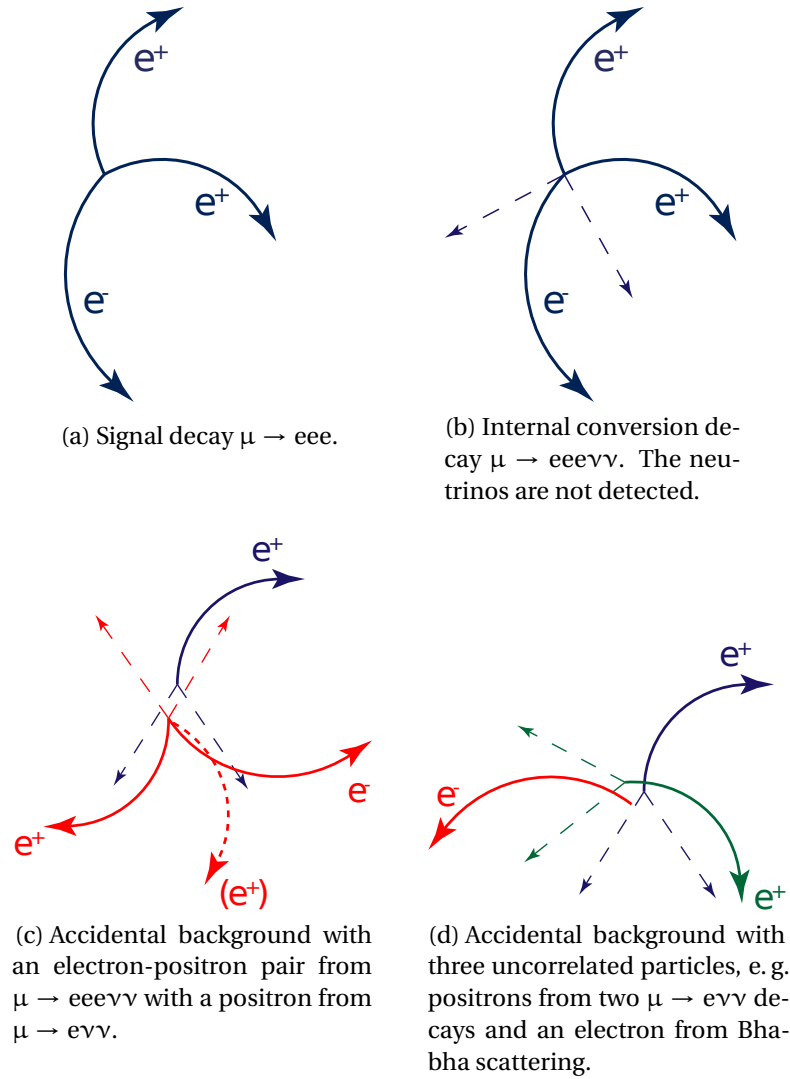


Figure 2.1: Topologies of $\mu \rightarrow eee$ signal decays and background.

with \mathbf{P}_i denoting the four-momenta of the electrons. In the centre of mass system of the muon, the energies of the decay particles sum up to the muon rest mass

$$\sum_{i=1}^3 E_i = m_\mu, \quad (2.2)$$

and the sum of momenta vanishes

$$\sum_{i=1}^3 \vec{p}_i = 0. \quad (2.3)$$

In *Mu3e*, muon decays at rest are observed so that above conditions also hold in the laboratory frame. The maximum momentum of a decay electron is slightly below half of the muon rest mass with

$$|\vec{p}_e| \leq \frac{m_\mu}{2} \sqrt{1 - 10 \frac{m_e^2}{m_\mu^2} + 9 \frac{m_e^4}{m_\mu^4}} \approx 52.8 \text{ MeV}. \quad (2.4)$$

There are two types of background to $\mu \rightarrow eee$ searches. One is background from radiative muon decays in which the photon internally converts to an electron-positron pair $\mu^+ \rightarrow e^+e^-e^+\bar{\nu}_\mu\nu_e$. This decay mode has a branching fraction of $3.4(4) \cdot 10^{-5}$ [1] (see table 1.1 for more SM muon decay modes). As in the case of the signal decay, the two positrons and the electron are coincident in time and share a common vertex. The two neutrinos are not detected which results in missing momentum compared to the signal decay. Hence, equations 2.1 to 2.3 are employed to suppress internal conversion background. The sensitivity of the experiment is thus limited by the momentum resolution. Figure 2.2 shows the integrated branching fraction for cuts on the missing energy. In order to suppress internal conversion background below the projected signal sensitivity, an invariant mass resolution and thus average momentum resolution of better than 1 MeV is required [136]. The second type of background is constituted by accidental combinations of positrons and electrons from multiple muon decays at a time. Plenty of combinations are possible. Positrons typically stem from the dominant Michel decay $\mu^+ \rightarrow e^+\bar{\nu}_\mu\nu_e$ which has a branching fraction close to 100%. These can mimic a signal decay when combined with an electron-positron pair from a Bhabha scattering event of another Michel positron, from photon conversion, or from an internal conversion decay. Combinations of three independent particles are of course also possible. In addition, also deficiencies in the reconstruction can con-




Figure 2.2: Integrated branching fraction of internal conversion decays $\mu \rightarrow eee\nu\nu$ as a function of the missing energy due to the undetected neutrinos. Matrix element courtesy of A. Signer et al. [137]. Original plot from [138].

tribute to the accidental background, i. e. via fake tracks, multiple track reconstructions per particle, and charge misidentification.

This type of background can be suppressed via kinematic constraints as in equations 2.1 to 2.3 and constraints on the vertex and relative timing of the particles. Reducing the amount of material in the active detector volume is another measure as many of the electrons arise from Bhabha scattering, photon conversion and Compton scattering.

The amount of accidental background depends on the rate of observed muon decays. As a rule of thumb, combinations with an electron-positron pair scale linearly with the muon rate and combinations of three uncorrelated tracks with the muon rate squared.

In essence, the kinematics of the signal and the sources of background dictate the design of the experiment and choice of technologies. The experiment must be able to reconstruct low momentum electrons from signal events with an excellent momentum resolution and to provide at the same time very precise timing and vertex information. Furthermore, a large number of muon decays needs to be observed. This requires high muon stopping rates and ideally a large acceptance and reconstruction efficiency in the experiment.

Moreover, all detector components have to be able to cope with high track rates.

2.2 EXPERIMENTAL CONCEPT

The *Mu3e* experiment is a magnetic spectrometer optimized for precise momentum measurements of low momentum electrons from muon decays at rest. The momentum and charge of the decay particles is measured via the curvature of the tracks in the magnetic field. Additional timing detectors provide precise time information. As the muon decays exclusively to electrons and neutrinos, no particle identification system is required.

Due to the low momentum, multiple scattering of the particles in the detector material becomes the dominant source of momentum measurement uncertainty. A low material detector is thus required. Measuring *recurling* particles (see below) further reduces scattering induced uncertainties.

MULTIPLE COULOMB SCATTERING When traversing material, electrons are deflected at nuclei in elastic Coulomb scattering. As a result, they exit the material with a displacement and under an angle with respect to the initial trajectory. For thin layers of material, the displacement is negligible and the scattering angle can be described by a Gaussian distribution with mean zero. The width is given by the Highland formula [139]

$$\sigma_{\text{MS}} = \frac{13.6 \text{ MeV}}{p\beta c} q \sqrt{\frac{x}{X_0}} \left(1 + 0.038 \ln \left(\frac{x}{X_0} \right) \right) \propto \frac{1}{p} \sqrt{\frac{x}{X_0}}. \quad (2.5)$$

Herein, p denotes the momentum and βc the velocity of a particle with charge q in units of electron charge. The path length in the material is x , and the radiation length of the material is X_0 .

As can be easily deduced from equation 2.5, the deflection is stronger for particles with low momenta and for thick detector layers. With momenta of only a few 10 MeV in *Mu3e*, a low amount of material in the active detector volume is thus crucial for precise momentum measurements. In addition to multiple scattering, also the spatial uncertainty of the individual measurement points of the track contributes to the momentum resolution. As illustrated in figure 2.3, the spatial uncertainties can be neglected if they are sufficiently small compared with the uncertainties induced by scattering.




Figure 2.3: Tracking in the case when scattering uncertainties (green) dominate compared to spatial uncertainties (red).

RECURLING PARTICLES In the case of dominating scattering uncertainties, the relative momentum resolution scales with the scattering angle θ_{MS} and the inverse of the track deflection Ω (see figure 2.4) [136]

$$\frac{\sigma_p}{p} \propto \frac{\theta_{MS}}{\Omega}. \quad (2.6)$$

A detector geometry with widely spaced layers is thus preferred for a good momentum resolution. This has to be balanced with requirements on the acceptance as a larger spacing reduces the acceptance for low momentum particles. In addition, smaller spacings facilitate track finding.

As can be seen in figure 2.4b, the trajectory that is distorted by multiple scattering overlaps with the initial undistorted trajectory after about a half turn thus at $\Omega \approx \pi$. At this point, the scattering uncertainties cancel to first order and the best possible momentum resolution is achieved. The geometry of the *Mu3e* detector is therefore optimized to measure at $\Omega \approx \pi$ for most of the tracks.

A sketch of the detector is provided in figure 2.5. The experiment has two detector systems: a tracking detector built from thin silicon pixel sensors and a timing detector consisting of a scintillating fibre and a scintillating tile system. Muons are stopped on a target in the central detector part and eventually decay. The experiment is placed in a solenoidal magnetic field in order to measure the momentum and charge of the decay products via the curvature of the tracks. The inner volume of the magnet is sufficiently large so that the decay electrons do not hit the inner magnet wall. Instead, the particles return back to the detector following their helical trajectory. Such




Figure 2.4: Tracking in a magnetic field with multiple scattering at the detector layers. The particle trajectories are shown in a plane perpendicular to the magnetic field.




Figure 2.5: Schematic of the $Mu3e$ detector in phase I with a $\mu \rightarrow eee$ signal event. The lateral profile (in muon beam direction) is shown.

particles will be referred to as *recurler* from now on. The dimension of the tracking detector is chosen such that between the measurements of the outgoing and recurling particle, the particle has performed about a half turn (see figure 2.6). Thus, the uncertainty of the momentum measurement is minimum. In addition, the acceptance for recurling particles is increased by installing a *recurl station* upstream and downstream of the central detector station.

In the following, the various parts of the phase I detector are described in more detail. The dimensions of the subsystems are summarized in table 2.1.

	Length [cm]	Radius [cm]
Target	10	1.9
Pixel detector	vertex layer	2.3
		12
	outer/recurl layer	34
		36
Fibre detector	29	6.4
Tile detector	36	6.3
In total	112	8.6

Table 2.1: Dimensions of the detector elements in the phase I *Mu3e* detector.




Figure 2.6: Schematic of the *Mu3e* detector in phase I with a $\mu \rightarrow eee$ signal event. The transverse profile (perpendicular to the muon beam direction) is shown. The scintillating fibre detector is not drawn to scale.

2.2.1 Muon Beam

A high muon decay rate is necessary in order to reach the projected sensitivity goal on a reasonable time scale of a few years. In the case of $\mu \rightarrow eee$, it is favourable to operate with continuous muon beams. Thus, the decays are evenly distributed in time which facilitates accidental background suppression compared to operating with pulsed beams.

At the Paul Scherrer Institute, high rate muon beams are available at secondary beam lines. The muons are produced from an intense proton beam impinging on a rotating carbon target. A lot of pions are generated which preferentially decay into muons and neutrinos as the decay to electrons is helicity-suppressed. The muon yield is largest for so-called surface and sub-surface muons. In this case, the pions decay at rest close to the surface of the production target. The momentum of surface beam muons is 29.8 MeV, and slightly less for sub-surface muons.

In the first phase of the experiment, *Mu3e* will operate at the $\pi E5$ beam line which is also used by the *MEG* and *MEG II* experiments. This beam line routinely provides muon rates of $\mathcal{O}(10^8 \mu\text{s/s})$ with muon momenta of 28 MeV. The compact muon beam line (CMBL, see figure 2.7) which is installed in the experimental area of *Mu3e* allows to operate *MEG* and *Mu3e* alternately with minimal and thus time efficient modifications to the $\pi E5$ beam line. During first commissioning runs of the CMBL, rates of $8 \cdot 10^7 \mu\text{/s}$ at the centre of the later *Mu3e* experiment were demonstrated [140]. This allows to observe about $2 \cdot 10^{15}$ muon decays in 300 days of data taking².

As a result of the production from pion decay, the muon beam is polarized with the spins of the μ^+ pointing in the direction opposite to the momentum. Depolarization occurs because of cloud muons from pion decays in flight, transportation losses and beam divergence. The *MEG* experiment measures a muon beam polarization of 86(2) % at the stopping target [142]. A similar polarization is expected for *Mu3e*.

² The following sensitivity studies assume a muon stopping rate of $1 \cdot 10^8 \mu\text{/s}$ and thus a slightly larger total number of observed muon decays. With modifications to the production target and an increase of the initial proton beam current from nominal 2.2 mA to 2.4 mA, increased muon beam rates are expected rendering $1 \cdot 10^8$ muon stops per second feasible [141].




Figure 2.7: Drawing of the CMBL in the experimental area of *Mu3e*.
Figure taken from [136].

2.2.2 Target

The target is a hollow double cone made out of Mylar foil with a thickness of about $70\mu\text{m}$ in the upstream and $80\mu\text{m}$ in the downstream part. In combination with a degrader in the beam line, the material in the target suffices to achieve a high stopping fraction of about 90 % for the low momentum muons. The target is about 10 cm long with a radius of 1.9 cm. The radius is chosen to match the width of the muon beam³. The target is supported by a carbon tube on the downstream side.

The particular shape of the target distributes the muon stopping positions and thus decay vertices over a larger surface. This simplifies the vertex separation and thus accidental background suppression. In addition, it is beneficial if the electrons have to transverse as little target material as possible in order not to be distorted by multiple scattering and to lower the chances for Bhabha scattering and photon conversion. This also motivates to choose a material with a long radiation length X_0 such as Mylar.

2.2.3 Magnet

A solenoidal magnetic field of 1 T is applied. The field lines are aligned with the beam axis. The magnetic field is required to be homogeneous throughout the whole volume that is accessible to the decay particles in order to enable a fast online track reconstruction (see section 2.2.7). The long term stability is specified with a relative deviation of less than 10^{-4} in 100 days of data taking [136].

In order to allow for detailed measurements of the detector acceptance, the magnetic field is adjustable between 0.5 T and 2 T.

³ The muon beam is collimated shortly before the end of the beam line in order to shield the inner pixel detector layers.




Figure 2.8: Geometry of the central pixel detector in lateral and transverse view. Figure adapted from [136].




Figure 2.9: Overlap at two adjacent pixel ladders. The active overlap is about 0.5 mm.

2.2.4 Pixel Tracker

The tracking detector provides precise three dimensional information about the path of the particles in the detector volume. In the case of *Mu3e*, a pixel detector is best suited for this purpose as the high granularity of pixel sensors enables not only precise track measurements but also allows to operate with low occupancies even at high track densities.

The pixel layers are cylindrically arranged around the beam axis. The geometry of the central detector station is shown in figure 2.8. Two so-called vertex layers surround the target. For a good vertex position resolution, it is desirable to have the innermost vertex layer as close as possible to the target. The minimum radius is determined by the target dimensions as well as the beam profile as irradiation of the pixel detector by the initial muons should be avoided. Therefore, also a collimator is installed shortly before the end of the beam pipe. The radii⁴ of the vertex layers are 2.3 cm and 3.0 cm. With a length of 12 cm, the vertex layers slightly extend over the target and thus offer a good coverage.

⁴ The geometry of the pixel layer is a prism. The quoted radii are always to be understood as the minimum radius of the prism.

The outgoing decay electrons are subsequently measured in two outer pixel layers. These four measurement points are sufficient for track reconstruction. Further pixel layers are not added in order to keep the overall material amount low.

The radius of the outer layers is a compromise. Having a large spacing between the vertex and outer layers improves on the one hand the momentum resolution. On the other hand, the radius of the outer layers determines the minimum transverse momentum necessary for a particle to cross the outermost layer and be detected, and therefore the momentum acceptance. Whereas there is no other possibility to restore the acceptance for low momentum particles, the momentum resolution is further improved by measuring recurling tracks in a good portion of events. The radii of the outer layers are thus chosen to be 7.4 cm and 8.6 cm with a length of 34 cm and 36 cm, respectively. This results in a transverse momentum threshold of 10 MeV.

Arranging the pixel layers in pairs yields not only a larger spacing between the vertex and outer layer as preferred in terms of momentum resolution, but aids also the track finding process. Two hits in the closely spaced double layers provide directional information and thus help to reduce the number of hit combinations.

An identical copy of the outer layers is installed in the recurl stations upstream and downstream of the central detector station in order to increase the acceptance of recurling tracks⁵. There are two cases in which a particle is not detected as recurler: either the recurling particle hits the non-instrumented gap between the central and recurl station, or it has a sufficiently large longitudinal momentum to pass over the recurl stations. The latter case is compensated in the phase II upgrade by installation of an additional recurl station upstream and downstream of the phase I detector. The gap is currently estimated to be 2 cm wide for outermost pixel layers, and accordingly 4 cm for the shorter outer layer directly underneath.

The detector layers are composed of *ladders* built from $2 \times 2 \text{ cm}^2$ large pixel sensors. The ladders are arranged in a prism-like geometry with a bit of overlap between two adjacent ladders in order to avoid blind regions (see figure 2.9). In the event that a particle crosses the overlap region, it can thus cause two hits in a pixel layer, one in each of the two adjacent ladders.

⁵ The maximum length of a pixel layer is dictated by mechanical stability and the trace density on the readout High Density Interconnect. Therefore, the detector is segmented into a central and two recurl stations.

In order to reduce multiple scattering of the decay electrons in the detector to a minimum, the pixel detectors need to be very thin, and also the support structure should consist of as less material as possible. A material thickness equivalent to 0.1 % of a radiation length X_0 should not be exceeded per pixel layer including the mechanical and electrical support structure [136].

Standard pixel sensors as typically installed in high energy physics experiments cannot fulfill this requirement. These sensors are usually *hybrid pixel sensors* consisting of a sensitive sensor chip for particle detection and an additional readout chip that is connected via bump bonds to the sensor chip. With sensor thicknesses of 200 μm to 300 μm and bump bonds containing heavy metals such as indium, the material thickness of hybrid sensors is typically a few % of X_0 [143, 144]—too much material for *Mu3e*.

The pixel sensors for the *Mu3e* experiment are therefore built in a novel technology as so-called *High Voltage Monolithic Active Pixel Sensors (HV-MAPS)*. This technology allows for sensors as thin as 50 μm and eliminates the need for an additional readout chip. Details are given below.

The support structure is built from polyimide foils. Electrical connections for readout, slow control and voltage supplies are realized with High Density Interconnects (HDI) consisting of polyimide foil and thin aluminum traces. In the current design, one pixel layer amounts to 0.115 % X_0 which results about equally from the pixel sensor itself and from the mechanical support and HDI combined [136].

The sensors dissipate heat and therefore need to be cooled. This is realized with a gaseous helium flow. As also the scattering in the surrounding gas deteriorates the momentum resolution, helium is chosen because of the low atomic number and thus long X_0 .

2.2.4.1 High Voltage Monolithic Active Pixel Sensors

The pixel sensors are required to be very thin in order minimize scattering. In addition, the sensors need to have a sufficient granularity and a fast signal processing and readout in order to cope with the expected high track rates in the detector. Furthermore, a high hit detection efficiency and a low noise rate are essential for a reliable track reconstruction.

Mu3e uses for this purpose silicon pixel sensors built in *HV-MAPS* technology [145–149]. The detection principle is illustrated in figure 2.10. The pixels are implemented as n-doped wells on a p-doped substrate. At the pn-junction, electrons from the n- and holes from the p-doped side diffuse into the other side of the junction and re-




Figure 2.10: Detection of ionizing particles in a *HV-MAPS* sensor. Figure taken from [145].

combine. A depletion zone without mobile charge carriers is created which forms the sensor diode. Ionizing radiation deposits energy when it traverses the sensor. This results in the excitation of electrons from the valence to the conduction band and thus formation of electron-hole pairs. In silicon, an energy of about 3.62 eV is required on average for the creation an electron-hole pair [150]. The decay electrons in *Mu3e* typically deposit a few keV in the depletion zone resulting in several hundreds of electron-positron pairs.

By applying a reverse bias voltage, the width of the depletion zone and thus the sensitive volume increases. This voltage is called *high voltage* in this context and is typically chosen between (50-90) V. The actual width depends on the doping concentration. For the latest two *HV-MAPS* prototypes for *Mu3e*—MUPIX7 [148] and MUPIX8 [151]—widths of (10-15) μm and (20-30) μm are expected for substrate resistivities of 20 Ωcm and 80 Ωcm at a high voltage of 85 V.

Within the depletion zone, the electric field is strong ($\mathcal{O}(\text{MV/m})$). Hence, the created charge carriers are collected via drift. This is a fast process ($< (1\text{ ns})$) compared to charge collection via diffusion if no high voltage is applied. The timing resolution is thus only given by time constants of the following electronics which range around 1 μs .

As the sensitive volume of the pixel sensors is only a few 10 μm wide, sensors as thin as 50 μm can be produced⁶. This fulfills the requirements of the *Mu3e* experiment in terms of material amount.

As the *HV-MAPS* technology uses a complementary metal-oxide-semiconductor (CMOS) process, it is possible to implement tran-

⁶ In addition to the sensor diode, also the logic inside of the n-wells as well as the metal layers on top of the sensor require some volume. Therefore, the thickness of the pixel sensors needs to be a bit larger than width the depletion zone itself. Within the metal layers, traces are realized which connect for instance the analog pixel electronics with the digital readout cell in the periphery.

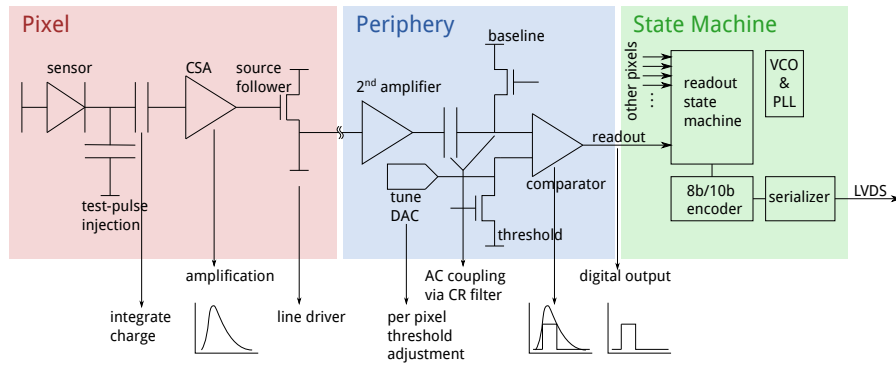


Figure 2.11: Schematics of the electronics implemented on the MUX7 prototype. There is in-pixel charge collection and amplification via a charge sensitive amplifier (CSA), digitization in the readout cells located in the periphery and readout control and data output by an on-chip state machine.




Figure 2.12: Layout of the MUX7 sensor. The active pixel matrix is highlighted in red, the digital periphery in dark blue and the readout control state machine in light blue.

sistors within the pixel n-well and thus build readout electronics directly on the sensor chip. Hence, no additional readout chip is required, a great benefit in terms of material amount.

A schematic of the on-sensor electronics is shown in figure 2.11. The sensor diode collects the charges generated by ionizing particles. Signal amplification and shaping is performed in the n-well itself. Further electronic circuits are hosted in the *periphery*, a small non-sensitive area at the bottom edge of the chip (see figure 2.12). Each pixel has a point-to-point connection to a digital processing unit (DPU) in the periphery. The analogue signal is driven to the DPU. In the case of MUPiX7, a second amplification is performed. MUPiX8 is submitted on a substrate with higher resistivity ($80 \Omega \text{ cm}$ instead of $20 \Omega \text{ cm}$). This results in a larger depletion zone, and thus more charge is collected. The second amplifier is therefore omitted [151].

Then, the signal is digitized in the comparator. If the signal amplitude exceeds a predefined threshold, the comparator output changes and the signal is registered. This is indicated by a flip-flop, the *hit latch*. In addition, the status of an internal counter is stored at the moment the comparator has registered the hit. This is the *time stamp*.

A readout state machine recognizes the hit latches and performs the readout of the sensor. Details are given in section 3.1. In essence, the MUPiX sensors feature digital, zero-suppressed data output at a fast serial link of 1.25 Gbit/s. Thus, the *HV-MAPS* technology also fulfills the high rate requirements.

For the final sensor, an active area of $2 \times 2 \text{ cm}^2$ is foreseen with a pixel size of $80 \times 80 \mu\text{m}^2$. The MUPiX7 is a small size prototype with an active area of $2.9 \times 3.2 \text{ mm}^2$ and a pixel size of $103 \times 80 \mu\text{m}^2$. Apart from the size, this prototype comprises already all the functionalities required for the final pixel chip. It is extensively tested with test beam and yields for instance a hit detection efficiency above 99 % and noise rate of just 2 Hz per pixel [148]. Prototypes of $50 \mu\text{m}$ thickness are tested. The latest prototype MUPiX8 is the first large scale MUPiX sensor with an active area of $2 \times 1 \text{ cm}^2$ and a pixel size of $80 \times 81 \mu\text{m}^2$, and is thus very close in dimension to the final chip. The characterization is currently ongoing.

When operating with a single threshold at the comparator, the timing of the hit detection varies depending on the amount of collected charge and thus signal amplitude. This effect is called *time walk* (see figure 2.13). As the rising edge of signals with a smaller




Figure 2.13: Time walk of signals with different amplitudes. When operating with a single fixed threshold, signals with smaller amplitudes are detected later than signals with larger amplitudes due to the flatter rising edge of the signal.

amplitude is flatter, these signals are detected later compared to signals with a large amplitude. This affects the time resolution of the detector.

In MUPIX8, various measures for time walk correction are implemented. Either, the sensor can be operated with two thresholds, one low threshold to give a precise timing and an additional higher threshold that confirms the detected hit and distinguishes actual hits from noise. Or the sensor provides additional information about the analogue pulse such as the time period in which the signal exceeded the threshold (time-over-threshold, ToT). The latter is a measure of the collected charge and can be used in a later stage of the data acquisition for time walk correction.

2.2.5 Timing Detector

Accidental background can be efficiently suppressed by timing measurements. The experiment therefore employs two timing detector systems.

In the central detector part where a low material amount is crucial, a thin timing detector consisting of scintillating fibres is placed directly underneath the outer pixel layers at a radius of 6.4 cm. The fibre detector is about 29 cm long. In the current design, four layers of scintillating fibres with round cross-section are foreseen [136]. The diameter of a fibre is about $250\mu\text{m}$. The fibres are arranged in ribbons with a silicon photo-multiplier (SiPM) column array at each end for photon detection. It is planned to use the same SiPM as foreseen for the *LHCb* fibre tracker upgrade [152].

The recurl stations are likewise equipped with timing detectors underneath the pixel layers. As no further measurement of the particle is performed once it has reached this timing detector, a large material amount is no longer problematic. Therefore, the recurl stations are equipped with scintillating tiles. The particle deposits significantly more energy in the large volume of a tile compared to a fibre. This results in a larger photon yield and thus better timing resolution. The tile detector is 36.4 cm long with an outer radius of 6.3 cm. It consists of individual tiles with a size of about $6.5 \times 6.5 \times 5.0 \text{ mm}^3$. Each scintillating tile has an individual SiPM.

The SiPMs of both timing detector systems are readout with a custom designed time-to-digital converter chip, the *MuTRiG* [153].

Detector simulation studies with input from prototype characterization studies have demonstrated a significant accidental background suppression by the combined information of the timing detectors [136, 154]. Accidental background with an electron-positron pair and an additional uncorrelated positron is suppressed by a factor of 108. For background of three uncorrelated tracks, the suppression factor is found to be 6400. In combination with further constraints on the kinematics and vertexing, this allows to operate free of accidental background in the first phase of *Mu3e*.

2.2.6 Data Acquisition

The subdetectors of the *Mu3e* experiment continuously stream data to the data acquisition (DAQ) system without the need for a hardware trigger. The DAQ system consists of front-end boards, switching boards, and a filter farm. A schematic of the data flow is shown in figure 2.14.

The complete detector information is collected and forwarded up to the filter farm level. Here, fast online track reconstruction and vertex finding is performed. Only for potential $\mu \rightarrow eee$ signal events, the detector information is stored to disk for a detailed offline reconstruction and analysis. In this way, the data rate can be kept at a reasonable level.

FRONT-END BOARDS Both subdetectors provide zero-suppressed hit information which is forwarded over fast serial links to the front-end boards. The hit information contains a time stamp indicating when the detector element was hit. The position of the detector element is known from the input channel to the front-end board. The pixel detector provides in addition information about the address of the hit pixel.

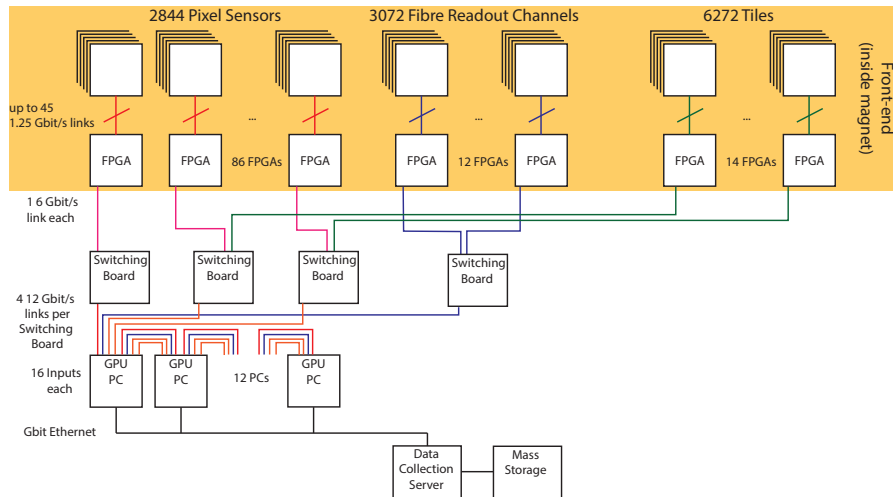


Figure 2.14: The data acquisition chain in the first phase of the *Mu3e* experiment [136].

There is a continuous communication of the pixel detectors with the front-end. Not only hit information is transmitted but also status information and fixed data patterns in the case no hit information is available.

Several detector elements of one type, i. e. either pixels, fibres or tiles, are connected to one front-end board. The front-end boards receive the data and remove unnecessary information. The data is packaged and forwarded to the switching boards over a fast optical link. Data processing is performed on a field programmable gate array (FPGA) on which custom designed firmware is implemented. In the case of the pixel detector, the time structure of the hits is not strictly conserved. It can be restored by using the time stamp information. This is also performed on the front-end FPGA. Details on the front-end firmware for the pixel detector and on hit sorting by time stamp are given in chapter 3.

In the fibre detector, it is expected that a particle induces a signal in several neighbouring fibres and can thus be identified by a coincidence between various channels of a fibre ribbon. Such *clusters* are built on the fibre front-end boards. As clusters are not expected for instance from dark counts, the dark count rate can thus be efficiently reduced.

In addition to data readout, the front-end boards also provide slow control and reset signals to the subdetectors and forward the clock.

The development of the front-end boards is ongoing and prototypes are currently tested in combination with the MUPIX8 sensor prototype.

SWITCHING BOARDS In phase I, four switching boards are foreseen, one for the central pixel detector, one for the fibre detector and one in each recurl station connecting to the recurl station pixel and tile front-end boards. On the switching boards, the data of several front-end boards is merged and partitioned into time slices (*reconstruction frames*) of 50 ns duration. Each of the DAQ computers of the following filter farm thus receives the complete detector information of a reconstruction frame.

For the switching boards, the *PCIe40* board will be used which is currently being developed for upgrades of the *LHCb* and *ALICE* experiment [155]. Again, data processing is performed on an FPGA. The data is transferred via optical links at 10 Gbit/s to the DAQ PCs in the filter farm.

FILTER FARM The filter farm consists of twelve DAQ PCs. Each PC is equipped with an FPGA board that connects via a PCIe (Peripheral Component Interconnect Express) [156] interface to the memory of the PC. The *PCIe FPGA* of the first DAQ PC receives the data from the switching boards. If busy, the data packets are forwarded in a daisy chain to the next DAQ PC. On the PCIe FPGA board, the packets from the four switching boards are merged. Hit information from the central detector is processed further, while the other data is buffered. A pre-selection is performed on the central pixel detector hits in order to identify hits that potentially stem from a common particle track. The coordinates of the data are transformed to the global coordinate system (instead of addresses of the individual chips and pixel addresses therein). This data is then written to the main memory of the PCs via direct memory access (DMA) and subsequently copied to the memory of a graphics processing unit (GPU). On the GPU, fast track and vertex finding is performed in parallel for all possible hit combinations [157].

If a $\mu \rightarrow eee$ signal candidate is identified, the complete data of the reconstruction frame is stored to disk for a detailed offline reconstruction and analysis. Therefore, data from the fibre detector and recurl stations is buffered on the PCIe board. Recent studies have demonstrated a data rate reduction from about 80 Gbit/s at the input to the filter farm to below 100 MB/s at the output to the data collection server [157].

2.2.7 *Track and Vertex Reconstruction*

The pixel layers represent the densest areas in the active detector volume and thus cause most of the scattering of the particles. Hence,




Figure 2.15: Triplet fit with three hits by varying the scattering angles at the central hit.




Figure 2.16: Joining two triplets to one track.

scattering at the pixel layers needs to be accounted for in the track and vertex reconstruction.

The track reconstruction relies on the fast three-dimensional *triplet fit* [158–162]. A homogeneous magnetic field is assumed and the energy loss in the pixel layers is neglected during the reconstruction. An energy loss correction is applied afterwards taking the traversed path length in the material into account.

In a first step, three hits from consecutive layers are joined to a *triplet*. The three-dimensional helix radius R is fitted by minimizing the scattering angles at the central hit. This is done with the method of least-squares by minimizing

$$\chi^2(R) = \frac{\Phi_{\text{MS}}^2(R)}{\sigma_\phi^2} + \frac{\Theta_{\text{MS}}^2(R)}{\sigma_\theta^2}. \quad (2.7)$$

Herein, the azimuthal scattering angle is denoted as Φ_{MS} , and the polar scattering angle as Θ_{MS} (see also figure 2.15). The variances are given by

$$\sigma_\phi^2 = \sigma_{\text{MS}}^2, \quad (2.8)$$

$$\sigma_\theta^2 = \frac{\sigma_{\text{MS}}^2}{\sin^2 \theta}, \quad (2.9)$$

in which θ is the polar angle and σ_{MS} the scattering uncertainty as given in equation 2.5.

In a next step, the individual triplets are joined to a track such that the two triplets share two hits (see figure 2.16). As the scattering in each layer is independent of the scattering in the other layers, this can be performed by minimizing the sum of the χ_i^2 of the individual triplets

$$\chi_{\text{global}}^2 = \sum_{n_{\text{triplets}}} \chi_i^2. \quad (2.10)$$

It can be shown that this is equivalent to taking the weighted average of the radii R_i of the triplets [162]

$$\bar{R} = \sum_{n_{\text{triplets}}} \frac{R_i}{\sigma_{R,i}^2}, \quad (2.11)$$

where $\sigma_{R,i}$ denotes the uncertainty estimate on R_i .

In the successive vertex fit, two positron track candidates and one electron track candidate are used [163–165]. The vertex position \vec{x}_v is fitted by minimizing the scattering angles of each track at the innermost vertex layer

$$\chi^2(\vec{x}_v) = \sum_{n_{\text{tracks}}} \left(\frac{\Phi_{\text{MS},i}^2(\vec{x}_v)}{\sigma_{\phi,i}^2} + \frac{\Theta_{\text{MS},i}^2(\vec{x}_v)}{\sigma_{\theta,i}^2} \right), \quad (n_{\text{tracks}} = 3). \quad (2.12)$$

In *Mu3e*, there are two types of reconstructed tracks: short tracks from four hits (two triplets) and long tracks from six or eight hits (four or six triplets), as shown in figure 2.17. Short tracks are reconstructed for outgoing particles using one hit from each layer of the central pixel detector. If the recurling particle hits the outer or recurl pixel layers again, the short tracks can be extended to long 6 hit-tracks. Particles with mostly transverse momentum can also cross the vertex layers again. In this case, a long 8-hit track can be reconstructed.

Figure 2.18 shows the acceptance and momentum resolution for the short and long track reconstruction. The acceptance for short tracks is limited by the dimension of the central detector station. Particles need about 10 MeV of transverse momentum to cross also the outer layers. Also particles that are emitted under a small angle compared to the beam axis are not detected.

For long tracks, additional acceptance losses arise from the gap between the central and recurl stations as well as from the total length of the detector. Particles can have sufficient longitudinal momentum to pass over the recurl stations and can therefore not be recon-

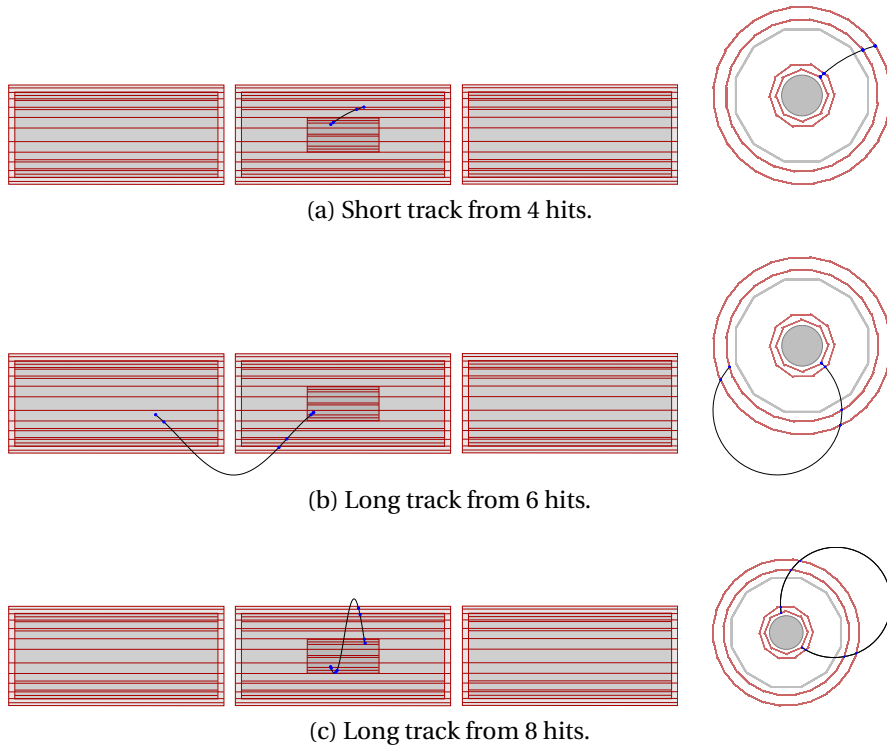


Figure 2.17: Event displays of reconstructed short and long tracks.

structed as long tracks.


In the case of 8-hit tracks, it is a priori not clear with which sense of rotation the track should be reconstructed, and the track reconstruction can thus mis-identify the charge of the particle. This ambiguity can be resolved with information from the timing detectors.

As discussed above, multiple scattering uncertainties cancel to first order if the particle has performed about a half turn in between the measurements. This is approximately fulfilled for long tracks. Therefore, a significantly improved momentum resolution is observed.


Short tracks yield momentum resolutions between 0.5 MeV and 3 MeV. For long tracks, the momentum resolution is improved with about 0.1 MeV to 0.45 MeV. The best momentum resolution is achieved around momenta of 20 MeV as for this momentum the condition of the half turn is fulfilled.

The performance of the vertex fit is discussed in the context of sensitivity studies for $\mu \rightarrow eee$ signal decays in chapter 4.


In the online reconstruction, only the information of the central pixel layers is processed and therefore only short tracks are reconstructed at this stage. This has an impact on searches for two-body muon decays $\mu \rightarrow eX$ which are studied in chapter 5. In principle, also long track reconstruction can be performed online. However,




(a) Efficiency for short tracks.



(b) Momentum resolution of short tracks.



(c) Efficiency for long tracks.



(d) Momentum resolution of long tracks.

Figure 2.18: Efficiency and momentum resolution for reconstructed short and long tracks. Figures taken from [136].

track finding is more complicated in this case. The search window for the additional two to four hits has to be chosen rather large due to the long flight path. This results in increased combinatorics which cannot be processed on the currently specified filter farm. With the constant improvement of the performance of commercially available GPUs, online reconstruction of long tracks might come into reach in a few years.

2.3 SIMULATION OF THE MU3E EXPERIMENT

A detailed simulation of the *Mu3e* detector is realized with the GEANT4 simulation toolkit [166]. The simulation is used to study the acceptance for $\mu \rightarrow eee$ signal decays and to optimize the detector geometry. The further *Mu3e* software framework comprises also algorithms for track and vertex reconstruction and thus enables efficiency and measurement resolution studies as well as sensitivity studies of $\mu \rightarrow eee$ decays and other processes. Up to now, only the pixel tracker information is used in the subsequent analysis.

The geometry of the phase I detector is implemented as described above with the exception that the fibre detector is currently modelled as three layers of square fibres instead of four layers of round fibres. The material amount is about the same for the two fibre geometries. The muon beam is generated 1 m in front of the target in the beam pipe and is modelled using measured beam parameters at this spot.

In GEANT4, the generated particles are propagated through the experimental volume and the detector response is simulated. In addition, interactions with matter such as energy loss due to ionization, bremsstrahlung, photon conversion, Bhabha scattering and multiple Coulomb scattering are modelled.

The simulation can be operated in two modes. In the normal operation mode, a muon beam with a chosen particle rate is generated. The muons decay according to their lifetime of about $2.2\ \mu\text{s}$. The decay modes and their relative occurrence are specified in a decay table (see below for the various muon decays). This is typically used to simulate SM muon decays.

In the special decay mode, only a selected muon decay or overlaps of specific processes are generated. One event of the specified process is generated per reconstruction frame of 50 ns. SM muon decays with a chosen muon beam rate can be simulated in addition. This operation mode is used to simulate specific signal decays as well as rare types of background.

2.3.1 *Generators of Physics Processes*

The simulated SM muon decay modes are the dominant Michel decay $\mu \rightarrow e\nu\nu$, the radiative muon decay $\mu \rightarrow e\gamma\nu\nu$, and the radiative muon decay with internal conversion $\mu \rightarrow eee\nu\nu$.

In a first step, four-momenta of the decay products are (pseudo-) randomly⁷ generated in the centre of mass system of the muon. The four-momenta are evenly distributed in the accessible phase space. Then, the accept-reject method [167, 168] is applied in order to distribute the events according to the respective matrix element. In a final step, the decay particles are boosted to the laboratory frame. Michel decays are simulated as implemented in GEANT4. The implementation is based on the matrix element calculations by Scheck and Fischer including radiative corrections [169, 170] and takes the polarization of the muons and electrons into account.

The phase space generator for the radiative decay mode is adopted from the GEANT4 implementation of the decay that stems from the *TWIST* collaboration [171]. Within the scope of this thesis, a bug in the phase space generator is corrected. By the means of the transformation method [167, 168], the photon energy E_γ is generated with a distribution proportional to $\frac{1}{E_\gamma}$ which is an envelope of the differential decay rate. An event distribution according to the matrix element is generated with the accept-reject method as usual. Because of the preceding transformation of the generated phase space, the number of trials during accept-reject is significantly reduced which results in a faster generation of larger samples of $\mu \rightarrow e\gamma\nu\nu$ decays. The matrix element is taken from Kuno et al [46]. As it is divergent for $E_\gamma \rightarrow 0$, a low photon energy cutoff needs to be introduced. Below the cutoff, the decay is treated as a normal Michel decay. The branching fraction with cutoff is derived by means of Monte Carlo integration [167].

The phase space for the radiative decay with internal conversion is generated with RAMBO (RANDOM MOMENTA BOOSTER) [172]. The generated events are then either distributed according to the matrix element by applying the accept-reject method, or they are weighted with the corresponding value of the matrix element. The latter is often used when large background samples are needed. The calculation of the matrix element stems from Signer et al. [137] and takes the polarization of the muon into account. It is further possible to generate events for an restricted phase space for instance for invariant masses of the three electrons above a certain value. As

⁷ Using a pseudo-random number generator allows to reproduce results by choosing the same seed.




Figure 2.19: Schematic of the *Mu3e* detector in phase II with a $\mu \rightarrow eee$ signal event. The longitudinal profile is shown.

in the case of the radiative decay, the branching fraction is derived by means of Monte Carlo integration. This is extensively used for background studies for $\mu \rightarrow eee$ signal decays.

The implementations of $\mu \rightarrow eee$ signal decays and other New Physics processes are presented in the corresponding chapters in the following.

2.4 UPGRADES TO THE PHASE I EXPERIMENT

2.4.1 Phase II of the *Mu3e* Experiment

For the ultimate sensitivity of one signal decay in 10^{16} muon decays, a muon beam rate of about $2 \cdot 10^9 \mu/s$ is needed which is up to now not available. The PSI is currently investigating on a high intensity muon beam line (HiMB) [173]. With changes to the muon production target, the capture solenoid and the beam line, muon beam rates of about $10^{10} \mu/s$ can be expected.

Correspondingly, the detector is upgraded (see figure 2.19). In the second phase of *Mu3e*, an additional recurl station is installed upstream and downstream of the phase I detector in order to increase the acceptance for recurling particles. The muon beam is collimated to a smaller cross-section. Thus, a target with a smaller radius can be installed. Accordingly, the vertex pixel detector is replaced with a version with smaller radius. In this way, the vertex position resolution is improved compared to the phase I detector. A further option is to reduce the size of the pixels in the vertex layers to account for the increased track density. The fibre detector cannot be operated at the expected track rates and thus needs to be replaced. One option under consideration are *Low Gain Avalanche Detectors (LGAD)* [174–176], a novel type of fast silicon semiconductor detectors with an expected timing resolution below 100 ps.




Figure 2.20: The potential *Mu3e-Gamma* upgrade [177]. The new detector station consists of two pixel layers, a photon conversion layer and/or a timing detector.

2.4.2 The *Mu3e-Gamma* Experiment

The *Mu3e* detector can also be upgraded for $\mu^+ \rightarrow e^+\gamma$ searches [177, 178]. A possible realization of the *Mu3e-Gamma* experiment is shown in figure 2.20. The positron is detected as usual in the *Mu3e* detector. The photon is reconstructed from the electron-positron pair from a conversion. For this purpose, an additional detector station is installed at a larger radius which comprises two pixel layers and a photon conversion layer. The converter can also be replaced by or combined with a timing detector. The experimental concept foresees to increase the magnetic field strength to 2 T so that the initial positron as well as electrons and positrons from background processes do not have sufficient transverse momentum to reach the outermost pixel detector.

Part II

DATA ACQUISITION

The advance in experimental particle physics goes along with a continuous improvement in measurement precision and ever increasing particle rates. This poses challenges to the data acquisition as large amounts of data need to be processed at high rates.

This applies also to the *Mu3e* experiment. With a highly granular tracking and timing detector operated at unprecedented muon stopping rates, a sizeable data rate arises. An elaborate data acquisition system is thus key to the success of the experiment.

The data acquisition in the *Mu3e* experiment operates without a hardware trigger. All detector components are read out and the complete data is streamed to the filter farm where $\mu \rightarrow eee$ signal candidates are identified and written to data storage for analysis. The data acquisition must thus be designed such that it operates efficiently even at high hit rates.

In the case of the pixel detector, data acquisition starts already on-chip as the sensors have in-built readout circuitries. The following front-end boards collect the hit data and sort it by time before the information of the pixel detector is merged with data from the other sub-detectors.

In the following chapter, an implementation of the hit sorting on the front-end is elaborated and evaluated in a telescope test setup consisting of the latest pixel sensor prototypes.

3.1 READOUT ON THE PIXEL SENSOR

The MUPIX sensors integrate particle detection and readout functionalities on a single device. The following considerations focus on the current prototypes MUPIX7 [148] and MUPIX8 [151].

Each pixel in the active area is connected to a readout cell in the periphery at the bottom edge of the sensor. Whenever the analogue signal from a pixel exceeds the threshold of the comparator in the corresponding readout cell, a flip-flop (hit latch) is set. In addition, the status of the on-chip Gray counter¹ is stored as time stamp indicating the time of arrival of the analogue pulse in the periphery and thus basically the time of the hit itself. The hit latch stays set until readout. Within this time period, no other hits can be detected in the corresponding pixel. However, the probability to hit the pixel twice within the typical readout time of $1\ \mu\text{s}$ is very low due to the high granularity and low occupancy of the sensor.

The readout of the hits is controlled by an on-chip state machine. At the beginning of a readout cycle, the hit latch is registered. This

¹ Gray encoding [179] is an alternative binary encoding such that the bit sequence of two successive natural numbers differs only in one bit. The bit sequence is never in an unintended state during the transition and thus the risk of failure is reduced.

is recognized by the pixel priority logic. The row address and time stamp of the first hit in the priority chain of each column are transferred to the end-of-column (EoC) logic.

If there is more than one hit in a column, only the first hit is copied, i. e. the hit with the lowest row address², and the other hits are read out in the following readout cycles (see figure 3.1). In this way, the time structure of the hits is distorted. Hits that happened at the same time do not necessarily end up in the same readout cycle. As the time information is contained in the time stamp, the time order can be restored at a later point in the data acquisition chain. This is performed on the front-end boards.

The priority logic in the EoC then transfers the column and row address and the time stamp of each hit one after the other to the serializer. In principle, the hits in the EoC logic are also prioritized. As the maximum number of hits per readout cycle is usually set to the number of columns, all hits in the EoC can be read out in one cycle and the time order of the hits is not further distorted.

In the serializer, the data is 8 bit/10 bit encoded³ and serialized. The serialized data is sent off-chip on a fast 1.25 Gbit/s low-voltage differential signaling (LVDS) link⁴.

The data output protocols of MUPix7 and MUPix8 are given in tables A.1 and A.2 in the appendix. The data output contains hit information, i. e. pixel address and time stamp, as well as time informa-

-
- 2 This is only strictly true for MUPix7. In the case of MUPix8, the mapping of the active pixels to readout cells in the periphery is a bit more complicated. Nevertheless, the problem that only one hit per column can be read in one cycle stays the same.
- 3 In 8 bit/10 bit encoding [180–182], an 8 bit data word is mapped on a 10 bit symbol. The two additional bits allow for some boundary conditions and thus for error checks of the data transfer. For instance, the disparity, i. e. the difference of zeroes and ones in one symbol, is bound and is either 0 or ± 2 . The running disparity, i. e. the disparity of the total transferred data after the transmission of each symbol, is in addition required to be ± 1 . In this way, the encoded data stream is direct current (DC) balanced meaning that also over longer data transfer periods the number of transmitted zeroes and ones are equal. Furthermore, bit sequences of more than five ones or zeroes are prohibited. This renders also clock recovery from the data stream possible. In addition, there are twelve control symbols called comma words that can be realized but do not correspond to an 8 bit data word. Comma words are used as synchronization pattern and to indicate different types of information in the MUPix data output.
- 4 In differential signaling, the electrical signal is transmitted as a differential pair of signals instead of a single signal in single-ended signaling. The voltage difference of the pair determines whether a zero or a one is transmitted. Differential signaling is necessary at fast serial links as it is more resistant against noise and cross talk. The low voltage standard of LVDS allows for a low power consumption. This is important as the signals are generated on the sensor chip.




Figure 3.1: Schematic of the readout sequence in the digital periphery of a MUPiX sensor.

tion, i. e. a 24 bit counter operating at the same frequency as the time stamps of the hits. This counter is called *binary counter* in the following. Just as the time stamp, the binary counter represents the time since the last reset. The hit and time information is interleaved with comma words from the 8 bit/10 bit encoding scheme which are used to flag the different types of information. The information about one hit or one binary counter is split over four clock cycles. The overhead, i. e. time information and comma words, is expected to be similar for the final sensor so that the bandwidth is about the same as for the current prototypes. Additional information, e. g. about the temperature of the sensor, can be transferred instead of comma words. MUPiX7 has a single data output link. In the case of MUPiX8, the pixel matrix is divided into three sub-matrices. There are four LVDS links per sensor, one link per sub-matrix and an additional fourth link that either transmits a copy of the data output of one of the other links or a merged output of all the three other links. An identifier for the sub-matrix is included in the header of the data output. A similar scheme will be implemented on the final sensors. For sensors of the vertex layers, the occupancy is higher than for the remaining sensors. Therefore, the three individual links are used for

data output. In the outer layers and recurl stations, the occupancy is lower. As there are eight to nine sensors in a ladder that are readout at one end, the density of traces on the readout flexprint becomes the limiting factor. Therefore, only the merged link is used for data output of these layers.

Both sensor prototypes feature a synchronous reset. This signal resets the state of the Gray and binary counter and is used to synchronize several MUPIX sensors.

An important difference in the on-chip readout of both prototypes is the length of the column. MUPIX8 is a large area prototype and has a column length of 200 pixels. The previous MUPIX7 is smaller. The physical pixel column is 40 pixels long and it is mapped on two readout columns so that a readout column is 20 pixel long—one order of magnitude less than for MUPIX8.

3.2 DATA ACQUISITION AT THE FRONT-END OF THE MU3E PIXEL DETECTOR

The data path in the front-end of the phase I pixel detector is shown in figure 3.2. The front-end boards collect and process the data of the pixel sensors. This is performed on a powerful ARRIA V 5AGX BA7D4F31C5N FPGA from ALTERA [183]. In addition, the clock, the fast reset for synchronization and the slow control data are forwarded to the pixel sensors.

The assignment of the sensors to the front-end boards is determined by the geometry of the detector. The sensor ladders of one layer are split into an upstream and downstream half. Sensor ladders that are geometrically close are connected to one front-end board so that for example ladders from the two vertex layers or from the two outer layers are connected to one board. In the case of the vertex layers, either twelve or fifteen sensors are connected to one front-end board. With three links per sensor, this amounts to 36 or 45 links per board. In the case of the outer layers and recurl stations, either 32 or 36 sensors are connected to one board. Here, the data of one sensor is transmitted on a single merged link.

The received data is deserialized and 8 bit/10 bit decoded. As it takes four clock cycles to transmit the data of one hit or one binary counter, the data stream of four sensors can be easily multiplexed to one. The output is then one full data word or binary counter per clock cycle. An additional multiplexing stage can be implemented that merges two data streams into one data stream at twice the frequency. The number of data streams is reduced and thus also the logic usage on the FPGA.

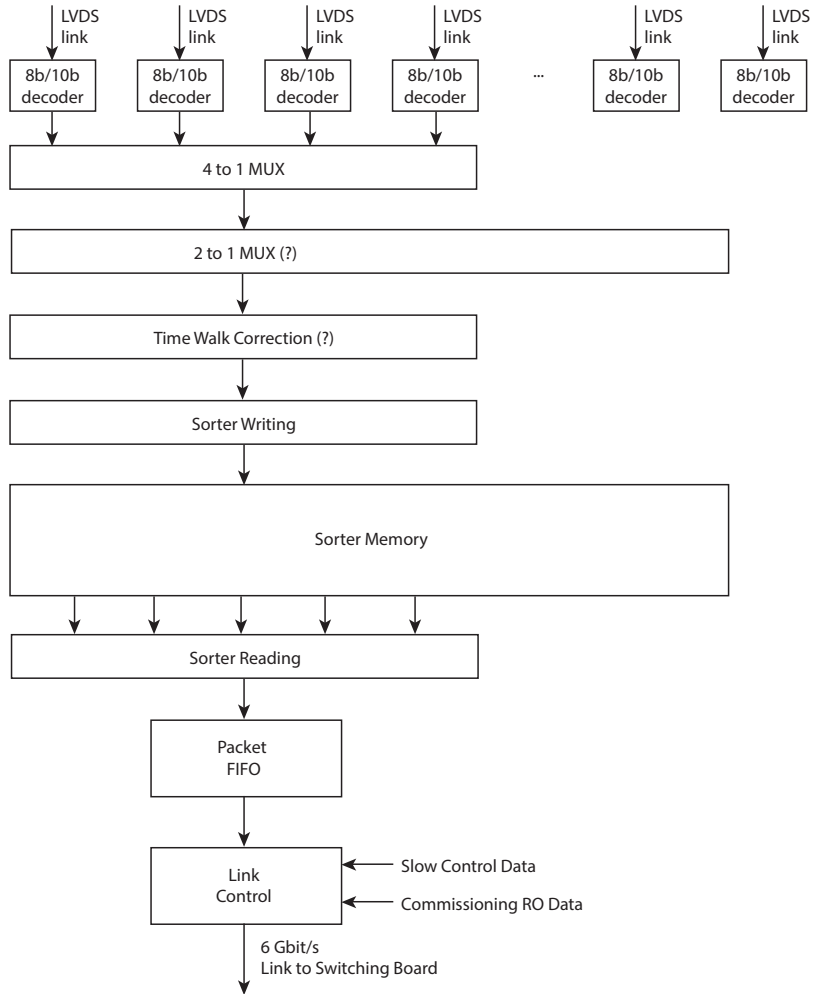


Figure 3.2: Current outline of the data path in the pixel detector front-end in phase I [136].

Layer	Hit rate [MHz]	Hit data rate [Mbit/s]
Vertex, inner	4.1	130
Vertex, outer	3.3	110
Outer, inner	0.67	21
Outer, outer	0.50	16
Recurl, US, inner	0.11	3.4
Recurl, US, outer	0.091	2.9
Recurl, DS, inner	0.071	2.3
Recurl, DS, outer	0.060	1.9

Table 3.1: Expected average hit and data rates per sensor in the various pixel layers of the phase I *Mu3e* detector. A muon stopping rate of $10^8 \mu/s$ is assumed. For the hit data rate, a 32 bit data word per hit is assumed and the header information is neglected. US denotes upstream and DS downstream.

The MUPIX8 sensor features information about the analogue pulse of each hit which can be used to correct the time stamps for time walk. The necessity of this feature has to be evaluated with current prototypes and the final sensor. If a time walk correction is needed, it has to be performed on the front-end board before the following time sorting step.

Sorting of the hits by time stamp is important for the later online reconstruction. The online reconstruction on the filter farm is performed on the hit data from the complete central detector station of a select time slice. If some hits would arrive too late at the filter farm, they are missing for the track finding and the online reconstruction would work less efficiently. This in turn reduces the efficiency to find a $\mu \rightarrow eee$ signal event. Therefore, the time order of the hits is restored on the front-end boards before data streams of different subdetectors are combined.

After sorting, the data is packaged and fed into a FIFO (first in - first out) buffer from where it is accessed by the output link controller. The data is transmitted to the switching boards via an optical link at 6 Gbit/s.

3.2.1 Data Rates at the Front-End

The bandwidth requirements of the front-end board and its components are evaluated with the GEANT4 based detector simulation

of the *Mu3e* experiment at phase I. In table 3.1, the expected average hit rates per sensor are listed assuming a muon stopping rate of about $10^8 \mu/s$. As expected, the sensors in the vertex layers are the busiest. The average data rate even of the inner vertex layer sensors stays well below the $3 \cdot 1.25 \text{ Gbit/s}$ of the three LVDS links. Fluctuations of the hit rate on the sensor should thus be easily buffered.

The rate capability of the hit sorter is evaluated in the following sections. All other components basically stream the data and do not pose a bottleneck. The maximum rate is only limited by the optical link to the switching board. With a data transfer rate of 6 Gbit/s , the optical link is generously designed, as even fifteen inner vertex layer sensors produce an average hit data rate of only 2 Gbit/s .


Besides the data rate, also the time structure of the incoming hits needs to be considered in the firmware design, more precisely the time differences between the incoming hits and the number of hits with the same time stamp.

For the latter, an estimate is derived from simulation and presented in the following. The former however can currently not be derived from simulation because the readout sequence on the sensors is not implemented and because a modelling of noise and other sensor related effects is missing. Instead, the implementation of the hit sorter is tested in test beam measurements with actual sensor prototypes. The results are given in section 3.4.4 and 3.5.4.


As the sorting is performed on a memory of fixed size, the depth of the sorting, i. e. the maximum number of hits with the same time stamp that can be processed, is limited.

How many hits occur in the interval of one time stamp depends on the muon beam rate and on the chosen time stamp size. The phase I detector is expected to operate at a muon stopping rate of $10^8 \mu/s$. The current developments of the data acquisition foresee a time stamp size of 8 ns . Figure 3.3 shows on the left-hand side the distribution of the number of hits per time stamp in the various layers. On the right-hand side, the remaining efficiency is shown for a limited number of hits with the same time stamp that can be processed. Time stamp sizes of 4 ns , 8 ns and 16 ns are chosen.


The number of hits per time stamp is derived as follows. In each layer, the number of hits of the busiest sensor is determined. A front-end board connects always to two layers, for example to the two vertex layers. If hits are caused by an actual particle track, the hit sensors in two neighbouring layers are geometrically close so that the chance is high that both sensors are connected to the same




(a) Number of hits per 4 ns time stamp.




(b) Efficiency for limits on the maximum number of hits per 4 ns time stamp.




(c) Number of hits per 8 ns time stamp.



(d) Efficiency for limits on the maximum number of hits per 8 ns time stamp.



(e) Number of hits per 16 ns time stamp.



(f) Efficiency for limits on the maximum number of hits per 16 ns time stamp.

Figure 3.3: Number of hits per time stamp and efficiency for limits on the maximum number of hits per time stamp. For the vertex layer, the number of hits of the busiest sensor in the inner and in the outer vertex layer are summed up. The same applies to the other combinations of layers. The upstream recurl layers are labelled US, and the downstream recurl layers DS. 10^6 frames of a length equal to the time stamp size are generated for a muon stopping rate of $10^8 \mu/s$ in the phase I detector.

front-end board. Therefore, the number of hits of the busiest sensor in the two vertex layers is added. The same is done for the outer layers and the upstream and downstream recurl layers. This estimate is thus rather conservative as it ignores the fact that the sensor in the other layer could also connect to another front-end board. This way of estimating the number of hits is chosen to compensate in part the fact that the simulation does not include hits caused by noise, clustering (i. e. a particle causes a signal in more than one pixel) and cross talk (a signal on one trace on the sensor induces a signal on another trace). These contributions are expected to be very small, but a further investigation and input from sensor characterization studies is needed to give an actual estimate.

The distribution of the number of hits in figure 3.3 is largest at small numbers of hits. At a rate of $10^8 \mu/s$, a muon decay happens about every 10 ns but due to the geometric acceptance of the detector, many of the decay events are not detected. Therefore, all of the distributions have many entries at zero hits. If a particle crosses two layers, two hits are recognized in the front-end board. Therefore, there is a maximum at two hits. In rare cases, the particle crosses the overlap region between two adjacent pixel modules of one layer so that two hits are detected in the layer. For this reason, the hit sorter should be able to process at least four hits per time stamp in order not to reject single particle events. The tail towards larger numbers of hits is caused by recurling particles with mostly transverse momentum. These particles perform many loops in the detector and if the loops are closely spaced many hits are created in one sensor. As it takes a bit of time to perform a loop, about 2 ns for a 30 MeV positron, the number of hits picked up from recurling particles decreases when a smaller time stamp size is chosen. As a result, the tail towards higher numbers of hits is longer for larger time stamp sizes. For all investigated time stamp sizes, the distributions for layers of the central detector station extend to larger numbers of hits than the distributions for the recurl layers due to the higher density of tracks. Although the outer layers see in general less hits per area, the distribution of the number of hits in the vertex and outer layers as defined above does not differ much. The outer layers are similarly affected by looping particles as the vertex layers. Looping particles are no problem in the recurl stations as the particle is absorbed in the material of the beam line before performing a turn. The tail towards higher numbers of hits stems from events where more than one particle hits the sensor. This contribution exists of course also for the vertex and outer layers.

On the right-hand side of figure 3.3, the efficiency is shown if only a

Time stamp size	4 ns		8 ns		16 ns	
	$\leq 10^{-2}$	$\leq 10^{-3}$	$\leq 10^{-2}$	$\leq 10^{-3}$	$\leq 10^{-2}$	$\leq 10^{-3}$
Vertex layer	5	13	7	16	10	19
Outer layer	5	12	7	15	9	17
Recurl layer (US)	3	9	3	12	5	14
Recurl layer (DS)	3	7	2	11	4	13

Table 3.2: Number of hits with the same time stamp that needs to be processed in order to keep the inefficiency below 1 % or 0.1 %.

limited number of hits with the same time stamp can be processed. As expected, the longer the tail towards higher numbers of hits per time stamp, the lower the efficiency for a given maximum number of hits. Therefore, a larger efficiency is achieved in the recurl layers than in the central detector layers. And of course also a smaller time stamp size yields better efficiencies.

Table 3.2 lists the number of hits that need to be processed in order to keep the inefficiency below 1 % or 0.1 %. Assuming, the counter for the number of hits per time stamp is realized with a certain number of bits, powers of two are a natural choice for the maximum number. For a time stamp size of 4 ns, eight hits per time stamp are sufficient for an inefficiency below 1 %, and sixteen for 0.1 %. The same numbers can be chosen for 8 ns time stamps. The chosen maximum inefficiency limits are just fulfilled. For 16 ns time stamps, the width of the counters needs to be increased by a factor of two.

3.3 THE MUPIX TELESCOPE

The MUPIX TELESCOPE [184] is a particle tracking telescope consisting of four or eight planes of MUPIX sensors for tracking and two additional scintillating tiles in the front and back for precise timing information (see figure 3.4). The telescope is used for efficiency and noise measurements as well as measurements of the spatial and time resolution of sensor prototypes at test beam.

In addition, the MUPIX TELESCOPE represents an integration test of the data acquisition at the front-end of the *Mu3e* pixel detector. Many of the later firmware components are already implemented on a smaller scale in the telescope data acquisition chain.

The telescope features a high rate capability. Furthermore, thin




Figure 3.4: Schematic of the MUIX TELESCOPE. Four layers of MUIX sensors are framed by a scintillating tile in the front and in the back [187].

MUIX sensors and thin, custom-built printed circuit boards (PCB) reduce multiple scattering of the beam particles and allow to operate at low momentum beams.

Four MUIX sensors connect to one STRATIX IV GX FPGA DEVELOPMENT BOARD [185] equipped with a STRATIX IV GX EP4SGX230KF40 FPGA from ALTERA [186]. The FPGA is connected via PCIe bus to the main memory of the data acquisition computer. In the case of the eight layer telescope, two FPGA development boards are installed in the PC. A first test using a front-end board prototype instead of the commercial development board is conducted in March 2018.

Throughout the time, the sensors were replaced with recent versions of the MUIX prototypes requiring amongst others changes in the data acquisition chain. In the following, only the latest two prototypes MUIX7 and MUIX8 are considered.

3.3.1 Data Path

The data acquisition chain of the MUIX TELESCOPE is displayed in figure 3.5. Four MUIX sensors are connected to the FPGA board housed in the data acquisition PC. The pixel sensors continuously transmit serialized and 8 bit/10 bit encoded data on a 1.25 Gbit/s LVDS link.

The serial data from the sensors is received at the FPGA and deserialized. The word boundary is determined on the basis of the comma word K.28.5 which is sent repeatedly after a reset or power-up of the sensor. Thereafter, the data is 8 bit/10 bit decoded and the data words are aligned to the phase of the FPGA internal clock at 125 MHz with a FIFO.

The sensors send hit data as well as status information and control words. This information is disentangled in the data unpacker based on the data protocol of the MUIX output and control words are

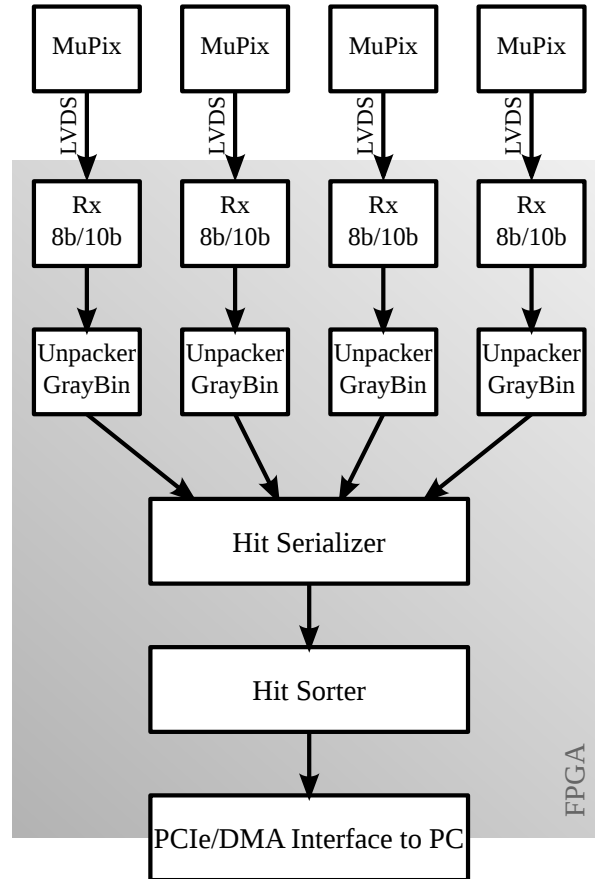


Figure 3.5: The data acquisition chain of the MUPIX TELESCOPE. Data from the MUPIX sensors is received (Rx) at the FPGA and 8 bit/10 bit decoded (8b/10b). The data passes the unpacker and is converted from Gray code to binary (GrayBin) before being multiplexed (hit serializer) and time sorted (hit sorter). Data transfer to the data collection computer happens via a PCIe interface with possibility for DMA transfer. The data chain explicitly shown corresponds to a telescope of MUPIX7 sensors. In the case of MUPIX8, each sensor has four outputs and thus all components before the hit sorter exist in four copies.

removed. The information of one hit or one binary counter is partitioned into three (MUPIX7) or four (MUPIX8) 8 bit data words in the MUPIX readout state machine. These are merged to one data word in the unpacker. After this, the time stamp in the data words is converted from Gray encoding to natural binary encoding.

As the MUPIX data output takes four clock cycles of the 125 MHz clock to transmit the information about one hit or one binary counter, the data stream of four links can be merged into one. This is performed in the hit serializer.

Afterwards, the data is sorted by time stamp in the hit sorter and fed into the PCIe interface that is the connection to the data acquisition computer. The telescope data acquisition uses three base address registers (BAR) to the PCIe bus for data transfer to and from the FPGA. The data stream is written to 256 KiB of 32 bit memory which can be read from the PCIe bus. In addition, 64 registers of 32 bit can be written by PCIe and read from the FPGA. This BAR is used to control the FPGA, for instance for reset or to adjust options. There are additional 64 registers of 32 bit writeable by the FPGA and readable from PCIe used to communicate status and debugging information to the computer. Data transfer can be either performed via polling at low hit rates or via direct memory access (DMA) which allows to process high data rates [157].

On the data acquisition PC, data is collected and stored to disk. Online data quality checks and online track reconstruction allow to constantly monitor the performance of the system during operation. A detailed analysis of efficiency, noise, etc. is performed offline.

In the context of this thesis, the hit serializer is adapted to the needs of the investigated prototypes as well as the data unpacker and hit sorter are developed. The implementation of the first two components is straightforward. The hit serializer is basically a standard multiplexer. The data unpacker is realized with a simple state machine that respects the data protocol of the MUPIX sensors. Both components are therefore not further discussed in the following. The hit sorter however is more complicated. As it represents also an important and non-standard component of the *Mu3e* front-end firmware, its implementation and evaluation is presented in detail in the following.

3.4 THE HIT SORTER ENTITY FOR MUPIX7

Hits are continuously sorted by time stamp as data comes in. The hit sorter has to be able to cope with the incoming data rate and the

spread of time differences between the hits.

For this purpose, a large memory is used as a circular buffer. Incoming hits are written to the memory in a time ordered manner. This is achieved by using the time stamp as part of the address at which the data is written to the memory. The *oldest* entries are then read back first. In this way, the time order is restored at the output of the sorter.

The memory has to be suited to the needs of the data acquisition. On the one hand, it should be chosen large enough to record all incoming hits and thus avoid inefficiencies. On the other hand, the usage of memory and logic resources should be kept at a reasonable level.

For the MUPIX TELESCOPE consisting of four planes with one MUPIX7 sensor each, the address is built by concatenating the 8 bit time stamp and an additional 4 bit counter. The counter allows to store up to fifteen hits with the same time stamp. The counter value of zero is used to indicate that no hits for the specific time stamp are received. In this way, empty entries can be easily skipped when reading back the data from the memory.

Each data word is 26 bit and contains information about the position of the sensor in the telescope (2 bit), the address of the hit pixel (16 bit), and the time stamp of the hit (8 bit).

The MUPIX7 is operated at a time stamp frequency of 62.5 MHz, thus one time stamp unit corresponds to 16 ns. So, the sorter memory spans a time range of 4.096 μ s. The sorter memory is implemented as a dual-port RAM (random access memory) allowing to have simultaneously one write and one read access to the memory per clock cycle. The memory is divided into eight blocks each spanning 32 time stamp units or 512 ns. A block can be either accessed in the writing or reading process. This prevents reading and writing at the same address and ensuing data corruption.

An outline of the hit sorter memory is given in table 3.3. Figure 3.6 shows the flow of blocks in the hit sorting algorithm. The relevant input and output signals to the hit sorter entity are shown in figure 3.7.

3.4.1 *Writing Process*

There are two types of information that are sent by the sensors to the data acquisition chain: hit information consisting of addresses and time stamps of the hits and a larger overall binary counter (24 bit) sampled at the same frequency as the hit time stamps. These two are disentangled in the data unpacker prior to the hit sorter. Only

Address	TS[7:5]	TS[4:0]	Hitcounter[3:0]					
Time stamp	Block		Hit number					
0	Block 0	0	1	2	3	4	...	15
1		1						
2		2						
3		3						
4		4						
5		5						
6		6						
7		7						
8		8						
...		...						
31		31						
32	Block 1	0						
...		...						
63		31						
Blocks 2-6								
224	Block 7	0						
...		...						
255		31						

Table 3.3: Outline of the hit sorter memory for MUIX7. TS stands for time stamp. The notation $[n_1 : n_2]$ denotes a vector of bits between the positions n_1 and n_2 with $n_1 > n_2$.

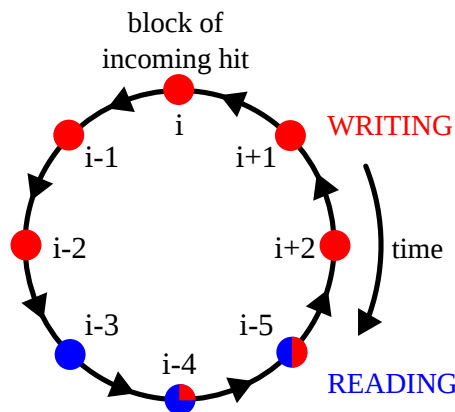


Figure 3.6: The block structure of the hit sorter memory for MUIX7. As soon as reading of a block is finished, it becomes available for writing again. This is indicated by the partially red dots.




Figure 3.7: Inputs and outputs to the hit sorter entity for MUPIX7. The various error counters for debugging purposes are not shown explicitly.

hit information is written to the hit sorter memory. Information about the binary counter is registered for the later data output. The writing process relies on a simple state machine. If the incoming information to the hit sorter is hit information, it is written to the sorting memory. In two cases, a hit cannot be stored and thus the information is lost: *overflow* of the hit counter and *old* hits. In the case of overflow, the counter for the number of hits with the corresponding time stamp has already reached its maximum. In this case, the last hit is overwritten. An overflow flag for the corresponding time stamp is raised and the overflow counter is incremented. The incoming hit can also have a time stamp that deviates significantly from the preceding hits. This happens usually when the occupancy in a column is high. In this case, it takes many cycles to read out all hits from the column and hits can thus be *old* when they arrive at the hit sorter entity. This is problematic as it is important to ensure that the writing process does not write into parts of the memory that are currently accessed by the reading process in order to avoid data corruption. Therefore, there is a set of flags indicating the status of the block. The writing side is not allowed to access blocks that are made accessible to the reading side but are not yet read, that are currently read from, and that are read but not yet reset. If the incoming hit belongs to one of these blocks, it must be rejected and a corresponding error counter is incremented. If there is no overflow and no conflict with the reading process, the hit counter for this time stamp is increased and the full hit information is written to the memory at an address consisting of the time stamp and counter. The block is marked *writing* and a counter for the number of hits written to memory is increased. If the event happens to be the first event in this block, the current status of an FPGA-internal 125 MHz counter is stored in addition as internal

time stamp. This can be regarded as arrival time to the hit sorter. If the incoming data is a binary counter, it is stored per block in an array. The entry of the current block is updated.

The *writing* flag is lifted for blocks that are three to five blocks older than the block of the incoming hit. These blocks are marked as *done* indicating that data is written to these blocks and that writing has finished (see figure 3.6). The reading side is now allowed to access these blocks. A safety margin of two blocks around the block of the incoming hit is chosen in order to be robust also against larger fluctuations of the time stamps of incoming hits. In this way, hits are rejected if they deviate by more than about $1\ \mu\text{s}$ from the previous hit.

Sometimes no hits arrive for a certain period, for instance when the beam is off. The remaining hits in the memory should be read so that no hits from very different times are mixed. In this case, lifting the *writing* and assigning the *done* flag by the means of incoming hits is circumvented by using the binary counter information.

On the MUPIX7, the status of the binary counter is sampled for output when a new readout cycle on the sensor starts, but the output happens only at the beginning of the next readout cycle. In this way, the relation between the timing of the hits and the binary counter depends on the number of hits in the previous readout cycle. In addition, the readout of an individual hit might have taken a while. Moreover, binary counter information and hits are received from four sensors. In order to account for this variable correlation in time between hits and binary counter, a period equivalent to two blocks is waited before the binary counter is used to release blocks for the reading process. Hits arriving after this waiting period are likely to be outside of the writing range defined by the last hit so that longer waiting periods are not justified. The waiting period is however adjustable during operation via a register.

The solution with the waiting period is to some extent unpleasant as the blocks could in principle be released earlier. In addition, the time relation between the hits and the binary counter becomes even weaker when more sensors are read out. Updating the hit sorter for the use with the MUPIX8 prototype, a different approach independent of the binary counter is adopted (see section 3.5).

When the reading side signals that a block is completely read, the hit counters and overflow arrays of this block are cleared and the block is ready to be written to again.

3.4.2 Reading Process

The number of addresses per block exceeds the number of data words that can be read on average during the time period of a block by a factor of 7.5. It is therefore important to identify addresses with data in advance and to skip empty entries during reading in order to prevent data congestion.

The reading process constantly monitors if blocks with data can be read, i. e. blocks are marked *done*. The *done* flag is only lifted in a reset after reading of the block has finished. Therefore, it is further ensured that the block is neither read from at the moment nor the reading has just finished but the reset not yet.

Within the blocks with data, the block closest to the currently read block (*current block*) becomes the next block to be read. In order to assemble the address for reading from the memory, the first time stamp with data and the next time stamp with data larger than the current time stamp of the reading address are identified for the current block. For this purpose, a copy of the hit counters from the writing process is scanned for counters larger than zero.

There are two clock cycles of latency until data of the requested address becomes available at the output of the sorting memory. This has to be accounted for during reading. In order to reduce delays, the reading address is therefore assigned unlocked⁵.

3.4.2.1 The Reading State Machine

Reading from the hit sorter memory and data output are controlled by a finite state machine. A flow diagram of the various states is shown in figure 3.8. The data output protocol is outlined in table 3.4.

After start-up and in the case no data is available, the state machine remains in the IDLE state. As soon as blocks with data become available for reading, the readout process and data output is started. The corresponding block becomes the *current block* and a flag is raised indicating that this block is now in *reading* status. The start of the data output is signalled by a dedicated BEGINOFEVENT data word.

The following states HEADER1 to HEADER5 write header information to the output while waiting for a valid read address and valid

⁵ This is no good practice for larger firmware designs. The status of the signal might be requested at a point where the signal is in a transitional state. Due to increased routing densities, this risk increases for larger designs. A clocked process guarantees on the other hand that the signal is valid at the clock transition. In the MUPIX8 hit sorter, the address is therefore assembled in a clocked process.

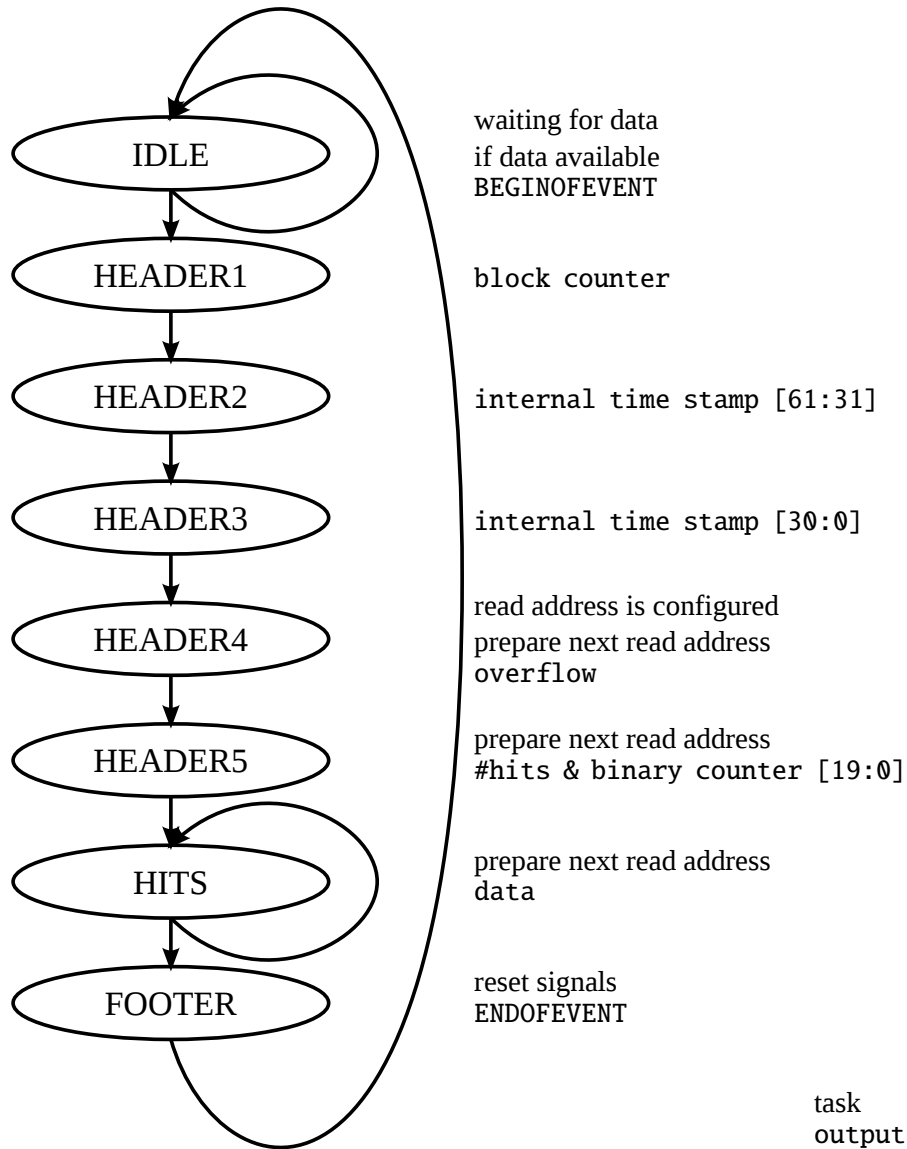


Figure 3.8: The reading state machine of the hit sorter for MUPIX7. For the sake of simplicity, states for the readout of trigger or hit bus information are not shown. The transition to these states is possible in IDLE.

Word offset	Value	Comment
0	0xFABEABBA ⁶	Start of data output
1	0b0 & Block counter	
2	0b0 & FPGA internal TS [61:31]	at 125 MHz
3	0b0 & FPGA internal TS [30:0]	
4	Overflown TS	
5	Time word	
6	Data word	
...	<i>Output of data words is repeated for a total of N hits</i>	
5+N	0xBEEFBEEF	End of data output
Time word	0b0 & N hits[8:0] & Chip[1:0] & Binary counter[19:0]	
Data word	Hit marker[5:0] & Chip[1:0] & Address[15:0] & TS[7:0]	

Table 3.4: Data output of the hit sorter entity for MUPIX7. The abbreviation TS stands for time stamp. N hits is the number of hits in this readout block. Address contains 8 bit column and 8 bit row address information. Chip is an identifier for the position of the sensor in the telescope. Hit marker is a fixed bit pattern.

data output at the sorting memory. In HEADER1, the hit counters and overflow signals of the writing process are copied for the current block. A counter of the number of sent data blocks is output. The following two states output the FPGA-internal time stamp of the moment the first data word of this block arrived at the hit sorter. This time stamp is 62 bit wide and therefore split into two output words. In HEADER2, the first time stamp with data in the current block is identified. And in HEADER3, the *current time stamp*, which is part of the reading address, is set to the value of the first time stamp with data. Thus, the reading address is valid in the following HEADER4 state.

From this point onwards, the reading address is updated every clock cycle until all data from the current block is read. If the hit counter of the current time stamp is larger than one, it is simply decremented. If it is one, i. e. the last entry for this time stamp is already requested by the current reading address, the current time stamp is moved to the next time stamp. If there is no possible next time stamp, reading of the block has finished. The state machine is not yet terminated in

⁶ Here and in the following, hexadecimal numbers are represented with the prefix 0x and binary numbers with 0b. The operator & denotes concatenation.

this case as it takes two additional clock cycles until the data is available at the output of the sorting memory. Instead, a flag is raised indicating that the block is completely read. This flag is pipelined so that it can be recognized two clock cycles later.

The output of state HEADER4 is an array indicating at which time stamp values an overflow has happened. In addition, there is a counter for the number of overflown hits. This number is not part of the data stream but can be accessed through the registers.

In the following state HEADER5, the data output contains information about the binary counter. The binary counter is stored the last time a binary counter with a time stamp belonging to the current block arrived in the writing process. The important information is contained in the upper sixteen bits of the binary counter. These bits allow to identify the timing of the hits on a longer time scale. The lower eight bits are equivalent to the time stamp of the hits. The upper three time stamp bits correspond to the block and are kept for the output data word. The lower five bits are set to zero as they do not provide additional information. Also, the sensor position in the telescope is given for the sensor from which the binary counter is received. As the binary counter is not generated at the same time as the hit time stamps, it might have overflown. This is corrected for. In addition, the output data word contains the number of hits in the current block. The four most significant bits of the binary counter are truncated in order to fit the binary counter and number of hits into a single 32 bit data word.

With the transition to the next state HITS, for the first time in this readout cycle a valid data word is available at the output of the sorter memory. The data word as received from the memory is output with a preceding constant bit pattern which can be used in the following data acquisition to identify the data word as hit information.

The state machine remains in HITS until it is signalled that all hits from the current block are read. In this case, it proceeds to the final FOOTER state. Here, the *reading* flag for the current block is lifted and the block is marked for a reset in the following clock cycle after which it becomes again accessible to the writing process. The state machine proceeds to the IDLE state and waits for new blocks to be read.

In addition to hit information of the MUPIX, also data from the hit bus and trigger⁷ are processed in the case of the MUPIX TELESCOPE. If the FIFOs of either hit bus or trigger contain data, those are read

⁷ Trigger information stems from the scintillating tiles. The hit bus information is the time-over-threshold of a select pixel.

and output every alternate cycle of the state machine. For this purpose, the empty signal of those FIFOs are checked in the IDLE state.

3.4.3 *Resource Usage*

The implementation of the hit sorter entity requires about 4.5 % of the logic resources available on the chosen STRATIX IV FPGA. The sorting memory can store up to 4096 26 bit words which amounts to a memory size of 13 KiB. This is less than 1 % of the available memory.

3.4.4 *Tests of the MuPix7 Hit Sorter*

During the design and before implementation in the telescope data acquisition chain, the hit sorter entity is tested in a behavioural simulation using MODELSIM ASE 10.4B [188]. For this purpose, a data emulator is developed that produces 8 bit data words according to the MUPIX7 data output scheme. The frequency of the data words is adjustable. Whereas the row and column address are implemented as simple counters, the lower time stamp bits can be pseudo-randomized with a 4 bit pseudo-random binary sequence [189]. Four emulators are implemented corresponding to four MUPIX7 sensors. The emulator output is connected to the data unpacker and the data acquisition chain up to the hit sorter is simulated.

The hit sorter is tested for a large range of data rates. No data loss is observed in the behavioural simulation.

In the synthesis and routing process, the various logic elements are assigned a fixed position in the FPGA. This generates small timing variations between the signals. Therefore, the data emulator is implemented into the telescope firmware. The hit sorter can be tested after routing with a known data input pattern. In addition, the data acquisition chain is extended up to the data transfer to the PC. In contrast to the behavioural simulation, tests with the FPGA-internal data emulator can be performed in real time and thus a large number of hits can be processed.

As in the simulation, no data loss is observed in this test.

The hit sorter is integrated in the firmware of the MUPIX7 telescope and is used at numerous test beams at PSI, DESY, MAMI and SPS and thus under very different conditions. Under normal

working conditions⁸ the hit sorter operates at efficiencies well above 99.99 %. Figure 3.9 shows a high rate scan of the efficiency of the hit sorter performed with an electron beam of 720 MeV at MAMI in August 2017. The efficiency is derived from a hit counter before and after the hit sorter entity. The sensors are operated under stable conditions. Only the beam rate is varied.

At smaller hit rates, a slightly higher efficiency is achieved. From about 0.5 MHz onwards, the efficiency stays about constant.

Inefficiencies can have various origins. Hits can be lost because of overflow or because they are too old. In addition, a small number of hits remains in the memory when the data taking is finished. When a long run with a large number of hits in the telescope is taken, the latter is negligible.

It seems that the decline in efficiency at small rates is due to the time structure of the hits. It happens, that the sensor has a few *hot* pixels which turn on with a high frequency even though not being hit by a particle. These hot pixels can block the readout of a column so that it can take a while for a real hit to be read out. As long as only a few hits are created on the pixel matrix, the chances for old hits are low as exactly the columns with the hot pixels need to be hit. With increasing hit rate, the probability increases but also flattens out. This matches the observation of the high rate scan.

Overflow is no problem over the whole range of investigated hit rates. Otherwise a clear decrease of the efficiency with the hit rate would be visible.

The average efficiency during this scan is 0.999 95(1), and the inefficiency of about $5 \cdot 10^{-5}$ is negligible in comparison to the efficiency of the sensor itself. Rates of up to 1.7 MHz are tested. This is about half the hit rate that is expected for one sensor of the vertex layers, and thus more than the expected hit rate per link.

3.5 THE HIT SORTER ENTITY FOR MUPIX8

As in the case of the MUPIX7 prototype, also the hit sorter for the MUPIX8 is tested and evaluated with the MUPIX TELESCOPE.

The MUPIX8 has 10 bit of time stamp information. Thus, the memory of the hit sorter entity is adapted accordingly. As in the case of MUPIX7, a dual-port RAM is used in a circular buffer scheme. There are 1024 possible time stamp values, corresponding to a range

⁸ As the telescope is used to characterize the pixel sensor prototypes, the device under test is often operated in far from optimum conditions with for instance large noise rates. In this case, the hit sorter finds its limitations as it is not designed for these conditions.




Figure 3.9: Hit sorter efficiency in a rate scan at MAMI in August 2017. Data courtesy of A. Tyukin.

in time of $8.192\mu\text{s}$ when operating at the nominal time stamp frequency of 125MHz . The counter per time stamp is reduced to 3 bit in order to keep memory and logic usage at a reasonable level. By introducing an additional non-empty flag, a maximum of 8 hits per time stamp can be written to the memory. A data word is 30 bit wide which results in a memory size of 30 KiB per input to the hit sorter. The memory is divided into eight blocks defined by the upper 3 bits of the time stamp. Each block spans a range of 128 time stamp values, which is about $1\mu\text{s}$ at 125MHz . The memory is outlined in table 3.5.

A schematic of the MUPix8 hit sorter entity is shown in figure 3.10.

Based on the fact that the MUPix8 data protocol does not provide binary counter information in the case when no hits are available, the hit sorting algorithm for the MUPix8 uses an FPGA-internal time stamp as reference for determining blocks. This allows in addition the use of multiple data inputs to the hit sorter entity.

The frequency by which the internal time stamp is incremented can be adjusted via a register in order to match the frequency with which the time stamps are sampled at the MUPix8. It can be set to fractions of the core clock of the FPGA, usually at 125MHz .

The assignment of blocks accessible to the writing and reading process of the hit sorter are shown in figure 3.11. The internal time stamp sets the *sum block* which defines the border between the

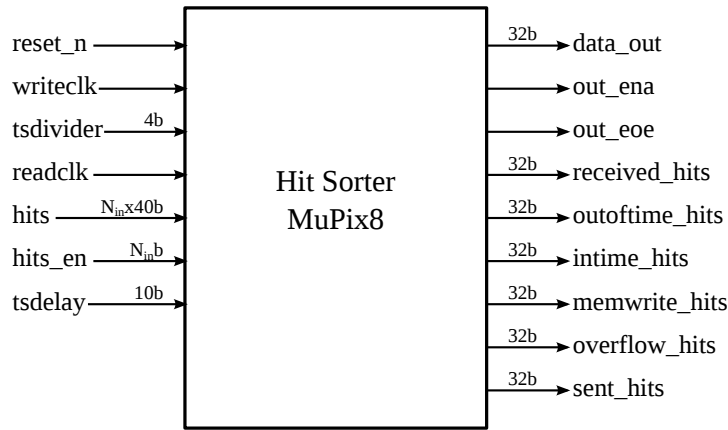


Figure 3.10: Inputs and outputs to the hit sorter entity for MUPIX8. N_{in} denotes the number of input links to the hit sorter which is set to four in the tested scenario.

Address	TS[9:7]	TS[6:0]	Hitcounter[2:0]							
Time stamp	Block		Hit number							
0	Block 0	0	0	1	2	3	4	5	6	7
1		1
2		2								
3		3								
4		4								
5		5								
6		6								
7		7								
8		8								
...		...								
127	127									
128	Block 1	0								
...		...								
255		127								
Blocks 2-6										
896	Block 7	0								
...		...								
1023		127								

Table 3.5: Outline of the hit sorter memory for MUPIX8. TS stands for time stamp. There is one memory instance per input to the sorter which corresponds to one memory instance per four input links to the FPGA.




Figure 3.11: The block structure of the hit sorter memory for MUX8.

blocks accessible to the writing and reading side. Therefore, it is necessary to be able to shift the actual value of the internal time stamp such that the time stamps of the incoming hits ideally all belong to blocks accessible to the writing side. The internal time stamp can be shifted over the whole range of possible time stamps, i. e. from 0 to 1023.

The hit sorter entity can handle more than one input data stream. The current implementation is tested with four data inputs but more inputs are also possible. As each input is fed by a 4-to-1 multiplexer, this implementation is capable to operate with a maximum of 16 input links on the FPGA board, thus satisfying the requirements of the MUX TELESCOPE.

For each input, a separate writing process with a corresponding memory is generated. Accordingly, there is a 3 bit counter per time stamp for each input. The writing process is described in the following.

3.5.1 Writing Process

A valid incoming hit is indicated by a hit enable flag⁹. The data is written to memory with an address built from the 10 bit time stamp

⁹ The data unpacker entity rises the hit enable flag both for hit information and for the binary counter. In the case of the binary counter, a counter enable flag is set as well. Thus, the hit enable flag at the input of the hit sorter is vetoed by the counter enable.

Received data word [39:0]	Chip[3:0] & Matrix identifier[3:0] & Address[15:0] & Charge[5:0] & TS[9:0]
Stored data word [29:0]	Chip[3:0] & Matrix identifier[3:0] & Address[15:0] & Charge[5:0]

Table 3.6: Data words as received at the input of the hit sorter entity for MUPIX8 and as written to the memory. The time stamp information is truncated when writing to the memory.

Chip is an identifier for the position of the sensor in the telescope. Matrix identifier encodes the sub-matrix of the sensor and takes the values 0xA, 0xB, and 0xC. Address contains 8 bit column and 8 bit row address information. Charge is information about the analogue pulse. TS stands for time stamp.

and a 3 bit counter, and the counter is incremented accordingly. In addition, a non-empty flag is set for this specific time stamp. In this way, the hit data is stored in a time ordered way. The actual time stamp bits are truncated and not written to the memory as this information is also contained in the memory address. The information contained in the hit data is shown in table 3.6.

As in the case of the MUPIX7 hit sorter, there are two cases in which a hit cannot be stored and thus the information is lost: the hit counter has already reached its maximum (overflow) or the incoming hit is out of the range of the writing process.

In the case of overflow, all following hits with the same time stamp have to be rejected. Only a per-block-counter for the number of overflow hits is kept as well as an overflow flag per time stamp.

It is important to ensure that the writing process does not write into parts of the memory that are currently accessed by the reading process in order to avoid data corruption. Therefore, hits with time stamps belonging to the sum block, reset block and reading block 1 and 2 are rejected (see figure 3.11). Again, the hit information is lost and only a counter for these out-of-time hits is kept. At this point, it is crucial to choose a reasonable shift for the internal time stamp in order to ensure that most or ideally all time stamps of the incoming hits belong to blocks accessible to the writing process. By comparing the number of in-time and out-of-time hits during operation, an appropriate shift can be easily chosen.

For MUPIX8, it is expected to observe larger differences between the time stamps of the incoming hits than for MUPIX7 as the readout column length has increased from 20 to 200. This is partially compensated by increasing the duration of one block from

512 ns to 1024 ns.

As the hit sorter entity is implemented as a circular buffer, blocks need to be cleared or reset before they become available for the writing process again. This happens for all time stamps in the *reset block* as it contains the oldest information. All counters, non-empty flags and overflow flags of the writing side belonging to the reset block are cleared.

The *sum block* marks the transition from the writing to the reading process. At this point, non-empty signals for the reading process are prepared which allow to skip over memory entries with no data. These signals include a non-empty signal per time stamp that indicates that at least at one input a hit with the corresponding time stamp was received. In addition, a per-block non-empty signal is derived. Only information belonging to the sum block is updated.

3.5.2 Reading Process

After the hits are written to the memory, they are read in time order. This is performed in the reading process. As the hits from all inputs need to be merged to one data stream, there is only one reading process that reads from all memory instances. Hence, the reading process needs to efficiently skip over entries without data in order to prevent congestion.

The reading process constantly monitors if blocks with data are available for reading. When a change in sum block¹⁰ is detected, i. e. when the current value of the internal time stamp has moved to the next block, the per-block and per-time stamp non-empty signals are copied for the new sum block.

By the means of the per-block non-empty signals the next block with data is determined. Starting with read block 2, the oldest block which is labelled non-empty and which is not identical to the currently read block becomes the next block to be read. The counters, non-empty and overflow signals for each input and time stamp in the range of the next block are copied. In addition, a signal is set indicating that a new block with data is available.

After reading of a block is completed, the per-block non-empty flag

¹⁰ The sum block signal in the reading process is actually delayed by five clock cycles with respect to the writing process in order to allow for enough time to derive the per-block non-empty signals.

for this block is reset so that this block will not be mistaken for a block with data when determining the next block.

If read block 2, read block 1 and sum block are empty, no new next block can be identified. In this case, the value for the currently read block is set to the value of the reset block once the data output has finished. This block is not read, but by setting the value to the reset block it is ensured that the entire range of blocks accessible by the reading process can be used to determine the next block and no block is occupied due to an abandoned value for the currently read block.

Just as the next block, the next time stamp to be read within this block is determined from the corresponding non-empty signal. The smallest time stamp which is not empty and not equal to the currently read time stamp becomes the next time stamp to be read. As in the case of blocks, this non-empty signal is reset once all hits with this time stamp are read.

If a block is completely read, the next time stamp signal cannot be updated to a new value. The non-empty flag of the next time stamp is however reset, and thus this flag can be used to indicate that reading of the block has finished.

3.5.2.1 *The Reading State Machine*

The reading process relies on a finite state machine (see figure 3.12). After a change in the reading address, it takes two clock cycles until the new data is available at the output of the memory. This is accounted for in the state machine. The address at which the memory is read is constructed from the 10 bit time stamp and a 3 bit counter. The time stamp in the address is the same for all the memory instances of the various inputs, whereas the counter is specific to the input. Thus, at some memory instances there might be no entry at this address. The state machine ensures that only the output of memory instances that have entries is read. In fact, also the next block and next time stamp indicate that there is at least one memory instance with entries.

As long as no blocks with data are available to the reading process, the state machine remains in the IDLE state with no output to the PCIe interface. As soon as data becomes available, the state machine moves on to the STARTREADOUT state in which the number of sent blocks is incremented and the BEGINOFFEVENT data word is output signalling the beginning of a data block to the data acquisi-

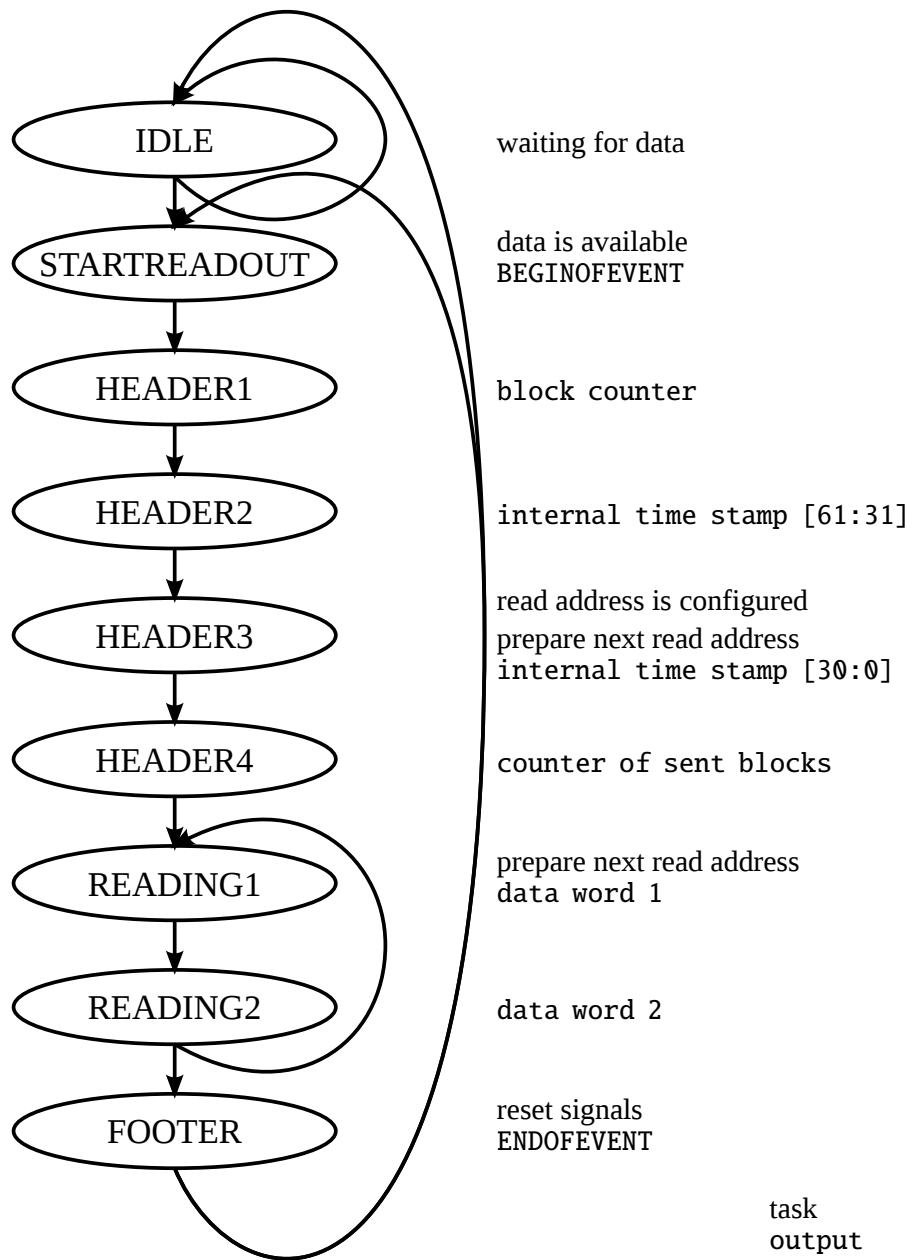


Figure 3.12: The reading state machine of the hit sorter for MUPIX8. For the sake of simplicity, states for the readout of trigger or hit bus information are not shown. The transition to these states can either happen in IDLE or FOOTER.

tion PC. After that, the state machine proceeds through the states HEADER1 to HEADER4 in which header information is written to the output before going to the states READING1 and READING2. The header information consists of the number of blocks elapsed since the last reset, 62 bit of internal time stamp split over two clock cycles, and the number of blocks actually sent out by the hit sorter. During the output of the header information, the read address is prepared. First, the value for the current time stamp to be read is updated. Then, the input is determined by checking if the non-empty flag is set for the specific time stamp and input. The lowest non-empty input is to be read first. Thus, the read address is ready when transiting from HEADER2 to HEADER3. The read address has to be configured two clock cycles in advance before the corresponding data is read from the output of the memory. Therefore, the first read address needs to be ready in HEADER3 in order to have the data available in READING1.

As one data word in the case of the MUIPX8 hit sorter spans more than 32 bit, it has to be split into two output words. Accordingly, data is read from the memory every other clock cycle and likewise the read address is updated. This happens once in HEADER3 and then repeatedly in READING1 given that there is still data available for the currently read block.

The read address is updated as follows. If the counter for the currently read time stamp and input is not yet zero, it is simply decremented. Otherwise, the next input with data at the same time stamp is used for the address. If there is none, the current time stamp proceeds to the next time stamp and the corresponding input is determined. In the latter case, the non-empty signal for the previous time stamp is reset so that it will not be used when determining possible next time stamps with data.

In READING1, the first data word is written to the output. Additionally, the second data word is prepared which is output in the subsequent state READING2. The data words are shown in table 3.7.

If all data corresponding to the currently read block has been retrieved, this is recognized in READING2, and the state machine proceeds to the state FOOTER. Otherwise, it returns to READING1.

It might occur that the currently read block is moved out of the scope of the reading process when the sum block changes. In this case, the state machine will directly move from state READING2 to FOOTER in order to prevent data corruption. These events are indicated by a dedicated counter.

Word offset	Value	Comment
0	0xFABEABBA	Start of data output
1	0b0 & Block counter	
2	0b0 & FPGA internal TS [61:31]	at 125 MHz
3	0b0 & FPGA internal TS [30:0]	
4	0b0 & Counter of sent blocks	
5	Data word 1	
6	Data word 2	
...	<i>Output of data words is repeated for a total of N hits</i>	
4+2N	0xBEEFBEEF	End of data output
Data word 1	0b00 & Chip[3:0] & 0b0 & Overflow & Hitlabel[3:0] & Matrix identifier[3:0] & Row[7:0] & Column[7:0]	
Data word 2	0b00 & Chip[3:0] & 0b0 & Overflow & Timelabel[3:0] & Matrix identifier[3:0] & Charge[5:0] & TS[9:0]	

Table 3.7: Data output of the hit sorter entity for MUPIX8. The abbreviation TS stands for time stamp. Chip is an identifier for the position of the sensor in the telescope. Matrix identifier encodes the sub-matrix of the sensor. Charge is information about the analogue pulse. Overflow flags if overflow is observed at the current time stamp. Hitlabel and Timelabel are fixed bit patterns. Data word 1 contains information about the hit address, data word 2 contains time information. Bits that are not used in the data words are set to zero. Except for the output word for start and end of the block, all data words have a zero as most significant bit. This way, they cannot be confused with the beginning or end of the block.

In the FOOTER state, the end of the data block is indicated by the output of the ENDOFEVENT data word. The up to then read block is reset as well as the corresponding per-block non-empty signal. The state machine proceeds either to IDLE or directly to STARTREAD-OUT if a new block with data is available.

The data output at the MUPIX8 hit sorter entity is outlined in table 3.7.

Again, there is the possibility to output trigger and hit bus information in the MUPIX TELESCOPE. The corresponding FIFOs are checked in the IDLE and FOOTER state. Output of trigger or hit

bus information happens at most in every second readout cycle thus alternating with hit information.

3.5.3 *Resource Usage*

The hit sorter in the MUPIX8 telescope can be operated with up to sixteen links. This corresponds to four inputs to the sorter and thus four sorting memory instances of 30 KiB each. In total, about 7% of the memory resources of the STRATIX IV FPGA of the development board are used. The increase in the size of the writing and reading address and the increase in inputs to the hit sorter entity however result in a significant usage of about 45% of the logic resources.

3.5.4 *Tests of the MuPix8 Hit Sorter*

As in the case of MUPIX7, various tests are performed to adjust the dimensions of the hit sorter and verify its suitability.

The behaviour of the entity and preceding components is tested in simulation using MODELSIM ASE 10.4B [188]. The data output of the MUPIX8 is generated using a data emulator [190]. The data words are already 8 bit/10 bit decoded. The data unpacker filters out hit data and creates corresponding flags and the time stamp information is further decoded from Gray to binary. The hit serializer merges four data streams to one. Then, the data enters the hit sorter. This data flow matches the data flow in the FPGA. The hit sorter runs without data loss in the behavioural simulation.

The same data path including an internal data emulator is also implemented on the FPGA. This allows to test the hit sorter in a controlled environment after synthesis. Due to the routing on the FPGA, small timing differences between the various signals might occur. In contrast to the simulation, these tests are performed in real time. This allows to test the processing of a large number of hits. Again, the hit sorter is operated without data loss.

A first test of the updated MUPIX8 hit sorter is performed at DESY in November 2017 at an electron beam of 4 GeV momentum. Only one of the three sub-matrices of the sensors is operated. The hit sorter is nevertheless implemented with four sorting memory instances in order to test this new feature. Each sensor connects to one memory instance.

Figure 3.13a shows the measured efficiency of the hit sorter. The


Efficiency	0.999 67(8)
Inefficiency	$33(8) \cdot 10^{-5}$
Out-of-time hits	$27(9) \cdot 10^{-5}$
Hits remaining in memory	$53(14) \cdot 10^{-6}$
Overflow hits	$9(29) \cdot 10^{-8}$

Table 3.8: Average efficiency of the MUPix8 hit sorter and sources of data loss during test beam measurements at DESY in November 2017.


efficiency is derived from the number of received hits at the sorter and the number of hits that are sent out to the PCIe interface. In addition, error counters are monitored. The relevant counters are the number of out-of-time hits and overflow hits. Furthermore, the number of hits written to the memory is monitored. The difference of this number and the number of hits sent out gives the number of hits that still remain in the sorter memory after the readout has finished. The sources of data loss are shown in figure 3.13b. Table 3.8 lists the average efficiency and the inefficiencies of the various origins.

The overall efficiency of the MUPix8 hit sorter is high. No inefficiency below 99.95 % is observed. An average efficiency of 99.967(8) % is achieved. Figure 3.13a indicates no dependence on the rate of the incoming hits. However, the beam rate at DESY is low compared to the test beam at MAMI used for the rate scan of the MUPix7 hit sorter in section 3.4.4. A hit rate dependence can thus not yet be excluded. In contrast to the tests at MAMI, the data used for the tests of the MUPix8 hit sorter include also runs in which settings of the device under test are scanned. Thus, one of the sensors is for instance at a high noise rate in some of the runs. This causes the spread of the data points.

About $33(8) \cdot 10^{-5}$ of the hits are lost on average. The largest contribution to this loss with about $27(9) \cdot 10^{-5}$ on average are out-of-time hits. In addition, a few hits remain in the memory at run stop and are not read out. This loss is about $5(1) \cdot 10^{-5}$. With the comparably low beam rate at DESY, overflow hits are no issue. In only about one third of the runs, overflow is observed. The contribution of $9 \cdot 10^{-8}$ to the data loss is negligible.



(a) Efficiency of the hit sorter.



(b) Sources of data loss.

Figure 3.13: Efficiency and sources of data loss of the hit sorter for MUPIX8 in a test beam measurement at DESY in November 2017.

3.6 SUMMARY ON THE PIXEL DETECTOR FRONT-END

In this chapter, a firmware-based hit sorting algorithm is elaborated for the use in the front-end of the *Mu3e* pixel detector. The algorithm is capable to sort a continuous stream of incoming hits by time stamp. For this purpose, the time stamp is used as part of the address to write to a circular buffer. The data is then read back in time order.

The test bed for the hit sorter is the data acquisition of the MUPIX TELESCOPE being either equipped with MUPIX7 or MUPIX8 pixel sensors.

In the case of the MUPIX7 hit sorter, blocks are assigned to the writing or reading process based on the time stamp of the incoming hits. As soon as reading of a block is finished, this block becomes available for writing again. The hit sorter is designed to operate with up to four MUPIX7 sensors.

It has proven very high efficiencies of above 99.99 % even at hit rates up to 1.7 MHz.

The hit sorter is updated and further improved for the latest prototype MUPIX8. The sorting memory is increased to fit the increased time stamp size. As the MUPIX8 can be operated at up to three output links per sensor, a simple 4-to-1 multiplexer is no longer sufficient to merge the data stream of four sensors. Therefore, four sorting memory instances are implemented which are read in time order in a single reading process. The blocks for writing and reading are now assigned with an internal time stamp generated on the FPGA.

A first test is performed at total hit rates between 20 kHz and 40 kHz. The high rate capability still needs to be assessed. An efficiency of larger than 99.95 % is observed. The largest inefficiency stems from out-of-time hits. In addition, always a few hits remain in the sorting memory when the readout is finished. Overflow turns out to be no problem at these rates.

The inefficiency is about one order of magnitude larger than in the case of the MUPIX7 hit sorter. This is mostly due to the fact that the column length in the readout is increased from 20 to 200 and likewise the probability of out-of-time hits increases.

It has to be noted, that problems with the time stamp bits are observed for the MUPIX8 sensor. Especially if the on-chip state machine is operated at high frequencies, some bits of the time stamps

seem to be *sticky*, i. e. they do not change their status in time. A possible explanation would be that the driver of the bits of the Gray counter is not sufficiently strong. For this reason, the MUPIX8 is so far only operated with a reduced state machine frequency.

Furthermore, a larger spread of the time stamps of the MUPIX8 is observed compared to MUPIX7. This is possibly due to the increased size of the sensor. The length of the signal lines from the active pixel diode to the comparator in the periphery varies more for the larger sensor and so does the transmission time and also the amplitude loss of the analogue signal.

In addition, the sensor is operated *un-tuned* in the test beam measurement at DESY. In tuning, the individual threshold at the comparator of each pixel is adjusted to create a homogeneous response to the signal created by the incoming particles over the whole sensor or to adjust the noise rate of the individual pixels to a chosen level. This is performed for MUPIX7 but still pending for MUPIX8.

Sticky bits and a larger variation of the time stamps of course influence the performance of the hit sorter. Some improvement can thus be expected from tuning and when the driver strength is increased.

There are a few possibilities to further reduce the inefficiency. In order to reduce the number of out-of-time hits, concepts of both hit sorter designs can be combined. In general, it is easier to define the boundary between blocks accessible for writing and reading based on an internal time stamp especially if a large number of sensors is connected. However, if a block is completely read, it should immediately become available for writing again. In addition, the memory can be divided into a larger number of blocks to allow for more flexibility. This can however only partially remove the out-of-time hits. In the end, it is the time structure of the incoming hits that determines the inefficiency.

One reason for large time differences between the incoming hits are also events with delta electrons or cosmic muons. In the former case, the incident particle transfers a large amount of its energy to an electron of the atoms of the sensor. If the delta electron or cosmic muon moves about parallel to the sensor surface, many pixels are hit and several cycles are required for the readout. In addition, the hits also have a time structure caused by time walk as well as energy deposition in the bulk underneath the depletion zone. In the latter case, the generated charges are collected via diffusion which takes in the order of 100 ns and causes latency. An example of an event with many hits caused by a cosmic muon is shown in figure 3.14. In such cases, the hit sorter records a large number of out-of-time hits.




Figure 3.14: Event with a cosmic muon on the MUPIX8. The muon enters from the right side and has a hard scattering interaction. It traverses the whole sensor. The time structure is caused by time walk as well as by energy deposition in the bulk underneath the depletion zone. Figure courtesy of A. Herkert.

Events with many hits are however not of greater interest in the later experiment as they give rise to ambiguities.

The inefficiency due to hits that remain in the memory can be easily eliminated. In the MUPIX TELESCOPE, the reset procedure needs to be adapted so that the reading process can finish to output the hits. This takes a few μs . In addition, the number of hits in the memory is about constant for constant beam rates so that this type of inefficiency becomes less and less important over long run times, for instance in the later experiment. Again, the reset procedure on the front-end boards should allow to output all data from the memory. Overflow has not yet been a problem in the MUPIX TELESCOPE. But the large scale MUPIX8 sensor is also not yet tested at high beam rates. In principle, the amount of overflow hits is expected to increase. As the time stamp size in MUPIX8 is half of the time stamp size in MUPIX7 and the hit counter size is reduced by a factor of two, the overflow rate should be comparable to the MUPIX7 hit sorter at the same hit rate on the sensor. As overflow is no problem even at high hit rates for the MUPIX7 hit sorter, it is also expected to be no problem for the MUPIX8 hit sorter.

In addition, overflow at the telescope can not be directly translated to overflow in the later experiment. In the telescope, sensors of four layers connect to one FPGA. In the later experiment, sensors from

two layers are connected to one front-end board.

The easiest measure against overflow is increasing the size of the hit counter per time stamp.

A very high efficiency to detect a particle crossing the pixel layers is essential to the sensitivity of the experiment. The overall hit inefficiency should not exceed 1-2 %. In order to reconstruct a $\mu \rightarrow eee$ signal event, each particle has to have at least one hit in each of the central detector layers. With a hit efficiency of 99 %, the reconstruction probability of a signal event is about 89 %. For 98 %, this decreases already to 78 %. Requiring the tracks to be reconstructed at least as 6-hit tracks, the reconstruction probabilities are even lower with 83 % and 70 %, respectively.

How much inefficiency is acceptable for the hit sorter is not yet determined. It depends also on other sources of inefficiency, for instance on the inefficiency of the pixel sensors itself as well as inefficiencies emerging from the small spacing between sensors. It is expected that a hit sorter inefficiency of about 0.1 % is acceptable. From the simulation studies in section 3.2.1, it follows that a hit counter of sixteen is needed for 8 ns time stamps in order to keep the overflow inefficiency below 0.1 %.

Overflow at the front-end hit sorter can be further reduced by a sophisticated assignment of the links to the sorter memory instances. Two sensors that are directly on top of each other should not be connected to the same memory instance as they are likely to have hits with the same time stamp as they are often hit by the same particle. A simple way to assign the links to the memory instances would be to avoid connecting sensors of two layers to one instance. Furthermore, directly neighbouring sensors are likely to be hit by the same looping particle so that it is better to connect sensors that are separated in space to the same memory instance. In the case of the vertex layers, this approach can be further extended to the sub-matrices as these are read out at individual links.

Of course, also a smaller time stamp size, e. g. 4 ns, reduces inefficiencies due to overflow.

Apart from the efficiency, also the bandwidth is important. In the case of the MUPIX8 hit sorter, a data word is fed out every second clock cycle. With a total of six header words, an average hit rate of about 60 MHz can be processed which is slightly less than half the nominal frequency of the FPGA of 125 MHz. In the case of the MUPIX TELESCOPE, this corresponds to a hit rate per sensor of 15 MHz which is sufficient for test beam purposes.

In the later experiment, a larger number of sensors connects to the front-end board. The current implementation would just be sufficient to process the data of the vertex layers. Assuming the most disadvantageous configuration with a fraction of 60 % inner and 40 % outer vertex layers and a total number of fifteen sensors connecting to one front-end board, the incoming hit rate is about 57 MHz on average (see also section 3.2.1). An operation at beam rates only slightly larger than $10^8 \mu/s$ would thus not be possible.

However, the need to use two clock cycles to output one data word is not given on the front-end board. In the telescope, this is required as one data word¹¹ consists of at least 34 bit, but the memory of the PCIe interface is configured in units of 32 bit.

On the front-end board however, the hit sorter is followed by the packet builder that re-arranges the data before it is output over the optical link. The packet builder can of course have a larger input width. At minimum, 38 bit are required if the size of addresses and time stamps on the final sensor are unchanged compared to MUPIX8: 16 bit for the address, (10 + 6) bit for time stamp and analogue pulse information, and 6 bit to encode the sensor number as up to 36 sensors connect to one front-end-board. The 6 bit analogue pulse information might be skipped for instance when time walk corrections are already performed on the front-end board. In this case, the minimum data word becomes 32 bit.

Assuming data output from the hit sorter in every 125 MHz clock cycle and the same amount of overhead as in the MUPIX8 implementation, an average hit rate of 119 MHz can be processed. This is about a factor of two more than the expected hit rate in the busy vertex layers.

A further measure to increase the bandwidth of the hit sorter is to operate at a higher frequency or to write and read two hits per cycle. However so far, there seems to be no need for this.

It has to be kept in mind, that the front-end boards have a single optical link at 6 Gbit/s to the switching boards. This limits the hit output rate to 150 MHz assuming 32 bit data words (40 bit 8 bit/10 bit encoded). This is thus close to the maximum average hit rate that the hit sorter can process.

The hit sorter for the MUPIX8 is designed such that the number of inputs to the hit sorter can be basically freely chosen. The vertex

11 This assumes that the position of the four chips in the telescope is encoded by 2 bit. There are further 16 bit for the pixel address, 10 bit for the time stamp, and 6 bit for the analogue pulse information. The identifier for the sub-matrix is omitted in this counting as it is also encoded in the address.

layers have the largest demands on the number of inputs. Fifteen sensors connect via 45 links to the front-end. This requires twelve sorting memory instances and thus scaling by a factor of three compared to the implementation for the telescope.

This upscaling is however not easily feasible as also the logic and memory usage increases. The memory needed to implement twelve sorting memory instances amounts to about 20 % of the available memory on the ARRIA V FPGA chosen for the front-end board. As the hit sorter is the entity with the largest memory usage, the memory resources of the FPGA are thus no limiting factor.

The logic utilization however is an issue. The current implementation already requires about 45 % of the resources and cannot be simply scaled by a factor of three. If the hit counters are increased from eight to sixteen, the logic utilization increases even further. A further downside of high logic utilization is that the synthesis and routing of the firmware design becomes complicated. This leads to timing problems between signals and increases the risk of errors.

A possible solution to this problem is to use a memory to store the hit counter. This implies changes to both, the writing and reading process. During writing a few cycles of latency are created as the counter first needs to be read from the memory before it can be increased. The same applies to the reading process. At first, the counter needs to be retrieved from the memory. Then the address to read back the actual hit data is assembled. In both cases, this is a pure latency. It is still possible to read and write one hit per clock cycle so that the bandwidth is not reduced, except for about two additional header words.

Part III

PHYSICS STUDIES

Modern particle physics experiments aim to unravel the nature of potential physics beyond the Standard Model in a plethora of signatures. New Physics can hide at very large mass scales or very weak couplings.

With an unprecedented number muon decays to be observed, the *Mu3e* experiment can probe both frontiers. Searches for the rare decay $\mu \rightarrow eee$ are indirect tests of models with new heavy particles that mediate lepton flavour violation at loop or tree level. Light particles with weak couplings can be directly produced.

THE SIGNAL DECAY MEDIATED VIA EFFECTIVE OPERATORS

Lepton flavour violation in the charged lepton sector and thus decays like $\mu \rightarrow eee$ are suppressed to unobservable levels in the SM but are naturally enhanced in many models for physics beyond the SM. Consequently, an observation of the $\mu \rightarrow eee$ decay would be a clear indication for New Physics.

Due to the large number of possible New Physics models, it is unfeasible to test each model individually. Therefore, sensitivities and limits on the branching fraction are typically derived with regard to events that are evenly distributed in the accessible phase space. This neglects the fact that the nature of the potential New Physics gives rise to very specific kinematics of the decay products. Consequently, the sensitivities or branching fraction limits derived for phase space distributed decays cannot necessarily be applied for a specific model of New Physics.

In this chapter, the sensitivity of the phase I *Mu3e* experiment is therefore not only derived for phase space distributed $\mu \rightarrow eee$ decays but also for decays mediated via operators of effective field theories (EFT) up to dimension six. In this way, a basis is provided for comparing the sensitivity for specific models.

Apart from $\mu \rightarrow eee$ searches, lepton flavour violation of charged leptons is searched for in many channels with the most stringent existing and expected limits stemming from processes involving muons. The EFT approach provides the means to evaluate similarities and complementarities of the various searches. Discoveries in some or all of the experiments can be related to certain operators and thus classes of New Physics models. Amongst the experiments searching for lepton flavour violation of muons, the *Mu3e* experiment has the unique possibility to further differentiate some of the effective operators and their chiral couplings if a sufficient number of $\mu \rightarrow eee$ events is observed. This is not possible in $\mu \rightarrow e\gamma$ searches, and requires measurements with different target materials in the case of $\mu N \rightarrow eN$ conversion experiments.

4.1 SEARCHES FOR $\mu \rightarrow eee$ EVENTS WITH MU3E

The *Mu3e* experiment is designed to measure all three tracks from a $\mu \rightarrow eee$ signal event with a high acceptance and excellent resolution. Thus, the full information about the final state can be exploited to separate signal events from background. As there is no other physics process with exactly three electrons in the final state that have an invariant mass equal to the muon rest mass, the search for $\mu \rightarrow eee$ can be performed background-free. The final sensitivity depends on the capability to suppress background and thus on vertex, timing and momentum resolution, as well as on the efficiency to reconstruct a signal event.

SIGNAL EVENTS A signal decay can be identified by two positrons and one electron stemming from a common vertex and appearing coincident in time. The invariant mass of the three particles equals the muon rest mass m_μ , and the sum of the momenta vanishes as muon decays at rest are observed in *Mu3e*.

BACKGROUND FROM INTERNAL CONVERSION DECAYS There are two sources of background: irreducible background from radiative muon decays with internal conversion $\mu \rightarrow eee\nu\nu$, and accidental background arising from combinations of tracks of various muon decays and other processes such as Bhabha scattering and mis-reconstruction.

Internal conversion decays are similar to signal decays as the three charged particles stem from a common vertex. The two additional undetectable neutrinos carry however some of the energy available in the decay. Cuts on the kinematics are thus the only handle to suppress internal conversion background.

BACKGROUND FROM ACCIDENTAL COMBINATIONS With high muon stopping rates, multiple muon decays are observed at the same time. This gives rise to combinations of tracks from different muon decays to one vertex. In order to be mistaken for a signal event, such events need to contain an electron track. Electron tracks can be created in Bhabha scattering of a positron in the detector material as well as in photon conversion and Compton scattering, or can stem from an internal conversion event. Moreover, a positron track from eight hits, i. e. the particle crosses the vertex layers for a second time, can be reconstructed with the wrong sense of rotation and is thus misinterpreted as an electron. In general, also tracks reconstructed from hits caused by different particles and noise hits

contribute to the accidental background.

The three tracks in accidental background events share no common vertex so that this background can be efficiently reduced by constraints on the vertex and relative timing of the tracks. In addition, with timing information the correct sense of rotation can be determined for eight hit tracks eliminating accidental background from misinterpreted positron tracks.

In contrast to internal conversion decays, invariant three-electron masses m_{eee} of about m_μ can indeed be achieved by accidental background although this is extremely rare. With a cut on the total momentum, the number of accidental background events in the signal region can be further reduced as it is very unlikely to fulfill both criteria at the same time.

Constraints on the vertex and timing are more effective for combinations of three uncorrelated tracks than for combinations with two correlated tracks, e.g. electron-positron pairs from photon conversion or Bhabha scattering. Thus, the latter is expected to contribute the most to the accidental background. This type of accidental background is dominated by events in which an electron-positron pair from Bhabha scattering is combined with another positron most likely from a Michel decay. Simulation studies show that such Bhabha events appear in the order of a hundred times more often than the other processes [136], at least in phase I. With an increased muon stopping rate in phase II, also other processes might have relevant contributions to the accidental background.

Simulating accidental background is a computationally intense task due to the vanishingly small production and reconstruction efficiencies. Numerous event frames and thus long run times are required to generate samples that have at least a few events with invariant masses in the vicinity of the muon rest mass. So far, these efforts have concentrated on the dominant Bhabha contribution. In figure 4.1, the results of a study of accidental background by Bhabha events are shown alongside background from internal conversion and $\mu \rightarrow eee$ signal events [136]. For the Bhabha sample, $1.1 \cdot 10^9$ weighted¹ events are generated in which the electron and positron from Bhabha scattering have at least an energy of 10 MeV—the minimum energy needed to traverse all four central detector layers and thus to be visible. The sample corresponds to $3 \cdot 10^{15}$ muon decays. After reconstruction and cuts, only four Bhabha events with an invariant mass of larger than 95 MeV remain and none of them

¹ The events are weighted with the differential decay rate. Thus, the weight corresponds to the number of events.




Figure 4.1: Spectra of the reconstructed invariant three-particle mass of $\mu \rightarrow eee$ signal decays and background from internal conversion and Bhabha events in phase I of the *Mu3e* experiment [136]. A muon stopping rate of $1 \cdot 10^8 \mu/s$ and a total number of $1 \cdot 10^{15}$ muon stops are assumed. The sum of weights of the Bhabha events amounts to 3.56 events with $m_{eee} > 95 \text{ MeV}$ in $1 \cdot 10^{15}$ muon decays.

would pass the invariant mass cut that is defined in the following. It is thus expected, that the contribution of accidental background is negligible for phase I with similar or tighter cuts as chosen in the existing study.

In the following, selection cuts are optimized considering signal events and background from internal conversion events. The existing Bhabha study is referred to where needed.

4.1.1 Simulation of Signal and Background Events

As a baseline, $\mu \rightarrow eee$ events are simulated evenly distributed in phase space. For all of the following considerations, a data set of $1 \cdot 10^7$ frames of signal events is used. These events are produced as *special decays*, meaning that only signal events and no other decays are generated. In addition, in *special decays* at most one decay per frame is generated. If decay products from the decay in the previous frame are still present, i. e. the decay happened late in the previous frame, no new decay generated. Therefore, the number of generated $\mu \rightarrow eee$ decays is smaller than the number of frames. In addition, the muons can also decay in flight in the beam line before reaching the target. Such events are not considered in the following as the

sensitivity will be determined with regard to the number of stopped muons. The total number of generated $\mu \rightarrow eee$ decays in the target region amounts to about $7 \cdot 10^6$.

Similarly to signal events, background from radiative decays with internal conversion $\mu \rightarrow eee\nu$ is also generated as special events. The phase space of the generated events is restricted in order to generate preferentially events in the signal region. Each decay electron is required to have a minimum energy of 10 MeV. With a smaller energy, the particle cannot have sufficient transverse momentum to cross all four central pixel layers and thus no track can be reconstructed. Furthermore, it is required that the invariant mass of the three electrons in $\mu \rightarrow eee\nu$ equals at least 90 MeV. The branching fraction for this selection is determined by means of Monte Carlo integration and amounts to $\mathcal{B}_{\mu \rightarrow eee\nu} = 1.4 \cdot 10^{-10}$. In order to produce a large sample of events, the events are weighted with the corresponding differential decay rate. In this way, the slower accept-reject approach can be bypassed. A weighted sample of $1 \cdot 10^8$ frames yielding about $7 \cdot 10^7$ $\mu \rightarrow eee\nu$ events is generated. The sum of weights corresponds to $2.5 \cdot 10^{19}$ muon decays.

Both signal and background generators are integrated in a detailed GEANT4-based detector simulation that represents the current status of the experimental design. The particles and decay products are propagated through the experiment and the detector response is simulated. All relevant interactions of particles passing through matter are taken into account, for example multiple scattering which is highly important in *Mu3e*.

4.1.2 Reconstruction

The simulated events are reconstructed using the multiple scattering track and vertex fit of the offline reconstruction as outlined in section 2.2.7.

Tracks are first reconstructed as short tracks built from four hits, and are then extended to long tracks of six or even eight hits whenever possible in order to benefit from the improved momentum resolution of recurling particles.

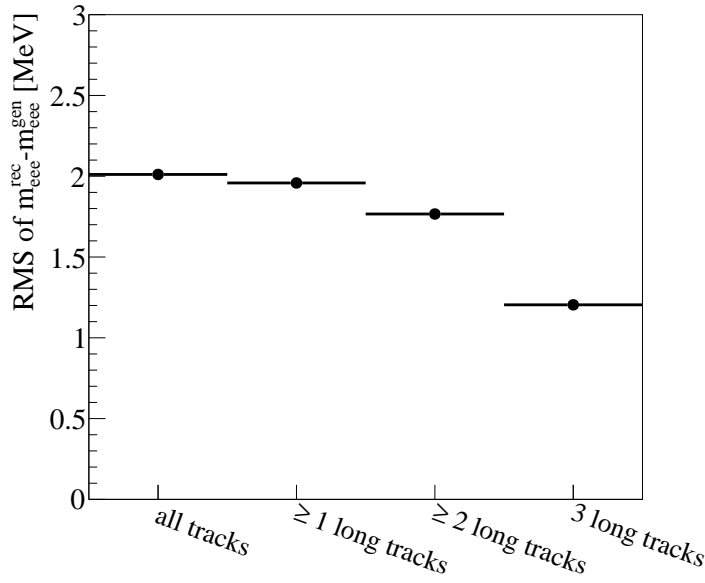

Vertex fits are performed with all possible combinations of one electron and two positron track candidates. A few selections on the tracks are applied. Short tracks with χ^2_{track} value of the track fit of larger than 32 and long tracks with χ^2_{track} larger 48 are omitted to ensure a certain quality of the reconstructed tracks. The inefficiencies in both cases range around 1% [136].

Sometimes, multiple tracks are reconstructed per particle for instance when the particle has high transverse momentum and performs several loops in the central detector station or when it crosses the overlap region of adjacent pixel modules and creates two hits in one layer. In addition, in the case of looping particles one turn can be misinterpreted as a track of a particle of the opposite sign. In order to avoid building vertices from tracks that partially originate from the same particle, selections are made on the opening angles and momentum difference. Two same sign tracks that differ less than 1 MeV in momentum are required to have an opening angle of larger than 0.14 rad and two opposite sign tracks of smaller than 3.00 rad.

In figure 4.2, the invariant mass resolution and signal reconstruction efficiency for $\mu \rightarrow eee$ events is shown requiring a certain minimum number of long tracks in the vertex fit. As expected, the mass resolution improves with a higher number of long tracks due to the better momentum resolution. Putting no restriction on the type of the reconstructed track (labelled *all tracks* in the following) yields a mass resolution of about 2.0 MeV whereas requesting only long tracks improves the resolution to about 1.2 MeV.

The improved mass resolution goes along with a decrease in reconstruction efficiency. The reconstruction efficiency for long tracks is simply smaller as the recurling particle can get lost in the non-sensitive gap between central detector and recurl stations as well as it can even overshoot the recurl station in the case of large longitudinal momentum. For vertices from all tracks, the geometric acceptance is 43 % and the reconstruction efficiency, i. e. the ratio of the number of reconstructed to generated signal events, is about 36 %. When three long tracks are requested, the geometric acceptance reduces to 26 % and the reconstruction efficiency to 20 % (for details see section 4.3).

In the later experiment, the events have to pass the online event selection before being reconstructed offline. The efficiency for signal decays of the preceding online vertex selection relative to the offline vertex reconstruction is assumed to be 95 % according to [157]. The same efficiency is adopted for internal conversion background as anyhow only events similar to $\mu \rightarrow eee$ signal decays are considered. In the following, slightly tighter cuts than in [157] will be defined in the offline vertex reconstruction so that this efficiency can be regarded as a conservative estimate.

(a) Resolution of the reconstructed invariant mass m_{eee} .

(b) Signal reconstruction efficiency.

Figure 4.2: Resolution of m_{eee} estimated as root mean square from reconstructed $\mu \rightarrow eee$ signal events and signal reconstruction efficiency of the vertex fit. No, one, two, or three recurling tracks are requested in the vertex fit. Cuts as outlined in table 4.1 are applied, but the mass cut is loosened to $95\text{MeV} \leq m_{eee} \leq 115\text{MeV}$.

4.2 BACKGROUND STUDY

A very effective background suppression is essential to searches for rare events like $\mu \rightarrow eee$. Only in absence of background events, the targeted sensitivity can be achieved. Cuts on reconstructed variables are derived with the aim to suppress events from $\mu \rightarrow eee\nu$ while keeping the efficiency for signal events as high as possible. The signal efficiency is determined for phase space distributed $\mu \rightarrow eee$ events. Accidental background is estimated elsewhere [136]. It is expected that it can be efficiently suppressed in phase I (see also section 4.1).

For the sake of simplicity, the following considerations are restricted to reconstructed vertices from all tracks or three long tracks and further combinations of different numbers of short and long tracks are not considered.

4.2.1 Pre-Selection

The following results are quoted for a certain number of muon stops. Therefore, in the efficiency analysis only events with a true vertex within a cylinder of radius 19 mm and length 100 mm around the target are considered. This volume is referred to as *target region* in the following discussions. This rejects muon decays in flight in the beam pipe.

All further selections are based on reconstructed variables. Some selections on the relative kinematics between pairs of tracks are already applied in the vertex fit (see section 4.1.2 for details). Furthermore, only events with a reconstructed invariant mass of the three electrons in a loosely defined signal region of $95\text{MeV} \leq m_{eee} \leq 115\text{MeV}$ are considered.

4.2.2 Selection and Cuts

There is a couple of variables that can be used to separate signal events from background. Constraints can be placed on the kinematics, the vertex, and the timing. In the following, cuts on the χ^2 value of the vertex fit, the distance d of the reconstructed vertex to the target surface, and on the total momentum p_{eee} and invariant mass m_{eee} of the three particle system are studied. Timing cuts are studied in [136, 154] and are only briefly presented.

In order to determine the efficiency, the number of generated and reconstructed signal and background events has to be determined. In some cases, the same event is reconstructed multiple times as un-

der certain conditions multiple tracks per particle are reconstructed. This happens for instance for recurling particles that cross the vertex layers again. Each turn is reconstructed as a separate track. Multiple tracks are also reconstructed for particles that cross the overlap of two adjacent pixel modules. Multiple vertex reconstructions then arise by replacing one or several tracks by secondary reconstructed tracks of the same particle.

Such events have to be taken out when determining the efficiency. Therefore, for a true reconstructed signal or background event, it is required that all tracks are associated with an individual true decay particle from the same $\mu \rightarrow eee$ or $\mu \rightarrow eee\nu\nu$ decay. The decay vertex has to be located in the target region. Each track is required to stem from the first loop of the particle and has to be unique, i. e. the track has not been used in a previous true event.

In the later experiment, the rate of multiple reconstructed events should be kept small. There are many measures to reduce it. Apart from selections on the opening angles and momentum differences as discussed above, timing information can be used to identify looping particles and to differentiate electrons and positrons. Furthermore, only a single event with similar kinematics is expected for a select time slice and volume. Therefore, it is possible to select only one event, e. g. the event with the best vertex fit χ^2 , in the case of multiple similar events without reducing the reconstruction efficiency. It is expected, that the multiple reconstructions are effectively reduced in the later experiment, therefore they are not particularly considered in the following.

CUT ON THE χ^2 OF THE VERTEX FIT Cuts on the χ^2 value of the vertex fit are placed to ensure a certain goodness of the vertex fit. In addition, reconstructed events of accidental background yield in general larger χ^2 values than signal events as there is no common vertex to be reconstructed.

The distribution of χ^2 for $\mu \rightarrow eee$ signal decays and background from $\mu \rightarrow eee\nu\nu$ decays and Bhabha events is shown in figure 4.3. The distributions for signal and internal conversion background are very similar. This is expected as in both cases the decay particles all stem from a common vertex. The distribution is maximum at small χ^2 and constantly decreases towards larger values. The χ^2 distribution for vertices reconstructed from all kind of tracks and from long tracks only do not differ qualitatively for both processes.

For accidental background from Bhabha events, the χ^2 distribution is almost flat as the tracks do not all originate from the same vertex. Signal efficiency and background rejection of internal conversion

decays are shown in figure 4.4 for cuts on χ^2 . As expected, no large background rejection for internal conversion decays is achieved in regions with a reasonable signal efficiency. The background rejection is larger for events with only long tracks compared to events without a selection on short or long tracks but of course selecting only long tracks also has a significant impact on the signal efficiency. As can be deduced from the χ^2 distribution of Bhabha events, the χ^2 cut needs to be optimized with regard to accidental background. Here, the same cut of $\chi^2 \leq 30$ is chosen as in the study of accidental background in [136]. This cut affects the signal efficiency only slightly.

CUT ON THE TARGET DISTANCE Muon decay vertices are mostly located on the target with the exception of a small amount of decays in flight. When combining tracks from different processes however the reconstructed vertex does not necessarily lie on the target. Thus, cuts on the distance of the reconstructed vertex position to the target are an effective measure to reduce accidental background.




The target distance d is defined as the shortest distance of the reconstructed vertex position to the surface of the target. The sign of d indicates whether the reconstructed vertex lies inside or outside of the target. As this differentiation is not important for background suppression, only cuts on $|d|$ are considered in the following.

The distribution of the target distance for $\mu \rightarrow eee$ signal decays and background from internal conversion decays and Bhabha events is shown in figure 4.5. As signal and internal conversion background both have a common vertex, the target distance distributions for the two are very similar. The distribution is maximum around $|d| = 0$ and rapidly decreases towards larger target distances. The slope is determined by the resolution of the target distance. An RMS of about 0.6 mm is achieved (see figure 4.6).

The reconstructed target distance of Bhabha events also decreases but with a larger slope than in the case of the other two processes. This is expected due to the missing common decay vertex.

The efficiency for signal decays and background rejection for internal conversion decays as a function of cuts on $|d|$ is shown in figure 4.7. As expected, the background rejection for internal conversion events vanishes in regions with adequate signal efficiencies. As in the case of χ^2 , the background rejection is stronger when three long tracks are required.

A suitable cut on $|d|$ can only be defined considering accidental background. Here, a cut is placed at $|d|$ larger than 3 mm in order to

(a) $\mu \rightarrow eee$ signal decays.(b) $\mu \rightarrow eee\nu\nu$ background decays.

(c) Background from Bhabha events (taken from [136]). Events with $\chi^2 < 200$ and $p_{eee} < 20 \text{ MeV}$ are selected. No selection on m_{eee} is in place. All kinds of reconstructed tracks are considered. Note that the x-axis range is changed.

Figure 4.3: Distribution of the χ^2 of the vertex fit for $\mu \rightarrow eee$ signal decays and background from radiative decays with internal conversion $\mu \rightarrow eee\nu\nu$ and Bhabha events. The generated $\mu \rightarrow eee$ sample contains about $7 \cdot 10^6$ events. $\mu \rightarrow eee\nu\nu$ events are generated for a restricted phase space as described in the main text. The sample contains about $7 \cdot 10^7$ weighted events. The sample of Bhabha events contains about $1.1 \cdot 10^9$ weighted events.







Figure 4.4: Signal efficiency and background suppression by cuts on the χ^2 of the vertex fit for $\mu \rightarrow eee$ signal decays and radiative decays with internal conversion $\mu \rightarrow eee\nu$ (IC).

account for the potential to suppress accidental background. The signal efficiency is only mildly affected. Although the cut is not particularly tight, it is considered to be in a reasonable range. In the previous accidental background study [136], this cut is even omitted and still a good suppression of accidental background is achieved.

CUT ON THE TOTAL MOMENTUM In $\mu \rightarrow eee$ signal decays, the total momentum p_{eee} of the three particles vanishes for muon decays at rest. This is neither expected for internal conversion background due to the additional neutrinos nor for accidental combinations. Thus, a cut on p_{eee} can efficiently suppress both types of background.

In figure 4.8, the distribution of p_{eee} is shown for signal decays and internal conversion background. In the case of $\mu \rightarrow eee$ decays, the events follow a Maxwell-Boltzmann like distribution. The distribution has a steeper decrease for vertices from long tracks only than for vertices without this restriction as a result of the improved momentum resolution. For long tracks, the RMS of the difference of reconstructed and generated p_{eee} —the latter is essentially zero except for decays in flight—is around 1.6 MeV when all tracks are used in the vertex fit, and 1.2 MeV when taking only long tracks (see figure 4.9).

The distribution of internal conversion background is broader. With two undetected neutrinos, the three electron system can have total momenta of up to half the muon mass. The cutoff at around 15 MeV

(a) $\mu \rightarrow eee$ signal decays.(b) $\mu \rightarrow eee\nu$ background decays.

(c) Background from Bhabha events (taken from [136]). Events with $\chi^2 < 200$ and $p_{eee} < 20 \text{ MeV}$ are selected. No selection on m_{eee} is in place. All kinds of reconstructed tracks are considered. Note that the x-axis range is changed.

Figure 4.5: Distribution of the distance $|d|$ to the target surface for $\mu \rightarrow eee$ signal decays and background from radiative decays with internal conversion $\mu \rightarrow eee\nu$ and Bhabha events. The generated $\mu \rightarrow eee$ sample contains about $7 \cdot 10^6$ events. $\mu \rightarrow eee\nu$ events are generated for a restricted phase space as described in the main text. The sample contains about $7 \cdot 10^7$ weighted events. The sample of Bhabha events contains about $1.1 \cdot 10^9$ weighted events.




Figure 4.6: Difference of the reconstructed and generated target distance d of the three electron system for $\mu \rightarrow eee$ signal decays.




Figure 4.7: Signal efficiency and background suppression by cuts on the target distance $|d|$ for $\mu \rightarrow eee$ signal decays and radiative decays with internal conversion $\mu \rightarrow eee\nu$ (IC).

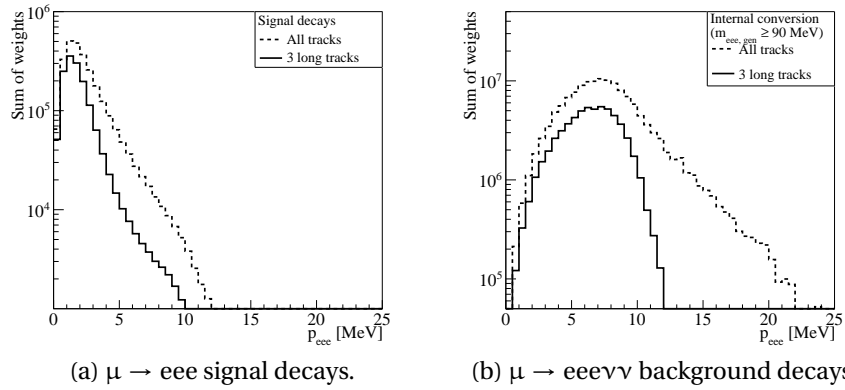


Figure 4.8: Distribution of the total momentum p_{eee} for $\mu \rightarrow eee$ signal decays and radiative decays with internal conversion $\mu \rightarrow eee\nu\nu$. The generated $\mu \rightarrow eee$ sample contains about $7 \cdot 10^6$ events. $\mu \rightarrow eee\nu\nu$ events are generated for a restricted phase space as described in the main text. The sample contains about $7 \cdot 10^7$ weighted events.

that is visible in figure 4.8 results from generating only internal conversion events with an invariant mass of the three charged particles of 90 MeV or larger. Without this restriction, the p_{eee} distribution of internal conversion events would expand further.

The signal efficiency and background rejection for internal conversion decays are shown in figure 4.10. The signal efficiency rises with loosening the p_{eee} cut and is about constant after 10 MeV. The background rejection is of course maximum at tight p_{eee} cuts and decreases with looser cuts. In this case, the background rejection curve itself is not much affected by generating internal conversion events with an invariant three-electron mass of 90 MeV or larger as anyhow only events with a reconstructed m_{eee} of larger 95 MeV are considered. As expected, the background rejection is stronger for events with three long tracks.

As the signal efficiency rises steeply, there is a range of possible p_{eee} cut values at which good signal efficiencies are achieved while having a considerable background rejection. A cut is placed at 8 MeV giving a moderate background rejection with a minor impact on the signal efficiency. As another kinematic cut will be placed on m_{eee} , maintaining a good signal efficiency is favoured over a strong background rejection in this cut.

CUT ON THE INVARIANT MASS The most effective cut in suppressing internal conversion background is on the reconstructed invariant mass m_{eee} of the three electron system. In $\mu \rightarrow eee$ signal




Figure 4.9: Difference of the reconstructed and generated total momentum p_{eee} of the three electron system for $\mu \rightarrow eee$ signal decays.




Figure 4.10: Signal efficiency and background suppression by cuts on the total momentum p_{eee} for $\mu \rightarrow eee$ signal decays and radiative decays with internal conversion $\mu \rightarrow eee\nu\nu$ (IC).

decays, the true m_{eee} is equal to the muon rest mass, but for internal conversion decays it is smaller due to the two additional neutrinos. Figure 4.11 shows the distribution of the reconstructed m_{eee} of $\mu \rightarrow eee$ signal decays and background from internal conversion decays.

As expected, $\mu \rightarrow eee$ signal events have a peak around m_μ . The distribution of events from all tracks and events with three long tracks vary notably. Events with three long tracks fall off steeply at around 107 MeV and have a tail towards m_{eee} smaller than m_μ . The tail is a result from the energy loss of the individual particles in the detector which is on average on the order of 0.1 MeV. As the material amount in the detector is low, the energy loss follows a Landau distribution [191] which has a tail stemming from rare events of single large energy loss by bremsstrahlung or ionization. The energy loss of the individual particles transfers to an energy loss in the three electron system and thus a smaller invariant mass is measured. In the m_{eee} spectrum of all track events, the tail is less pronounced due to the larger invariant mass resolution. This can be deduced from figure 4.12 where the difference of the reconstructed and generated m_{eee} of $\mu \rightarrow eee$ signal events is compared for the two cases. The distribution for three long tracks is narrower and yields an RMS of about 1.2 MeV. For all tracks, the RMS is larger with about 2 MeV.

The m_{eee} distribution of internal conversion events has a steep decline towards the muon rest mass especially for events with three long tracks. The events with all kinds of reconstructed tracks have a tail towards high m_{eee} that even extends over the muon rest mass. This is again a result of the worse momentum resolution of short tracks and thus worse m_{eee} resolution.

Figure 4.13 shows the signal efficiency and internal conversion background rejection for cuts on the reconstructed invariant mass. The reconstruction efficiency is at a high level until the cut approaches the muon rest mass where it quickly drops. This is similar for all track and three long track events but the edge of the first ones is softer. The difference of the two is even more apparent in the background rejection. For three long tracks, the background rejection rises steeply and reaches almost 100% already a few MeV below m_μ . The rise of the background rejection curve of all tracks is flatter. A high background rejection is only achieved for cuts past the muon rest mass.

The cut on the reconstructed invariant mass is set considering only long tracks. In order to suppress internal conversion background efficiently, a rather tight cut is placed at 103.5 MeV. The same cut is also used for the all track events.

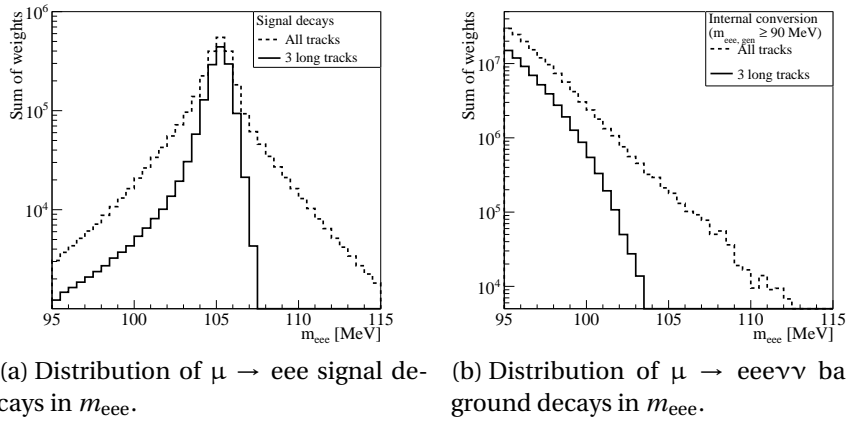


Figure 4.11: Distribution of the invariant mass m_{eee} of the three electron system for $\mu \rightarrow eee$ signal decays and radiative decays with internal conversion $\mu \rightarrow eee\nu\nu$. The generated $\mu \rightarrow eee$ sample contains about $7 \cdot 10^6$ events. $\mu \rightarrow eee\nu\nu$ events are generated for a restricted phase space as described in the main text. The sample contains about $7 \cdot 10^7$ weighted events.

An upper invariant mass cut has not much background rejection power. Internal conversion events cannot yield invariant masses above $m_\mu - 3m_e$, and simulations of accidental background have so far not yielded a single event above a reconstructed m_{eee} of 102 MeV. A loose cut is set to $m_{eee} \leq 115$ MeV in order to remove potential misreconstruction but not to harm the reconstruction efficiency.

TIMING CUTS The relative timing between the particles is an effective measure to suppress accidental background. Tracks from a signal event are expected to appear coincident in time. This is not expected for accidental combinations as the three tracks are at least partially uncorrelated. In a dedicated analysis of timing suppression [136, 154], a selection is defined that suppresses background events with two correlated tracks, e. g. with an electron-positron pair from Bhabha scattering, by a factor of 108 [136, 154]. Background events with three uncorrelated tracks are even stronger suppressed by a factor of 6400. At the same time, the efficiency for signal decays is mildly reduced to 98%. The same holds for background from internal conversion decays which also feature three coincident tracks. It is expected that the timing suppression factorizes from kinematic and vertex suppression.




Figure 4.12: Difference of the reconstructed and generated invariant mass m_{eee} of the three electron system for $\mu \rightarrow eee$ signal decays.




Figure 4.13: Signal efficiency and background suppression by cuts on the invariant mass m_{eee} of the three electron system for $\mu \rightarrow eee$ signal decays and radiative decays with internal conversion $\mu \rightarrow eee\nu$ (IC). Note that the axis range for the background rejection has changed compared to previous figures.


Cut
$0 \leq \chi^2 \leq 30$
$ d \leq 3 \text{ mm}$
$p_{eee} \leq 8 \text{ MeV}$
$103.5 \text{ MeV} \leq m_{eee} \leq 115 \text{ MeV}$

Table 4.1: Data selection cuts as applied in the sensitivity studies to $\mu \rightarrow eee$ decays.


COMBINING THE CUTS The cuts for the $\mu \rightarrow eee$ analysis are summarized in table 4.1. Figure 4.14 shows the reconstructed invariant mass spectrum for $\mu \rightarrow eee$ signal and background from muon decays with internal conversion before and after applying cuts. Except for the loss of signal events at lower masses, the signal distribution is only marginally affected by the cuts. The distribution for vertices from three long tracks is narrow and centred around the muon rest mass. For vertices from all tracks, there is a tail towards large m_{eee} due to the larger mass resolution. The advantage of the improved mass resolution by using only long tracks becomes apparent when the m_{eee} distributions of internal conversion background are compared. The resulting background is about two orders of magnitude smaller than in the case of vertex fits from all kinds of tracks.

In table 4.2, the reconstruction efficiencies for $\mu \rightarrow eee$ signal and internal conversion background events are listed before and after applying the cuts. A remarkable reduction of internal conversion background by four orders of magnitude is achieved. The resulting efficiency of $1.22 \cdot 10^{-6}$ corresponds to 0.44 expected internal conversion events during the estimated run time of the phase I *Mu3e* experiment of 300 days at $10^8 \mu/s$. In contrast, for events with all tracks in the vertex fit about 39 events from internal conversion background are expected. A background-free operation is thus only possible when only long tracks are considered in the vertex fit. The reconstruction efficiency for $\mu \rightarrow eee$ signal decays is at 17.1 % in this case.


In figure 4.15, the mass spectrum of $\mu \rightarrow eee$ signal decays and internal conversion background is shown for three long tracks vertices after cuts. For the signal decays, branching fractions between 10^{-15} to 10^{-12} are assumed. At least at larger branching fractions, the signal is clearly distinguishable from the background. The sensitivity is derived in the following.




(a) Distribution of $\mu \rightarrow eee$ signal decays in m_{eee} for vertex fits with all kinds of reconstructed tracks.



(b) Distribution of $\mu \rightarrow eee$ signal decays in m_{eee} for vertex fits with long tracks only.



(c) Distribution of $\mu \rightarrow eee\nu\nu$ background decays in m_{eee} for vertex fits with all kinds of reconstructed tracks.



(d) Distribution of $\mu \rightarrow eee\nu\nu$ background decays in m_{eee} for vertex fits with long tracks only.

Figure 4.14: Reconstructed m_{eee} spectra for $\mu \rightarrow eee$ signal decays and radiative decays with internal conversion $\mu \rightarrow eee\nu\nu$ before and after applying cuts. On the left hand side, vertex fits with all kinds of reconstructed tracks are shown, and on the right hand side the vertex is fitted with only long reconstructed tracks. The $\mu \rightarrow eee$ sample contains about $7 \cdot 10^6$ events. $\mu \rightarrow eee\nu\nu$ events are generated for a restricted phase space as described in the main text. The sample contains about $7 \cdot 10^7$ weighted events.

Decay	Reconstruction	$\epsilon_{\text{No cuts}}$	ϵ_{Cuts}
Signal $\mu \rightarrow eee$	All tracks	0.364	0.286
	3 long tracks	0.198	0.171
Background $\mu \rightarrow eee\nu\bar{\nu}$	All tracks	$3.89 \cdot 10^{-2}$	$1.08 \cdot 10^{-4}$
	3 long tracks	$1.66 \cdot 10^{-2}$	$1.22 \cdot 10^{-6}$

Table 4.2: Reconstruction efficiency of $\mu \rightarrow eee$ signal events and background from $\mu \rightarrow eee\nu\bar{\nu}$ decays before and after applying cuts. An online selection efficiency of 95 % is considered. The efficiencies after cuts also take a timing efficiency of 98 % into account. $\mu \rightarrow eee\nu\bar{\nu}$ events are generated for a phase space as described in the main text. The relative statistical uncertainty is on the 10^{-3} level.




Figure 4.15: Spectra of the reconstructed invariant three particle mass of $\mu \rightarrow eee$ signal decays and background from internal conversion in phase I of the *Mu3e* experiment. A muon stopping rate of $1 \cdot 10^8 \mu/s$ is assumed with a total number of $2.592 \cdot 10^{15}$ muon stops corresponding to 300 days of data taking. Cuts as outlined in table 4.1 are applied. The cut on the reconstructed invariant mass is indicated by the green line.

	Events	
Decays in target region	6969034	
	All tracks	3 long tracks
All particles in acceptance	3018696	1825485
All tracks reconstructed	2713805	1469410
Reconstructed vertices	2711068	1468294
Reconstructed vertices, cuts	2380214	1315767
Reconstructed vertices, cuts, online selection and timing	2215979	1224979

Table 4.3: The generated signal sample. Given are the number of events in the acceptance as well as after track and vertex reconstruction for $1 \cdot 10^7$ generated frames. The efficiency of the online selection is assumed with 95 % and the efficiency of the timing cuts with 98 %.

4.3 KINEMATICS OF THE SIGNAL DECAY

The kinematics, acceptance and reconstruction efficiency of signal decays are discussed for $\mu \rightarrow eee$ decays that are generated evenly distributed in phase space. The numbers of events left after each step in simulation and reconstruction are summarized in table 4.3. In total, about $7 \cdot 10^6$ signal events are generated.

The kinematics within the decay plane of the three particles in the centre of mass system of the muon can be fully displayed by a Dalitz plot of the invariant masses square of the two possible electron-positron combinations. A short description of Dalitz plots is given in section B.2 in the appendix. The orientation of the decay plane with respect to the spin of the muon is shown in $\cos(\tilde{\vartheta})$ and $\tilde{\varphi}$. The definition of the angles can be found in figure 4.16. In essence, $\tilde{\vartheta}$ denotes the angle between the momentum of the electron and the orientation of the muon spin. The angle $\tilde{\varphi}$ is the angle between the plane of the electron momentum and muon spin and the $\mu \rightarrow eee$ decay plane. The same parametrization is used in the differential branching fraction for $\mu \rightarrow eee$ decays in effective theories (see equation B.3 in the appendix).

ACCEPTANCE A Dalitz plot of the invariant mass square m_{ee}^2 of the two possible electron-positron combinations of generated $\mu \rightarrow eee$ signal events as well as the angular distribution are shown





Figure 4.16: Definition of the angles $\tilde{\vartheta}$ ranging from 0 to π and $\tilde{\varphi}$ ranging from 0 to 2π as in equation B.3 in the appendix. The decay positrons with momenta \vec{p}_1 and \vec{p}_2 and the electron with momentum \vec{p}_3 lie in plane I. Plane II is spanned by the electron momentum and the muon polarization vector \vec{P} . The figure is adapted from [46, 62].


in figures 4.17a and 4.17b. As expected, the events are evenly distributed over the accessible phase space.

A significant number of signal decays remains unobserved in the detector because of geometric acceptance. Particles emitted under a shallow angle with respect to the beam pipe are not detected as well as particles with transverse momentum p_T of lower than 10 MeV. The latter do not cross all four central pixel layers, and thus their trajectories cannot be reconstructed. It suffices to miss a single particle in order to fail to reconstruct the event. Figures 4.17c and 4.17d illustrate where events are lost. The corners of the triangular Dalitz plot are not populated. In these regions of phase space, one of the three particles has no or very small kinetic energy, and thus not sufficient transverse momentum to be detected. On the other hand, the regions around the centre of each of the edges are favoured. Here, one particle carries about half of the available energy with the rest being about equally shared between the other two. In this configuration, it is very likely that if one of the particles is detected also the others can be detected as the particles with lower energy are emitted anti-parallel to the particle with the higher energy.


The angular distribution shows an acceptance loss for $|\cos(\tilde{\vartheta})| \approx 1$.




(a) Dalitz plot of the generated sample.




(b) Angular distribution of the generated sample.




(c) Dalitz plot of generated events in the acceptance.



(d) Angular distribution of generated events in the acceptance.



(e) Dalitz plot of generated events with three reconstructable long tracks.



(f) Angular distribution of generated events with three reconstructable long tracks.

Figure 4.17: Dalitz plots of the truth invariant mass of the e^+e^- system and the truth angular distribution of phase space distributed events.


In these cases, the electron is emitted under a small angle with respect to the beam line. In cases with $\tilde{\varphi} \approx 0$ or π , the beam axis lies in the decay plane and thus it is more likely that one particle escapes through the beam pipe than in cases where the angle between the decay plane and the beam axis is larger. The angular distributions appear slightly asymmetric in shape because of the elliptical muon beam profile and an asymmetric stopping distribution on the target. Requiring all tracks to be reconstructable as long tracks adds more structures to the m_{ee}^2 and angular distributions (see figures 4.17e and 4.17f). There are two momentum thresholds for long tracks. The first is when the momentum is large enough to reach the gap between central and recurl station and the second is when the momentum is large enough so that the particle can reach further than the recurl stations. This leads to the steps visible in the Dalitz plot. The efficiency is largest in the centre of the plot where all particles have about the same momentum. Likewise, the gap can also be observed in the angular distribution. The elliptically shaped areas in figure 4.17d have got a second grading. The centre is populated by events where all particles recurl in the central station. In the surrounding area, one or more of the particles reaches the recurl stations. In addition, the gap also causes inefficiencies at $|\cos(\tilde{\vartheta})| \gtrsim 0.2$.

About 43 % of the signal events lie within the geometric acceptance. The fraction of events with all three particles being reconstructable as long tracks is 26 %.


TRACK RECONSTRUCTION Dalitz plots and angular distributions for events in which the tracks of all particles are reconstructed are shown in figure 4.18. The Dalitz plot of reconstructed tracks does not differ much from the one of generated tracks in the acceptance, even not when three long tracks are required.

In the angular distributions, an inefficiency at $\cos(\tilde{\vartheta}) \approx 0$ is observed. Here, the electron is emitted perpendicular to the beam axis and will perform multiple turns that are very closely spaced. It can occur that tracks built from hits of different turns yield a better χ^2 value in the track reconstruction and are thus kept instead of the track from the first turn. As the truth-matching applied in this study requires the tracks to stem from the first turn, an inefficiency appears in the angular distributions. The same is true for the positron tracks. Therefore, inefficiencies are observed at $\tilde{\varphi} = \frac{\pi}{2}$ and $\frac{3\pi}{2}$.


The efficiency to reconstruct all tracks from a signal event is 39 %,




(a) Dalitz plot of events in which all particles are reconstructed.



(b) Angular distribution of events in which all particles are reconstructed.




(c) Dalitz plot of events in which all particles are reconstructed as long tracks.




(d) Angular distribution of generated events in which all particles are reconstructed as long tracks.


Figure 4.18: Dalitz plots of the truth invariant mass of the e^+e^- system and the truth angular distribution of phase space distributed events after track reconstruction.




(a) Dalitz plot of events with reconstructed vertex



(b) Angular distribution of events with reconstructed vertex



(c) Dalitz plot of events with reconstructed vertex from three long tracks.




(d) Angular distribution of generated events with reconstructed vertex from three long tracks.

Figure 4.19: Dalitz plots of the truth invariant mass of the e^+e^- system and the truth angular distribution of phase space distributed events after vertex fit.


and 21 % if all of the tracks are reconstructed as long tracks.

VERTEX RECONSTRUCTION As can be seen in figures 4.19, the vertex fit does not affect a specific phase space region. Inefficiencies of the vertex fit stem from non-converging fits and configurations in which an intermediate reconstructed vertex position lies outside the target region, though the success rate of the vertex fit is very good with 99.9 %. In order to suppress background, cuts as derived in the previous section are applied (see table 4.2). Again, no specific phase space region is affected (see figure 4.20).


Figure 4.21 shows the Dalitz plot of the reconstructed invariant mass m_{ee} and the reconstructed angular distribution. Due to the finite




(a) Dalitz plot of events with reconstructed vertex.



(b) Angular distribution of events with reconstructed vertex.



(c) Dalitz plot of events with reconstructed vertex from three long tracks.



(d) Angular distribution of generated events with reconstructed vertex from three long tracks.

Figure 4.20: Dalitz plots of the truth invariant mass of the e^+e^- system and the truth angular distribution of phase space distributed events after vertex fit and cuts.

resolution, the features of the distributions of the truth variables are slightly blurred out but the characteristic acceptance of the detector is very well visible. In the Dalitz plot, there are a few entries that extend over the actually accessible phase space. These result from the m_{ee} resolution on the one hand, and from decays in flight on the other hand. In the case of vertex fits from three long tracks, fewer entries are found outside the accessible phase space due to the improved mass resolution.

The inefficiency of the online selection as well as by cuts on the relative timing are not studied here. An online selection efficiency of 95 % [157] and a timing cut efficiency of 98 % [136, 154] is assumed. After applying the cuts, the efficiency to reconstruct the vertex of a signal event is thus 32 %, and 18 % if three long tracks are required. An overview of the efficiency of the individual steps of reconstruction is given in table 4.4.

4.4 SENSITIVITY TO EFFECTIVE OPERATORS

4.4.1 *Effective Operators*


In the following, the sensitivity to $\mu \rightarrow eee$ signal decays mediated via different effective operators is studied. The energy scale μ of the investigated process is the muon rest mass m_μ and thus well below the scale Λ of the effective theory. At the scale μ , no mixing effects with operators involving quark or tau fermion fields are expected. Keeping only relevant operators, the SM Lagrangian can be extended by an effective Lagrangian [60, 61] mediating $\mu \rightarrow eee$ decays²

$$\mathcal{L}_{\text{eff}} = \frac{1}{\Lambda^2} [C_D^L \mathcal{O}_D^L + C_D^R \mathcal{O}_D^R + C_S^{LL} \mathcal{O}_S^{LL} + C_S^{RR} \mathcal{O}_S^{RR} + C_V^{LL} \mathcal{O}_V^{LL} + C_V^{RR} \mathcal{O}_V^{RR} + C_V^{LR} \mathcal{O}_V^{LR} + C_V^{RL} \mathcal{O}_V^{RL}] + \text{h.c.} \quad (4.1)$$


which is valid at a scale $m_W \geq \Lambda \gg m_b$. The operators \mathcal{O} are listed in table 4.5. The coupling is encoded in the dimensionless Wilson coefficients C and the scale Λ . The operators can be classified into *dipole* operators \mathcal{O}_D involving (virtual) photons and scalar or vector operators \mathcal{O}_S and \mathcal{O}_V corresponding to contact interactions of four leptons (*4-fermion*).

For the study of effective operators, a weight is applied on an event-


² The operators explicitly shown mediate $\mu^- \rightarrow e^- e^+ e^-$. The decay $\mu^+ \rightarrow e^+ e^- e^+$ investigated in *Mu3e* is mediated by the operators of the hermitian conjugate part of \mathcal{L}_{eff} .




(a) Reconstructed Dalitz plots of phase space distributed $\mu \rightarrow eee$ signal events.



(b) Reconstructed angular distribution of phase space distributed $\mu \rightarrow eee$ signal events.



(c) Reconstructed Dalitz plots of phase space distributed $\mu \rightarrow eee$ signal events from vertex fits with three long tracks.



(d) Reconstructed angular distribution of phase space distributed $\mu \rightarrow eee$ signal events from vertex fits with three long tracks.

Figure 4.21: Dalitz plots of the reconstructed invariant mass of the e^+e^- system and the reconstructed angular distribution of phase space distributed $\mu \rightarrow eee$ signal events after vertex fit and cuts.

Phase space, all tracks	ϵ [%]
Events in acceptance	43.3
Events after track reconstruction	38.9
relative to events in acceptance	89.9
Events with reconstructed vertex	37.0
relative to events in acceptance	85.3
relative to events after track reco.	94.9
Events with reconstructed vertex after cuts	29.2
relative to events in acceptance	67.3
relative to events after track reco.	74.9
relative to events with reco. vertex before cuts	74.9
Events with reco. vertex after cuts incl. online sel. & timing	28.6
Phase space, only long tracks	ϵ [%]
Events in acceptance	26.2
relative to all events in acceptance	60.5
Events after track reconstruction	21.1
relative to events in acceptance	80.5
Events with reconstructed vertex	20.0
relative to events in acceptance	76.4
relative to events after track reco.	94.9
Events with reconstructed vertex after cuts	17.5
relative to events in acceptance	66.8
relative to events after track reco.	83.0
relative to events with reco. vertex before cuts	83.0
Events with reco. vertex after cuts incl. online sel. & timing	17.1

Table 4.4: Efficiency after reconstruction and vertex fit for phase space distributed signal events. An online selection efficiency of 95 % and an efficiency of the timing cuts of 98 % is assumed.

Interaction		Operator	
Dipole		\mathcal{O}_D^R	$em_\mu (\overline{\mu}_R \sigma^{\mu\nu} e_L) F_{\mu\nu}$
		\mathcal{O}_D^L	$em_\mu (\overline{\mu}_L \sigma^{\mu\nu} e_R) F_{\mu\nu}$
4-fermion	Scalar	\mathcal{O}_S^{RR}	$(\overline{\mu}_R e_L) (\overline{e}_R e_L)$
		\mathcal{O}_S^{LL}	$(\overline{\mu}_L e_R) (\overline{e}_L e_R)$
4-fermion	Vector	\mathcal{O}_V^{RR}	$(\overline{\mu}_R \gamma^\mu e_R) (\overline{e}_R \gamma_\mu e_R)$
		\mathcal{O}_V^{LL}	$(\overline{\mu}_L \gamma^\mu e_L) (\overline{e}_L \gamma_\mu e_L)$
		\mathcal{O}_V^{RL}	$(\overline{\mu}_R \gamma^\mu e_R) (\overline{e}_L \gamma_\mu e_L)$
		\mathcal{O}_V^{LR}	$(\overline{\mu}_L \gamma^\mu e_L) (\overline{e}_R \gamma_\mu e_R)$

Table 4.5: Effective operators up to dimension six that mediate $\mu \rightarrow eee$ decays. $\mu_{R,L}$ and $e_{R,L}$ are the fermion fields in right- and left-handed projection. $F_{\mu\nu}$ is the photon field strength tensor, and e the electron charge.

by-event base in the analysis step. The weight corresponds to the respective differential branching fraction as given by equation B.3 [46, 62] (see section B.1 in the appendix for more details). The sensitivity is estimated for each operator individually meaning that the coefficients of all other operators are set to zero. This approach is legitimate as long as the energy scale is kept at about m_μ .

For simplicity, for each operator only one chirality is shown. The corresponding distributions with the opposite chirality can be found in section B.2 in the appendix. Diagrams of reconstructed $\mu \rightarrow eee$ events are shown for the case that all tracks are reconstructed as long tracks as only in this case a sufficient background suppression is achieved. Cuts as listed in table 4.1 are applied. A summary of the acceptances and efficiencies for the various operators is given in table 4.6.

4.4.2 Dipole Operator

The Dalitz plot and angular distribution for the dipole operator \mathcal{O}_D^R are shown in figure 4.22.

In the Dalitz plot of the generated events, most events are located in regions where either of the electron-positron pairs has zero or almost zero invariant mass. This is expected from the virtual photon. The generated angular distribution has no evident peculiarities. As expected for this operator, the positron with the higher energy is

preferentially emitted oppositely to the direction of the muon spin and the other positron as well as the electron in the same direction (see figure 4.23).

As the majority of the events are generated with nearly zero invariant mass, a large fraction of the events is located in the corners of the triangular Dalitz plot, and is thus not within the acceptance of the detector. This affects the efficiency. The fraction of events in the acceptance for three long tracks reduces from 26.2 % in the case of phase space distributed events to 18.9 % in the case of the dipole operator.

Also, the performance of the track and vertex reconstruction is affected. For phase space distributed events, 65 % of the events in the acceptance are reconstructed and pass the cuts. For the dipole operator, this fraction reduces to 56 %. In events with a small invariant e^-e^+ mass, the opening angle of the pair is very small and the two tracks are almost parallel in the vicinity of the decay vertex. The third particle is emitted anti-parallel to the direction of the electron-positron pair. This is an unfavourable configuration for the vertex fit as all tracks intersect under small angles. The position of the fitted vertex can be easily shifted in the direction parallel to the momentum of the electron-positron pair without large changes of the χ^2 value. Therefore, the resolution of the target distance deteriorates (see figure 4.24). As a result, events with small m_{ee} are stronger affected by a cut on the target distance.

The efficiency to reconstruct a $\mu \rightarrow eee$ event mediated by a dipole operator of type \mathcal{O}_D^R amounts to 10.6 %.

The same considerations are valid for the operator \mathcal{O}_D^L except that the angular distribution is inverted due to opposite chiralities of the involved spinors (see figure B.2). The efficiency is 10.5 %.

4.4.3 4-Fermion Operators

Figure 4.25 shows the Dalitz plot and angular distribution for the scalar 4-fermion operator \mathcal{O}_S^{RR} and the vector 4-fermion operator \mathcal{O}_V^{RR} . In equation B.3 for the differential branching fraction of $\mu \rightarrow eee$ decays in effective theories, the same kinematic function is used for both operators. Therefore, the two operators are indistinguishable.

In the Dalitz plot of the generated events, a broad diagonal band can be identified. This is a typical signature of a resonance in the two particle system that is not explicitly shown. In this case however, there is not a real resonance in the positron-positron system. The truth invariant mass of the two positrons has a rather broad distribution

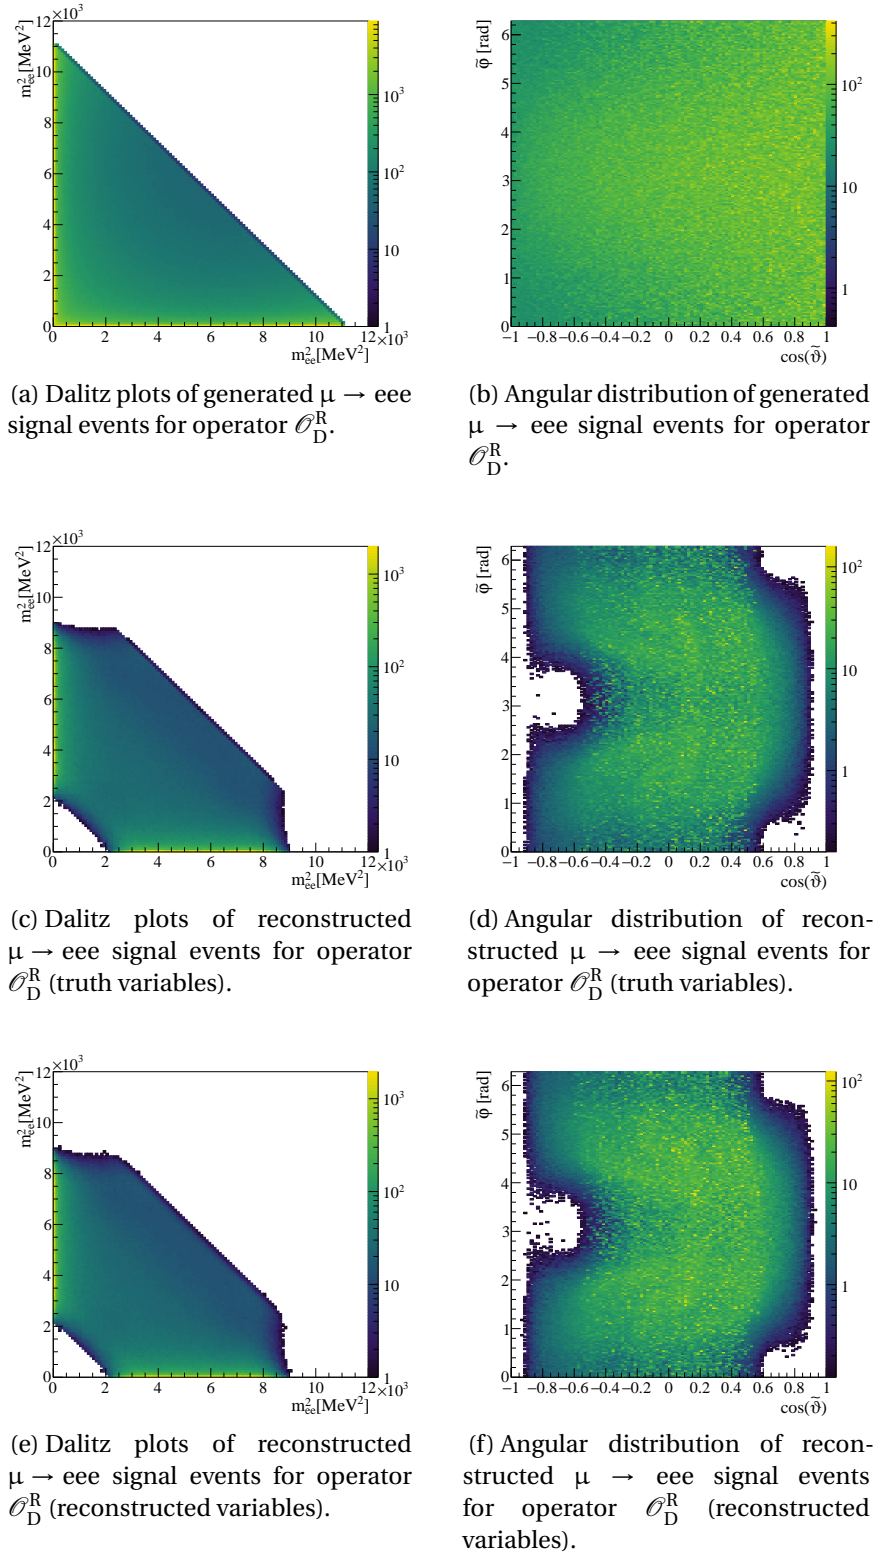


Figure 4.22: Dalitz plots of the truth and reconstructed invariant mass of the e^+e^- system and the truth and reconstructed angular distribution of $\mu \rightarrow eee$ signal events mediated via the dipole operator \mathcal{O}_D^R on generator level and after vertex reconstruction. Similar diagrams for \mathcal{O}_D^L are shown in figure B.2 in the appendix. Note the logarithmic scale.




Figure 4.23: Orientation of the momentum of electrons and positrons in $\mu \rightarrow eee$ decays mediated via the operator \mathcal{O}_D^R with respect to the muon spin direction. The positron with the higher energy (high E) and lower energy (low E) are differentiated.




Figure 4.24: Resolution of the reconstructed distance to the target d for $\mu \rightarrow eee$ signal events with different invariant masses m_{ee} of the electron-positron pair. The resolution is estimated as root mean square of the difference of the reconstructed and generated target distance.

(see figure 4.26). This is expected, as the underlying scalar or vector particle of the effective operator is typically very heavy and thus produced off-mass shell.

The electron is preferentially emitted parallel to the muon spin direction in accordance to the chiral couplings of the operators (see also figure 4.27).

The broad band in the Dalitz plot is in a region with high acceptance. About 29.0 % of the generated events are in the acceptance for three long tracks. This exceeds the fraction of accepted events in the case of phase space distributed events which is about 26.2 %. The track and vertex reconstruction performs equally well for these 4-fermion operators than for phase space distributed events. About 66 % of the events in the acceptance can be reconstructed and pass the cuts. In total, the efficiency to reconstruct an event mediated by a scalar operator of type $\mathcal{O}_S^{\text{RR}}$ or a vector operator of type $\mathcal{O}_V^{\text{RR}}$ amounts to 19.2 %. The same efficiency is achieved for the operators $\mathcal{O}_S^{\text{LL}}$ and $\mathcal{O}_V^{\text{LL}}$ (see also figure B.3).

The Dalitz plot and angular distribution for the 4-fermion vector operator $\mathcal{O}_V^{\text{RL}}$ are shown in figure 4.28.

In the Dalitz plot of the generated events, the centre of the diagonal edge is preferred. In this phase space region, the electron energy is maximum with about half the muon rest mass.


The angular distribution is dominated by the lower energy positron. It is emitted preferentially parallel to the muon spin direction (see also figure 4.29).

As observed in section 4.3, the acceptance is high for events in which one of the particles has about the maximum momentum. Therefore, also in the case of the vector operator $\mathcal{O}_V^{\text{RL}}$ a high acceptance is achieved. About 29.1 % of the generated events are in the acceptance for three long tracks. The reconstruction efficiency for events in the acceptance is comparable to phase space distributed events with 66 %. In total, a $\mu \rightarrow eee$ event mediated via $\mathcal{O}_V^{\text{RL}}$ is reconstructed with an efficiency of 19.2 %. The same holds for the effective operator $\mathcal{O}_V^{\text{LR}}$ (see also figure B.4).


4.4.4 Selection Cuts

In section 4.2.2, cuts are defined in regard to a phase space distributed $\mu \rightarrow eee$ sample. As the various operators change the kinematics of $\mu \rightarrow eee$ events, also the reconstruction efficiency is in principle differently affected by the selection cuts.


Figure 4.30 shows the resulting signal efficiency if one of the cuts is




(a) Dalitz plots of generated $\mu \rightarrow eee$ signal events for operator \mathcal{O}_S^{RR} or \mathcal{O}_V^{RR} .




(b) Angular distribution of generated $\mu \rightarrow eee$ signal events for operator \mathcal{O}_S^{RR} or \mathcal{O}_V^{RR} .




(c) Dalitz plots of reconstructed $\mu \rightarrow eee$ signal events for operator \mathcal{O}_S^{RR} or \mathcal{O}_V^{RR} (truth variables).



(d) Angular distribution of reconstructed $\mu \rightarrow eee$ signal events for operator \mathcal{O}_S^{RR} or \mathcal{O}_V^{RR} (truth variables).



(e) Dalitz plots of reconstructed $\mu \rightarrow eee$ signal events for operator \mathcal{O}_S^{RR} or \mathcal{O}_V^{RR} (reconstructed variables).



(f) Angular distribution of reconstructed $\mu \rightarrow eee$ signal events for operator \mathcal{O}_S^{RR} or \mathcal{O}_V^{RR} (reconstructed variables).

Figure 4.25: Dalitz plots of the truth and reconstructed invariant mass of the e^+e^- system and the truth and reconstructed angular distribution of $\mu \rightarrow eee$ signal events mediated via the scalar 4-fermion operator \mathcal{O}_S^{RR} or the vector 4-fermion operator \mathcal{O}_V^{RR} on generator level and after vertex reconstruction. Similar diagrams for \mathcal{O}_S^{LL} and \mathcal{O}_V^{LL} are shown in figure B.3 in the appendix.




Figure 4.26: Truth invariant mass of both electron-positron pairs as well as of the positron-positron pair in the case of $\mu \rightarrow eee$ signal decays mediated via the operator \mathcal{O}_S^{RR} or \mathcal{O}_V^{RR} .







Figure 4.27: Orientation of the momentum of electrons and positrons in $\mu \rightarrow eee$ decays mediated via the operator \mathcal{O}_S^{RR} or \mathcal{O}_V^{RR} with respect to the muon spin direction. The positron with the higher energy (high E) and lower energy (low E) are differentiated.




(a) Dalitz plots of generated $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{RL}}$.




(b) Angular distribution of generated $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{RL}}$.




(c) Dalitz plots of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{RL}}$ (truth variables).



(d) Angular distribution of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{RL}}$ (truth variables).



(e) Dalitz plots of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{RL}}$ (reconstructed variables).



(f) Angular distribution of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{RL}}$ (reconstructed variables).

Figure 4.28: Dalitz plots of the truth and reconstructed invariant mass of the e^+e^- system and the truth and reconstructed angular distribution of $\mu \rightarrow eee$ signal events mediated via the vector 4-fermion operator $\mathcal{O}_V^{\text{RL}}$ on generator level and after vertex reconstruction. Similar diagrams for $\mathcal{O}_V^{\text{LR}}$ are shown in figure B.4 in the appendix.




Figure 4.29: Orientation of the momentum of electrons and positrons in $\mu \rightarrow eee$ decays mediated via the operator \mathcal{O}_V^{RL} with respect to the muon spin direction. The positron with the higher energy (high E) and lower energy (low E) are differentiated.

Operator		a [%]	ϵ [%]
Phase space		26.2	17.1
Dipole	\mathcal{O}_D^R	18.9	10.6
	\mathcal{O}_D^L	18.9	10.5
4-fermion	Scalar \mathcal{O}_S^{RR}	29.0	19.2
	\mathcal{O}_S^{LL}	29.1	19.2
4-fermion	Vector \mathcal{O}_V^{RR}	29.0	19.2
	\mathcal{O}_V^{LL}	29.1	19.2
	\mathcal{O}_V^{RL}	29.1	19.2
	\mathcal{O}_V^{LR}	29.1	19.2

Table 4.6: Acceptance a and efficiency ϵ for $\mu \rightarrow eee$ signal decays mediated via different effective operators. Three long tracks are required. The relative statistical uncertainties are on the 10^{-3} level.

varied for the investigated effective operators. All other cuts are kept at values as given in table 4.1.

All 4-fermion operators are likewise affected, as well as the two dipole operators basically yield the same efficiencies for every cut. The difference in efficiency between 4-fermion and dipole operators as well as to the phase space distributed decays is the same as observed before. For cuts on the χ^2 of the vertex fit and on the total reconstructed momentum p_{eee} , all curves behave similarly. But in the case of cuts on the target distance d , the signal efficiency of dipole operator mediated decays has a flatter increase towards looser cuts than 4-fermion operator mediated events. As discussed above, the resolution of the reconstructed vertex position is deteriorated in events in which one pair of tracks has a small opening angle. As this is the preferred kinematics of dipole mediated decays, cuts on the target distance generally affect decays mediated by dipole operators stronger than those mediated by 4-fermion operators.

In figure 4.31, the distribution of the reconstructed invariant mass m_{eee} is shown for the various operators. As expected, the shape of the signal peak does not change much with the operators. For dipole operators, the distribution is slightly broader due to the worse performance of the vertex fit.

4.4.5 Sensitivity

BRANCHING FRACTION With an expected number of 0.44 background events from internal conversion decays in the signal region, the experiment is assumed to operate background-free in phase I up to about 300 days of data taking at a muon stopping rate of $10^8 \mu/s$ which corresponds to $N_\mu = 2.592 \cdot 10^{15}$ observed muon decays (see section 4.2).



The single-event sensitivity (SES)

$$\mathcal{B}_{\text{SES}}(\mu \rightarrow eee) \leq \frac{1}{\epsilon N_\mu} \quad (4.2)$$


depends on the efficiency ϵ to reconstruct a $\mu \rightarrow eee$ signal event and thus on the kinematics of the decay.

In the absence of a signal, an upper limit on the $\mu \rightarrow eee$ branching fraction can be derived. Assuming a Poisson distribution of the reconstructed m_{eee} masses, this yields an expected upper limit of

$$\mathcal{B}_{(1-\beta)\text{CL}}(\mu \rightarrow eee) \leq \frac{-\ln(\beta)}{\epsilon N_\mu} \quad (4.3)$$


 (a) Cut on the χ^2 of the vertex fit.


(b) Cut on the distance of the reconstructed vertex to the target.



(c) Cut on the reconstructed total momentum.

Figure 4.30: Signal efficiency to reconstruct the vertex of a $\mu \rightarrow eee$ event mediated via different effective operators for cuts on the χ^2 of the vertex fit, the reconstructed total momentum p_{eee} and the distance to the target of the reconstructed vertex $|d|$. The curves for 4-fermion operators (square markers) overlap as well as the two curves for the dipole operators (triangular markers) overlap. Inefficiencies from the online selection and from timing cuts are not taken into account.




Figure 4.31: Distribution of the reconstructed invariant mass m_{eee} of $\mu \rightarrow eee$ signal events mediated via different effective operators. The curves for phase space distributed decays and decays mediated via various 4-fermion operators overlap.

at a confidence level (CL) of $(1 - \beta)$.

The expected limits are summarized in table 4.7. In the case of phase space distributed events, a single event sensitivity of $2.3 \cdot 10^{-15}$ and an upper branching fraction limit of $5.2 \cdot 10^{-15}$ at 90 % CL is expected. With a smaller efficiency, the expected single event sensitivity for $\mu \rightarrow eee$ decays mediated via dipole operators reduces to $3.7 \cdot 10^{-15}$ and the expected upper limit at 90 % CL to about $8.4 \cdot 10^{-15}$. On the contrary, the efficiency for $\mu \rightarrow eee$ signal decays mediated by effective 4-fermion operators is enhanced. All investigated types yield a similar efficiency and thus expected branching fraction limits. The single event sensitivity is expected to be about $2.0 \cdot 10^{-15}$, and the upper limit at 90 % CL at $4.6 \cdot 10^{-15}$.

In figure 4.32, the evolution of the expected upper limits on the $\mu \rightarrow eee$ branching fraction with the run time of the experiment are shown for phase space distributed decays and decays mediated via effective operators. For the efficiency in the case of effective operators, an average over all dipole and an average over all 4-fermion operators is taken. In addition, the current best limit of $\mathcal{B} \leq 1 \cdot 10^{-12}$ at 90 % CL by the *SINDRUM* experiment is shown [64]. The high muon stopping rate in *Mu3e* allows to improve on this limit within a couple of data taking days. The final sensitivity is subject to the

Operator		$\mathcal{B}_{\text{SES}} [10^{-15}]$	$\mathcal{B}_{90\% \text{ CL}} [10^{-15}]$
Phase space		2.3	5.2
Dipole	$\mathcal{O}_{\text{D}}^{\text{R}}$	3.7	8.4
	$\mathcal{O}_{\text{D}}^{\text{L}}$	3.7	8.5
Scalar	$\mathcal{O}_{\text{S}}^{\text{RR}}$	2.0	4.6
	$\mathcal{O}_{\text{S}}^{\text{LL}}$	2.0	4.6
Vector	$\mathcal{O}_{\text{V}}^{\text{RR}}$	2.0	4.6
	$\mathcal{O}_{\text{V}}^{\text{LL}}$	2.0	4.6
	$\mathcal{O}_{\text{V}}^{\text{RL}}$	2.0	4.6
	$\mathcal{O}_{\text{V}}^{\text{LR}}$	2.0	4.6

Table 4.7: Expected upper limit on the $\mu \rightarrow eee$ branching fraction as single event sensitivity and at 90 % CL for $\mu \rightarrow eee$ signal decays mediated via different effective operators. Three long tracks are required. The branching fraction limits are expected for 300 days of data taking at $10^8 \mu/\text{s}$. The relative statistical uncertainties are on the 10^{-3} level.

efficiency to detect and reconstruct a signal event and thus depends on the nature of the New Physics.

WILSON COEFFICIENTS With equation B.6 (see section B.1 in the appendix), the expected upper limits on the branching fraction of $\mu \rightarrow eee$ decays can be translated to expected upper limits on the Wilson coefficients of the effective operators. In table 4.8, the expected limits are listed for the effective Lagrangian as formulated by Kuno et al. [46] and as formulated by Crivellin et al. [60, 61]. The limits are given in a one-at-a-time approach assuming that only a single operator mediates the signal decay. The limits are valid at the energy scale of the observable, thus in this case a scale $\mu \approx m_{\mu}$ is applicable. At other scales, contributions from all kinds of operators—also involving quarks and taus—arise through the renormalization group evolution of the Wilson coefficients.

4.5 SUMMARY ON $\mu \rightarrow eee$ SEARCHES

In this chapter, the sensitivity of the *Mu3e* experiment in phase I to the signal decay $\mu \rightarrow eee$ is estimated.

In a first step, a working point is defined at which background-free operation is feasible. As expected, a sufficient background suppres-

Operator	Coeff. of eq. B.1	Coeff. of eq. 4.1 at $\mu = m_\mu$ [GeV^{-2}]	
Dipole	\mathcal{O}_D^R	$ A_R \leq 1.9 \cdot 10^{-8}$	$\frac{ C_D^R }{\Lambda^2} \leq 2.1 \cdot 10^{-12}$
	\mathcal{O}_D^L	$ A_L \leq 1.0 \cdot 10^{-8}$	$\frac{ C_D^L }{\Lambda^2} \leq 2.1 \cdot 10^{-12}$
Scalar	\mathcal{O}_S^{RR}	$ g_1 \leq 1.9 \cdot 10^{-7}$	$\frac{ C_S^{RR} }{\Lambda^2} \leq 6.4 \cdot 10^{-12}$
	\mathcal{O}_S^{LL}	$ g_2 \leq 1.9 \cdot 10^{-7}$	$\frac{ C_S^{LL} }{\Lambda^2} \leq 6.3 \cdot 10^{-12}$
Vector	\mathcal{O}_V^{RR}	$ g_3 \leq 4.8 \cdot 10^{-8}$	$\frac{ C_V^{RR} }{\Lambda^2} \leq 1.6 \cdot 10^{-12}$
	\mathcal{O}_V^{LL}	$ g_4 \leq 4.8 \cdot 10^{-8}$	$\frac{ C_V^{LL} }{\Lambda^2} \leq 1.6 \cdot 10^{-12}$
	\mathcal{O}_V^{RL}	$ g_5 \leq 6.8 \cdot 10^{-8}$	$\frac{ C_V^{RL} }{\Lambda^2} \leq 2.2 \cdot 10^{-12}$
	\mathcal{O}_V^{LR}	$ g_6 \leq 6.8 \cdot 10^{-8}$	$\frac{ C_V^{LR} }{\Lambda^2} \leq 2.2 \cdot 10^{-12}$

Table 4.8: Expected upper limits on the coefficients of the effective operators at 90% CL for 300 days of data taking at $10^8 \mu/s$. The coefficients C are derived at $\mu = m_\mu \ll \Lambda$.




Figure 4.32: Evolution of the expected upper limit on the $\mu \rightarrow eee$ branching fraction at 90% CL with the days of data taking at $1 \cdot 10^8 \mu/s$ for effective operators. The efficiencies of the various operators of the same type are averaged.

sion is only achieved if all three decay particles are reconstructed as long tracks. The improved momentum and thus mass resolution is key to separate background and signal. For a prospective data taking period of 300 days at a rate of $1 \cdot 10^8 \mu/s$, a total of 0.44 background events from the radiative muon decay with internal conversion are expected in the signal region. Accidental background is estimated in a separate study [136], and is found to be subdominant to internal conversion decays even at looser cuts.

Signal events are studied in two scenarios. At first, $\mu \rightarrow eee$ events are investigated that are evenly distributed in phase space in order to study the acceptance and reconstruction efficiency for $\mu \rightarrow eee$ decays in *Mu3e*. Requiring that all tracks of the event can potentially be reconstructed as a long track, the acceptance amounts to 26%. Including inefficiencies of track and vertex reconstruction as well as the signal rejection due to cuts results in an overall efficiency to detect a $\mu \rightarrow eee$ signal event of 17%. This corresponds to a single event sensitivity of $2.3 \cdot 10^{-15}$, or an expected upper limit on the branching fraction of $5.2 \cdot 10^{-15}$ at 90% CL if no event is observed. The *Mu3e* experiment in phase I is thus sensitive to branching fractions about two hundred times smaller than what is excluded by the currently most stringent bounds set by the *SINDRUM* experiment [64].

The sensitivity is estimated by the means of simulation. The detector simulation reflects the current best knowledge on the design of the phase I experiment, but as the experiment is not yet constructed, some unknowns remain at this early stage. Some simplifications and idealizations are incorporated to overcome these unknowns.

This concerns in particular the pixel detector. The position of the pixel layers—as well as the position of all other detector parts—is known to full precision in the simulation. In the later experiment, there will be some uncertainty due to misalignment which affects the resolution of all reconstructed variables for instance the mass resolution. Moreover, the pixel sensors, and their readout, are fully efficient and produce no noise hits in the simulation. The sensors are placed without a gap between the individual chips. Although the performance of current prototypes hints at hit efficiencies $\gtrsim 99\%$ at low noise rates [148], the reconstruction efficiencies of all considered processes are thus slightly overestimated. In addition, noise hits and the remaining hits of tracks with a missing hit due to inefficiency enhance the risk to reconstruct fake tracks.

The timing suppression is assumed to factorize from the suppres-

sion of kinematics and vertexing. However, the performance of the timing suppression depends on the timing resolution and thus on the involved timing detector. In some configurations, the timing information stems only from the fibre detector as the scintillating tiles in the recurl stations are not reached by the particle. It has thus a larger uncertainty. Therefore, the timing suppression depends also on the kinematics of the decay.

Furthermore, the implemented physics processes of the radiative muon decay and the radiative muon decay with internal conversion are calculated to leading order. Both processes have recently been calculated to next-to-leading order [192–195]. Fortunately, the prediction for the rate of internal conversion background in the signal region reduces by about 17 % when next-to-leading order calculations are taken into account³. Therefore, the amount of background from internal conversion background is slightly overestimated and operating free of internal conversion background is well feasible.

The rate of accidental background is more difficult to estimate. Many types of background depend on the material distribution in the detector which can only be unsatisfactorily be modelled in simulation. Accidental background consisting of an electron-positron pair from Bhabha scattering paired with a positron of another muon decay is expected to dominate. This type is simulated up to sample sizes in the range of the run time of the phase I experiment, and no noteworthy contribution to the background in the signal region is found [136]. Another source of accidental background is studied in the context of dark photon searches (see chapter 6). The photon in the radiative muon decay $\mu \rightarrow e\gamma\nu\nu$ can convert in the target material and mimic a signal decay. In a sample equivalent to $1.3 \cdot 10^{11}$ muon decays, not a single event is left after reconstruction and cuts, even if no cut is placed on the reconstructed invariant three-particle mass. Although the sample size is too small to exclude contributions from this type of background, it seems very unlikely to have radiative decay events with nearby photon conversion in the signal region. So far, other types of accidental background are not yet studied with comparable sample sizes.

Apart from accidental background stemming from combinations of particle tracks from different origins, there is also the risk to create fake events due to deficiencies in the reconstruction, or by noise hits and inefficiencies in the pixel detector. Investigations on these

³ For a cut on the invariant mass at $m_\mu - 5m_e$, the next-to-leading order prediction is about 17 % lower than the leading order prediction [195]. For a cut at $m_\mu - m_e$, it is 23 % lower. The cut in this study is placed at 103.5 MeV. Therefore, a reduction of slightly more than 17 % is expected.

effects have just started.

The distribution of accidental background can be determined from data once the experiment is constructed and running. As the tracks are at least partially uncorrelated in time, the timing cuts requiring the coincidence of the tracks can be inverted in order to select only accidental background. It will require however a significant run time to collect a suitable amount of events in the signal region.

The cuts that are effective against accidental background are not chosen very strict in this sensitivity study. Thus, should the accidental background rate turn out to be significantly higher than assumed, cuts on the vertex fit χ^2 and the distance of the reconstructed vertex to the target can be tightened without too much loss of signal efficiency.

In the later experiment, there will be a further uncertainty stemming from the muon beam rate and thus the total number of stopped muons. This directly translates in an uncertainty on the branching fraction and has furthermore influence on the amount of accidental background. The number of stopped muons can be deduced from the online histograms and by random triggers during normal data taking. In the case of the *MEG* experiment, the number of muon decays has an uncertainty of 3.5 % [63]. A similar uncertainty can be expected for *Mu3e*.

The simple cutting strategy performed here will not be applied in the later data analysis. Instead, a similar strategy is pursued as in the analysis of the *MEG* experiment [63]. The signal region is blinded and the background is estimated in the sidebands. Then, a maximum likelihood fit is performed on the signal region and the number of signal and background events is extracted. In this way, correlations between the different observables can be included.

Using phase space distributed events is a suitable approach to study the acceptance of the detector as well as to compare to the preceding *SINDRUM* experiment but it does not account for the specific kinematics of the potential New Physics. Therefore, $\mu \rightarrow eee$ decays are studied in the context of effective theories in order to provide a model-independent perspective of the possibilities of *Mu3e*. The relevant effective operators that mediate muon to electron transitions are studied in a one-at-a time approach at an energy scale of the muon rest mass. Dipole-type operators typically produce an electron-positron pair with nearly zero invariant mass. It is shown that the acceptance for such events is in general lower, and also the

performance of the vertex reconstruction is decreased. As a result, the efficiency is reduced to about 11 % compared to 17 % in the case of phase space distributed events. For $\mu \rightarrow eee$ decays purely mediated via dipole operators, a single event sensitivity of $3.7 \cdot 10^{-15}$ and an upper branching fraction limit of about $8.4 \cdot 10^{-15}$ at 90 % CL is expected.

In contrast, all of the effective 4-fermion operators result in enhanced acceptances and thus efficiencies for $\mu \rightarrow eee$ decays. The efficiency to reconstruct a $\mu \rightarrow eee$ event mediated via a 4-fermion operator is found to be 19 %. Single event sensitivities of $2.0 \cdot 10^{-15}$ and upper branching fraction limits of $4.6 \cdot 10^{-15}$ at 90 % CL are expected.

In general, it is observed that the different operators leave a characteristic footprint in the Dalitz plots. The angular distributions contain in addition information about the chiralities of the involved fermion fields. Thus, given a sufficient number of events is observed, conclusions on the type of operator can be drawn.

The developed framework to study effective operators is also suited to study combinations of various operators. For combinations of 4-fermion operators, interference effects are negligible [61, 196], and thus the Dalitz plots and angular distributions are just superpositions of the distributions of the partaking operators. For combinations of 4-fermion and dipole operators, interferences are indeed expected. Such combinations are not studied here due to the large parameter space. In general, the expected sensitivities for these combinations range between the efficiency for 4-fermion operators and the efficiency for dipole operators.

The expected limits on the Wilson coefficients in this study are valid on the energy scale of the muon mass m_μ . Moving on to another scale, the Wilson coefficients of the various operators mix due to the evolution of the renormalization groups, and thus sensitivities to operators arise that not directly mediate $\mu \rightarrow eee$ decays. Performing this transition is however beyond the scope of this thesis.

The effective theory approach allows to compare between various experiments. Dipole type operators are naturally best tested in $\mu \rightarrow e\gamma$ searches. Assuming only dipole interactions, the branching fractions of $\mu \rightarrow eee$ and $\mu \rightarrow e\gamma$ searches can be related [46]

$$\frac{\mathcal{B}_{\mu \rightarrow eee}}{\mathcal{B}_{\mu \rightarrow e\gamma}} \simeq \frac{\alpha}{3\pi} \left(\ln \left(\frac{m_\mu^2}{m_e^2} \right) - \frac{11}{4} \right) \simeq 0.006 . \quad (4.4)$$

For both processes, the relevant scale is the muon rest mass. The current most stringent limit is set by the *MEG* experiment with $\mathcal{B} < 4.2 \cdot 10^{-13}$ at 90 % CL [63]. For the *MEG II* upgrade that is currently under construction, a single event sensitivity of $6 \cdot 10^{-14}$ is expected [73]. Already the existing limit gives a stronger constraint on dipole interactions than the phase I *Mu3e* experiment will be able to obtain.

Comparing $\mu \rightarrow eee$ searches to muon conversion on nuclei experiments requires to change the energy scale as the typical scale of the latter is at about 1 GeV. This is performed in [61]. A few results are summarized here. The prospects for future $\mu N \rightarrow eN$ searches as at the *COMET* or *Mu2e* experiment are very good with single event sensitivities in the range of 10^{-17} [81, 82]. For this reason, future muon conversion searches will outreach current and planned $\mu \rightarrow e\gamma$ searches in terms of sensitivity to dipole operators. 4-fermion operators mediating interactions between muons and electrons are however expected to be best constrained by *Mu3e*, because these operators enter at tree level in $\mu \rightarrow eee$ decays but at loop-level in $\mu N \rightarrow eN$. The remaining operators involved in muon to electron transitions are also best tested in the future $\mu N \rightarrow eN$ searches.

In the case of a signal in some or all of these searches, the effective theory approach is a tool to pin down possible interaction types. *Mu3e* is in a sense unique as it can potentially distinguish various operators given a sufficient number of events is observed. This is in parts also true for the $\mu N \rightarrow eN$ searches, as the sensitivity to the various operators differs with the choice of the nuclei and thus target material [61]. Whether runs with different targets are however feasible has to be seen. The target material influences the decay time of the captured muon. A sizeable decay time is necessary for $\mu N \rightarrow eN$ experiments as events in a certain time window after stopping of the muons in the target are vetoed in order to suppress beam-related background. Therefore, only a limited choice of materials remains.

With the phase II upgrade to the *Mu3e* experiment, the detector will be equipped with an additional recurl station upstream and downstream of the existing detector. The muon stopping rate is increased to $2 \cdot 10^9 \mu/s$. The beam profile is expected to be smaller so that the radius of the target can be reduced.

In this way, the efficiency for the reconstruction of long tracks increases and thus in general the efficiency to detect a three particle event. Only minor losses are expected for the cases in which the particles cross the gaps between the stations. Overshooting the recurl

stations is no longer possible.

With the smaller target size, the performance of the vertex fit becomes better. An improvement in the vertex position resolution results in an improved angular resolution and thus also mass resolution.

The increase in the number of observed muon decays requires to tighten the lower m_{eee} cut in order to suppress background from internal conversion decays. As shown before, this cut has a strong rejection power on internal conversion background and affects the signal efficiency moderately.

The accidental background might no longer be negligible. In the case of events with two correlated tracks, e. g. involving Bhabha scattering, the background rate scales with the muon stopping rate. And in the case of three uncorrelated tracks, the rate scales even with the muon stopping rate to a power of two. Thus, processes aside from Bhabha events should be considered. The increased rate of accidental background requires stricter cuts on timing and vertexing.

If Bhabha events turn out to be problematic, events with an invariant electron-positron mass m_{ee} of around 7 MeV can be excluded as the reconstructed m_{ee} spectrum of Bhabha events is about maximum in this region. Similar measures can also be taken against internal conversion background which is maximum at a slightly smaller m_{ee} . However, such a selection would also reduce the efficiency for dipole-mediated $\mu \rightarrow eee$ decays.

LEPTON FLAVOUR VIOLATING TWO-BODY DECAYS OF THE MUON

Aside from the decay $\mu \rightarrow eee$, a range of other muon decay channels involving physics beyond the SM can be investigated with *Mu3e*. Among the muon LFV experiments, *Mu3e* has a uniquely large geometric and momentum acceptance for electrons and positrons. In combination with a precise tracking of electrons and positrons and an immense number of observed muon decays, unprecedented sensitivities can be accomplished.

In this chapter, the search strategy for LFV two-body decays of the muon with the *Mu3e* experiment is elaborated and the sensitivity to these decays in phase I is estimated.

5.1 SEARCHES FOR $\mu \rightarrow eX$ WITH MU3E

In this study, it is assumed that the boson X in the decay $\mu^+ \rightarrow e^+ X$ escapes the experiment without being detected. Being long-lived, the width of the mass m_X of X is narrow. The decay is therefore identified by a mono-energetic positron whose energy is fully determined by m_X . Such a search has an irreducible background from SM decays of the muon, i. e. the dominant Michel decay, the radiative decay and the radiative decay with internal conversion, as well as from photon conversion events. Potential $\mu \rightarrow eX$ events would emerge as a sharp peak on the continuous momentum spectrum of the SM processes. This is illustrated in figure 5.1 for the dominant Michel decay and for $\mu \rightarrow eX$ decays with a mass m_X of 60 MeV. In measured data, the peak will be smeared out because of the finite momentum resolution. In fact, the sensitivity to $\mu \rightarrow eX$ is ultimately determined by the momentum resolution. As outlined in section 2.2.7, it is not possible to store the raw detector information for all events due to data rate restrictions. Searches for $\mu \rightarrow eX$ can thus only be performed on data from the online reconstruction.

The kinematically allowed mass range for X in $\mu \rightarrow eX$ decays is given by

$$0 \text{ MeV} \leq m_X \leq m_\mu - m_e \approx 105.15 \text{ MeV}, \quad (5.1)$$




Figure 5.1: Momentum spectrum of decay electrons in $\mu \rightarrow eX$ searches. The background shown in blue is represented by the dominating Michel decay $\mu \rightarrow e\nu\nu$. Superimposed in red is a potential $\mu \rightarrow eX$ signal for a mass m_X of 60 MeV.

but the experimentally accessible range is further limited.

In the *Mu3e* experiment, only electrons with a transverse momentum of at least 10 MeV can be tracked when operating at the nominal magnetic field strength of 1 T. With a lower transverse momentum, the particle cannot cross all four pixel layers in the central detector. Thus, electrons and positrons with momenta below 10 MeV are not detected, and in general, the acceptance for low momentum particles is lower than for high momentum particles. Because of the geometric acceptance, m_X above 95 MeV cannot be measured using the default reconstruction concept (see figure 5.2).

Low masses are not affected by the geometric acceptance, but these might need to be excluded if the characteristic edge of the dominant Michel decay of the muon at 52.83 MeV is used for momentum calibration. This affects momenta above 50 MeV, and thus sets a lower limit on the accessible mass range of $m_X \geq 25$ MeV. Other processes such as Mott or Bhabha scattering could potentially be used for the calibration instead of the Michel decay. This would allow to test the whole low mass range. These calibration methods are under investigation and are also employed in the *MEG* experiment [197].

In the following, results obtained for m_X between 0 MeV and 25 MeV are given under the condition that an alternative calibration method can be employed. In section 5.6, a different approach is studied in




Figure 5.2: Momentum of the positron in $\mu \rightarrow eX$ as a function of the mass of X . In red is the momentum range indicated that is not accessible in $Mu3e$, and thus limits the accessible mass range. The green region needs to be excluded in the case the Michel edge is used for momentum calibration.

which the calibration is simultaneously extracted from the momentum spectrum. In this case, a region around the expected peak from $\mu \rightarrow eX$ decays is left out in the momentum calibration fit.

5.2 ESTIMATING THE SENSITIVITY

The $\mu \rightarrow eX$ events can be identified by a peak on the reconstructed momentum of positrons from muon decays. The mean of the peak p_X is specific to the mass m_X and the width of the peak is determined by the momentum resolution Δp . In the following, the number of signal and background events in a certain momentum window $p_X \pm n\Delta p$ are counted. The width is chosen to be proportional to the momentum resolution Δp .

The smallest number of reconstructed signal events that can be identified as an excess depends on the number of reconstructed background events under the peak. Given a large number N_μ of muon decays is observed, a Poisson distribution can be assumed for the number of background events. The number of signal events thus becomes

$$N_{\mu \rightarrow eX}^{\text{rec}} = c \sqrt{N_{\text{BG}}^{\text{rec}}}. \quad (5.2)$$

The proportionality factor c depends on the chosen confidence level.

The width of the momentum window determines the selection efficiency a by which a reconstructed signal event is accepted. As the size of the window is chosen to be proportional to the momentum resolution, a is constant. The number of reconstructed signal events is thus

$$N_{\mu \rightarrow eX}^{\text{rec}} = \mathcal{B}_{\mu \rightarrow eX} \epsilon_X a N_\mu \quad (5.3)$$

with the signal branching fraction $\mathcal{B}_{\mu \rightarrow eX}$ and the reconstruction efficiency ϵ_X for the positron of $\mu \rightarrow eX$.

For the number of background events, the differential branching fraction has to be integrated over the range of the momentum window. As long as the momentum window is small and the background distribution smooth, this can be approximated as follows

$$\int_{p_X - n\Delta p}^{p_X + n\Delta p} \frac{d\mathcal{B}_{\text{BG}}(p)}{dp} dp \approx 2n \left. \frac{d\mathcal{B}_{\text{BG}}(p)}{dp} \right|_{p=p_X} \Delta p. \quad (5.4)$$

The expected number of reconstructed background events becomes

$$N_{\text{BG}}^{\text{rec}} \approx 2n \left. \frac{d\mathcal{B}_{\text{BG}}(p)}{dp} \right|_{p=p_X} \Delta p \epsilon_{\text{BG}}(p_X) N_\mu \quad (5.5)$$

with a background reconstruction efficiency ϵ_{BG} that depends in general on the investigated momentum range.

Using equation 5.2 and 5.5, the smallest signal branching fraction to which the experiment is sensitive can be calculated

$$\mathcal{B}_{\mu \rightarrow eX}^{\text{limit}} = \frac{N_{\mu \rightarrow eX}^{\text{rec}}}{N_\mu} \quad (5.6)$$

$$= \frac{\sqrt{2n} c}{a} \sqrt{\left. \frac{d\mathcal{B}_{\text{BG}}(p)}{dp} \right|_{p=p_X}} \frac{\sqrt{\Delta p \epsilon_{\text{BG}}(p_X)}}{\epsilon_X} \frac{1}{\sqrt{N_\mu}}. \quad (5.7)$$

The first factor in equation 5.7 purely depends on analysis choices and can be optimized. The distribution of background events enters with the square root. This is dominated by the matrix elements of the contributing background processes but includes in general also reconstructed fake events. As expected, a better sensitivity is achieved in momentum ranges with low background. The third term includes the reconstruction efficiencies and momentum res-

olution. A good signal efficiency has a stronger impact on the sensitivity than the background rejection. Finally, the sensitivity improves with the square root of the number of observed muon decays. The run time can be increased to achieve better limits up to a point where systematic uncertainties dominate.

The acceptance, reconstruction efficiencies, and momentum resolution are estimated with the detector simulation.

5.2.1 Toy Monte Carlo Approach

The branching fraction limits for $\mu \rightarrow eX$ decays are derived by the means of toy Monte Carlo studies using ROOFIT [198].

As input, histograms of the reconstructed momentum for SM muon decays and $\mu \rightarrow eX$ decays are used.

The generated background sample contains about $5 \cdot 10^8$ muon decays and is thus significantly smaller than the $\mathcal{O}(10^{15})$ muon decays that are expected to be observed in phase I of *Mu3e*. In order to even out potential fluctuations stemming from the smaller sample size, smoothing is applied to the SM histogram using the algorithm provided by ROOT [199] (based on [200]) with a single iteration. In one iteration, the median is taken successively over three, five, and three bins.

The numbers of the generated SM decays $N_{\text{gen}}^{\text{SM}}$ and the generated $\mu \rightarrow eX$ decays $N_{\text{gen}}^{\mu \rightarrow eX}$ are derived from the simulation data. The number of reconstructed tracks $N_{\text{rec}}^{\text{SM}}$ for SM decays and $N_{\text{rec}}^{\mu \rightarrow eX}$ for $\mu \rightarrow eX$ are the sum of the entries in the input histograms. The efficiencies to reconstruct a $\mu \rightarrow eX$ signal event or a background event are thus

$$\epsilon_{\mu \rightarrow eX} = \frac{N_{\text{rec}}^{\mu \rightarrow eX}}{N_{\text{gen}}^{\mu \rightarrow eX}}, \quad \epsilon_{\text{SM}} = \frac{N_{\text{rec}}^{\text{SM}}}{N_{\text{gen}}^{\text{SM}}}. \quad (5.8)$$

The background probability density function (p.d.f.) $f_b(p)dp$ is derived from the reconstructed momentum histogram of SM processes, and the signal p.d.f. $f_s(p)dp$ from the histogram of $\mu \rightarrow eX$ decays (see figures 5.3a and 5.3b). From f_b a total of N_{toyMC} events is derived, to which a signal-plus-background p.d.f.

$$f_{s+b} = (1 - f_{\text{sig}}) \cdot f_b + f_{\text{sig}} \cdot f_s, \quad -1 \leq f_{\text{sig}} \leq 1, \quad (5.9)$$

is fitted with the signal fraction f_{sig} as free fit parameter. Fits are performed using the maximum likelihood method on binned data. Figure 5.3c shows an exemplary fit for a mass of 60 MeV. This is re-

Confidence level	Quantile n_σ
0.90	1.282 σ
0.95	1.645 σ
0.99	2.326 σ

Table 5.1: Quantiles of a Gaussian distribution with mean 0 for various confidence levels for one-sided intervals (taken from [201]).

peated for a large number of samples per mass. The resulting fitted signal fractions are shown in figure 5.3d. A fit of a Gaussian distribution with fixed mean of zero¹ and width σ_{toyMC} as fit parameter is applied to the signal fractions. In order to obtain the signal fraction that corresponds to the 90% (95%; 99%) confidence level (CL), the fitted and normalized Gaussian distribution is integrated from $-\infty$ up to f_{sig} such that the result equals 0.90 (0.95; 0.99). In the case of the Gaussian distribution, these quantiles are multiples n_σ of σ_{toyMC} (see table 5.1).


Thus, an observation of $n_\sigma \sigma_{\text{toyMC}} N_{\text{toyMC}}$ reconstructed events from $\mu \rightarrow eX$ corresponds to the 90% (95%; 99%) CL in a sample of N_{toyMC} reconstructed SM muon decays. As not every $\mu \rightarrow eX$ event is reconstructed, this number is multiplied by the inverse of the reconstruction efficiency $\epsilon_{\mu \rightarrow eX}$ in order to obtain the actual number of $\mu \rightarrow eX$ decays. Thus, the total number of $\mu \rightarrow eX$ events corresponding to a certain CL for a toy Monte Carlo study of N_{toyMC} events is

$$n_\sigma \sigma_{\text{toyMC}} N_{\text{toyMC}} \frac{N_{\text{rec}}^{\mu \rightarrow eX}}{N_{\text{rec}}} \frac{1}{\epsilon_{\mu \rightarrow eX}}. \quad (5.10)$$


This can be scaled to the desired number N_μ of muon decays. The N_{toyMC} toy Monte Carlo events are replaced by $N_\mu \epsilon_{\text{SM}}$. With a larger number of muon decays also σ_{toyMC} becomes smaller. Studies with N_{toyMC} ranging over several orders of magnitude have proven that σ_{toyMC} scales with $\frac{1}{\sqrt{N_{\text{toyMC}}}}$ as expected (see section C.2 in the appendix). Therefore, σ_{toyMC} is replaced

$$\sigma_{\text{toyMC}} \rightarrow \sigma_{\text{toyMC}} \sqrt{\frac{N_{\text{toyMC}}}{N_\mu \epsilon_{\text{SM}}}}. \quad (5.11)$$


¹ See section C.2 in the appendix for more details.




(a) Histogram of the reconstructed momentum from SM background processes. About $5 \cdot 10^8$ decays in the target region are generated.



(b) Histogram of the reconstructed momentum from $\mu \rightarrow eX$ signal events with $m_X = 60 \text{ MeV}$. About 75 000 decays in the target region are generated.



(c) A single sample of events derived from f_b . The solid blue line is a fit of f_{s+b} to the events. The fitted amount of signal (solid red line) and background (dashed blue line) are also shown. Please note, that the solid and dashed blue line are indistinguishable on the logarithmic scale.



(d) Histogram of the signal fraction f_{sig} for a complete study with fit of a Gaussian distribution (solid blue line).

Figure 5.3: Exemplary toy Monte Carlo study for short track reconstruction and $\mu \rightarrow eX$ events with $m_X = 60 \text{ MeV}$. Similar studies for the other reconstruction modes can be found in section C.2 of the appendix.

The limit on the branching fraction at a certain CL becomes

$$\mathcal{B} \leq n_\sigma \sigma_{\text{toyMC}} \sqrt{\frac{N_{\text{toyMC}}}{N_\mu \epsilon_{\text{SM}}}} N_\mu \frac{\epsilon_{\text{SM}}}{\epsilon_{\mu \rightarrow eX}} \quad (5.12)$$

$$\leq n_\sigma \sigma_{\text{toyMC}} \sqrt{\frac{N_{\text{toyMC}} \epsilon_{\text{SM}}}{N_\mu}} \frac{1}{\epsilon_{\text{SM}}} . \quad (5.13)$$

Equation 5.13 is the basis for the following sensitivity studies. It is similar to equation 5.7. As the toy Monte Carlo studies are performed on the full momentum spectrum, the momentum dependent efficiencies $\epsilon_{\text{BG}}(p)$ and ϵ_X are replaced with the overall reconstruction efficiencies ϵ_{SM} and $\epsilon_{\mu \rightarrow eX}$. The signal and background distribution as well as the acceptance of the detector are contained in σ_{toyMC} .

The toy Monte Carlo method as applied here is a suitable measure to estimate the sensitivity of the experiment at this early stage. It has however some deficiencies.

The Michel decay of the muon is implemented with next-to-leading order calculations for the matrix element. Correspondingly, the higher-order processes $\mu \rightarrow e\gamma\nu\nu$ and $\mu \rightarrow eee\nu\nu$ are implemented with a leading-order matrix element. Recent next-to-leading order calculations of $\mu \rightarrow e\gamma\nu\nu$ and $\mu \rightarrow eee\nu\nu$ have shown corrections in the order of 1 % [192–195]. As the background distribution enters only with the square root, neglecting higher-order calculations does only marginally affect the sensitivities derived in the following.

The detector simulation comprises some idealizations. The position of all active detector parts is known with full precision. Misalignment in the later detector will degrade the momentum resolution. The simulated pixel detector is idealized with fully efficient sensors that are placed without a spacing. In this way, the overall efficiency in the study presented here is slightly overestimated. Furthermore, noise and hit clusters from charge sharing are not simulated. These increase the risk to reconstruct fake tracks. In addition, the target position and beam profile and rate as well as the material distribution will only be known with some uncertainty in the later experiment.

All of the above described effects influence the reconstructed momentum spectra. As the simulation cannot fully reflect the real experiment, using a fixed background and signal shape from simulation is not an optimum solution. A possible analysis strategy is described at the end of this chapter in section 5.9.

The following toy Monte Carlo studies are performed with 10^4 samples with $N_{\text{toyMC}} = 10^9$ events each². 300 days of data taking with a total of $N_\mu = 2.592 \cdot 10^{15}$ muon decays are assumed.

5.3 SIMULATION OF $\mu \rightarrow eX$ EVENTS

5.3.1 Event Generator

As $\mu \rightarrow eX$ is a two-body decay, the kinematics are fully constrained by the muon (m_μ) and electron mass (m_e) and the mass of the boson X (m_X). The energies of the decay products in the centre of mass system of the muon are

$$E_e = \frac{m_\mu^2 - m_X^2 + m_e^2}{2m_\mu}, \quad (5.14)$$

$$E_X = m_\mu - E_e = \frac{m_\mu^2 + m_X^2 - m_e^2}{2m_\mu}, \quad (5.15)$$

and the absolute values of the momenta are

$$p_e = p_X = \sqrt{E_e^2 - m_e^2} = \sqrt{\left(\frac{m_\mu^2 - m_X^2 + m_e^2}{2m_\mu}\right)^2 - m_e^2}. \quad (5.16)$$

The momentum of the positron in dependence of m_X is shown in figure 5.2. At $m_X = 0$ MeV, the momentum is maximum with $\frac{m_\mu^2 - m_e^2}{2m_\mu} \approx 52.8$ MeV.

For a purely scalar interaction, the decay is isotropic. The positron and X boson are generated in a back-to-back topology with a random orientation with respect to the beam axis.

Taking also pseudo-scalar couplings into account, an angular dependence arises. The differential decay rate $d\Gamma$ is no longer constant but depends on the angle θ between the positron momentum \mathbf{p}_e and the muon polarization \mathbf{P} (see also equation 1.21).

$$\frac{d\Gamma}{d\cos\theta} \propto 1 + hP \cos\theta, \quad -1 \leq h \leq 1 \quad (5.17)$$

Herein, P denotes the net polarization and the coupling parameter h describes the strength of the angular coupling. In this study, the muon polarization is parallel to the beam axis and thus θ is the polar

² A total of $\mathcal{O}(10^{15})$ events would be desirable, however in the implementation of toy Monte Carlo studies in ROOFIT, the maximum number of events is limited to the range of a 32 bit integer, i. e. up to 2.147.483.647. Generating notably larger samples than $1 \cdot 10^9$ events would in addition take a lot of computation time.

angle of the positron momentum. A fully polarized muon beam is assumed.

The distribution of events in $\cos\theta$ is generated by using an accept-reject method [167, 168] based on equation 5.17.

Background events are generated as SM muon decays at a stopping rate of $1 \cdot 10^8 \mu/\text{s}$. The various decay modes occur at a rate corresponding to their respective branching fraction (see section 2.3.1).

A sample of about $5 \cdot 10^8$ SM decays is generated, as well as about 75 000 $\mu \rightarrow eX$ events per mass m_X ranging from 0 MeV to 99 MeV in steps of 1 MeV.

5.3.2 Reconstruction

In the baseline design of the *Mu3e* experiment, *short* tracks are reconstructed online during the event filtering stage from four hits—one in each layer of the central pixel detector (see section 2.2.7 for details). A histogram of the reconstructed momentum and the cosine of the polar angle θ is stored. Storage of further information like hit positions is not foreseen in order to reduce the data rate to a reasonable level. Hence, offline reconstruction of recurling tracks is not possible for the two-body decay searches.

With an upgrade to the filter farm, it might become feasible to reconstruct also recurling tracks from six or eight hits online (*long tracks*). The sensitivity to $\mu \rightarrow eX$ is estimated for both reconstruction modes in the following.

A fine binning of 100 keV is chosen for the histograms of the reconstructed momenta.

5.3.3 Selection and Cuts

Searches of the type $\mu \rightarrow eX$ have an irreducible background from SM muon decays. But also shortcomings in the reconstruction of SM processes add to the background. The amount of such *fake tracks* can be reduced by a proper data selection.

A fake track can have various origins. For instance, it can be a track that is reconstructed from hits caused by different particles, or a track of an electron that is mistaken for a positron. Also particles that perform many loops in the central detector are prone to produce fake tracks, either because a secondary loop is interpreted as a different particle's trajectory or because tracks are built from hits belonging to different loops.

Based on truth information from the simulation, *true* $\mu \rightarrow eX$ or SM

muon decays are identified. A track is regarded to be *true* if it is built from hits of the first loop of a positron. In the case of $\mu \rightarrow eX$ events, this positron is in addition required to stem from a $\mu^+ \rightarrow e^+ X$ decay. For SM background, no requirement on the decay mode is placed. All other tracks are interpreted as *fake* tracks.

5.3.3.1 Pre-Selection

The sensitivities resulting from this study are quoted for a certain number of stopped muons. Therefore, only events are selected whose true vertices lie in a cylinder surrounding the target. This is the only selection cut on truth information applied in this analysis.

A further pre-selection is made based on reconstructed variables.

In order to ensure a certain track reconstruction quality, short tracks are required to have a χ^2 value of the track fit of smaller or equal to 32, and long tracks of smaller or equal 48. This pre-selection cut is considerably looser than the cut that will be placed in the following. The $\mu^+ \rightarrow e^+ X$ decays have only positrons and no electrons in the final state. The charge of the particle can be identified by the sense of rotation of the tracks. This is encoded in the sign of the reconstructed momentum. Only tracks with a negative reconstructed momentum—as compatible with a positron—are taken into account. Nevertheless, a small amount of electron tracks cannot be avoided this way. These tracks stem from recurling tracks with eight hits that can be reconstructed in both rotational directions. In the following, always the absolute value of the reconstructed momentum is used. In addition, a loose selection cut on the reconstructed momentum is applied. Only tracks with $p_{\text{rec}} \leq 100 \text{ MeV}$ are kept. This selection cut rejects obvious cases of mis-reconstruction but no relevant physics processes.

5.3.3.2 Selection Cuts

As single tracks are studied, the choice of cut variables is limited. For this analysis, cuts on three variables are applied: the χ^2 of the track fit, the dip angle λ_{01} and the intersection z_0 of the positron track with the target region.

In the following, the effect of the various cuts on SM background and $\mu \rightarrow eX$ signal events are shown. Each cut is applied individually on the pre-selected data without cuts on the other two variables being present.

For the discussion of $\mu \rightarrow eX$ events, a mass m_X of 60 MeV is chosen exemplarily. For this mass, the resulting peak in the positron spec-

trum is about centred above the spectrum of the SM decays. Similar figures for other masses can be found in the appendix C.1.


VERTEX ESTIMATE z_0 The muon decay vertices are expected to be located in the target region. But with the X boson being invisible to the detector, the decay vertex cannot be reconstructed in the case of $\mu \rightarrow eX$. It can however be estimated by propagating the reconstructed track backwards onto a cylinder around the beam axis that has the radius of the target. The z coordinate of the first intersection of the track with the cylinder is z_0 .

On the left-hand side of figure 5.4, the distribution of reconstructed true and fake tracks in z_0 is shown for SM processes in the three reconstructions. Similar distributions for signal decays with m_X of 60 MeV are shown in figure 5.5. The true tracks merely reflect the stopping distribution of the muons. The distribution steeply falls off at both edges at around 50 mm as expected from the dimensions of the target. The peak at $z_0 \approx 0$ is an artefact from the reconstruction and derivation of z_0 (see section C.1.1 in the appendix for more details). The fake tracks however are flatly distributed up to a $|z_0|$ of 60 mm. This is caused by the dimensions of the vertex pixel layers.


On the right-hand side of figure 5.4 and figure 5.5, the efficiency ϵ , purity p , as well as the product of the two quantities are shown for various cuts on $|z_0|$. The efficiency in this context is to be understood as the fraction of reconstructed true events that remain after cuts compared to all reconstructed true tracks. Accordingly, the purity is the fraction of true tracks within all reconstructed tracks after cuts. As can be directly seen from the distribution of true and fake tracks in z_0 , a sole cut on $|z_0|$ is not very effective in suppressing fake tracks. The purity does not improve much for tighter cuts on $|z_0|$. The improvement of the purity is a bit more pronounced for reconstructed short tracks. As short tracks have less constraints compared to long tracks, this reconstruction type simply yields higher amounts of fake tracks hence the stronger effect of $|z_0|$ cuts. A cut is set at $|z_0| \leq 50$ mm where the efficiency reaches a plateau for both types of reconstruction. This is in good agreement with the length of the target.

χ^2 OF THE TRACK FIT The track quality of the selected event sample can be improved by cutting on the χ^2 value of the track fit. In general, true tracks are expected to have smaller χ^2 values than fake tracks.


Figures 5.6 and 5.7 show the χ^2 distributions for SM muon decays




(a) Distribution in z_0 for reconstructed short tracks.



(b) Efficiency and purity for cuts on z_0 using reconstructed short tracks.



(c) Distribution in z_0 for reconstructed long tracks.



(d) Efficiency and purity for cuts on z_0 using reconstructed long tracks.

Figure 5.4: Selection cuts on z_0 for SM processes reconstructed as short and long tracks. On the left hand side, the z_0 distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of z_0 . The sample contains about $5 \cdot 10^8$ positrons from SM muon decays.

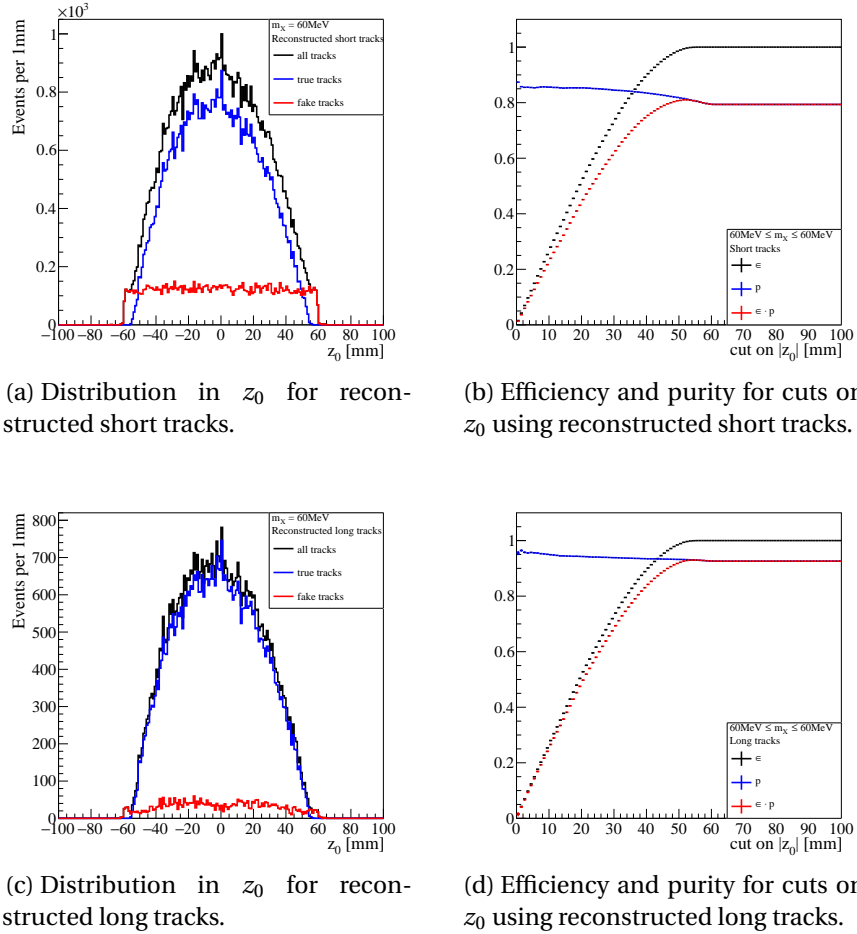





Figure 5.5: Selection cuts on z_0 for $\mu \rightarrow eX$ decays with a mass m_X of 60 MeV reconstructed as short and long tracks. On the left hand side, the z_0 distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of z_0 . The sample contains about 75 000 $\mu \rightarrow eX$ decays.




(a) Distribution in χ^2 for reconstructed short tracks.



(b) Efficiency and purity for cuts on χ^2 using reconstructed short tracks.

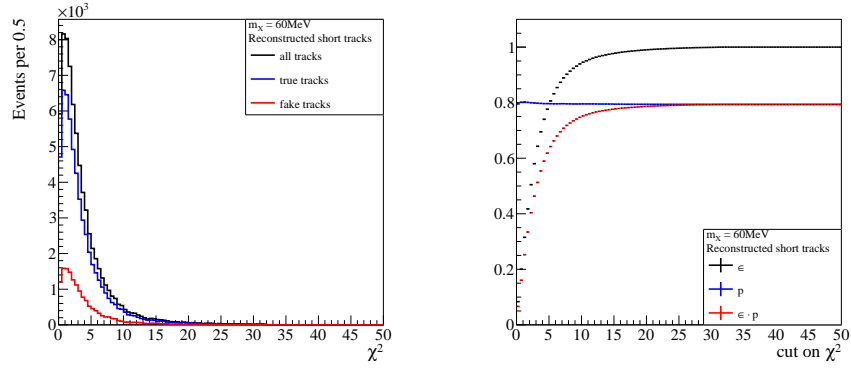


(c) Distribution in χ^2 for reconstructed long tracks.



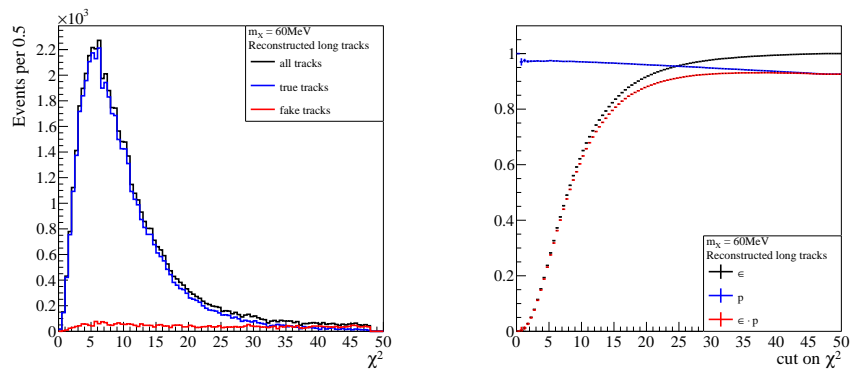
(d) Efficiency and purity for cuts on χ^2 using reconstructed long tracks.

Figure 5.6: Selection cuts on the χ^2 of the track fit for SM processes reconstructed as short and long tracks. On the left hand side, the χ^2 distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of χ^2 . The sample contains about $5 \cdot 10^8$ positrons from SM muon decays.



(a) Distribution in χ^2 for reconstructed short tracks.

(b) Efficiency and purity for cuts on χ^2 using reconstructed short tracks.



(c) Distribution in χ^2 for reconstructed long tracks.

(d) Efficiency and purity for cuts on χ^2 using reconstructed long tracks.

Figure 5.7: Selection cuts on the χ^2 of the track fit for $\mu \rightarrow eX$ decays with a mass m_X of 60 MeV reconstructed as short and long tracks. On the left hand side, the χ^2 distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of χ^2 . The sample contains about 75 000 $\mu \rightarrow eX$ decays.

and $\mu \rightarrow eX$ events, respectively, as well as efficiency and purity for cuts on χ^2 . Short tracks are reconstructed with fewer parameters than long tracks and therefore yield smaller χ^2 values. This is clearly visible in the χ^2 distributions. Hence, the cut for short tracks can be chosen at a smaller χ^2 compared to long tracks.


The purity is only slightly affected by the χ^2 cut, but a tight cut strongly reduces the efficiency. Therefore, a rather loose cut on χ^2 is chosen. For short track reconstruction, the cut is set at $\chi^2 \leq 15$, and at $\chi^2 \leq 30$ for long track reconstruction³.

DIP ANGLE λ_{01} The largest contribution to fake tracks stems from recurling particles that are emitted nearly perpendicular to the beam axis. As illustrated in figure 5.8, such particles perform many turns in the central detector because they do not traverse much material that would slow down and finally stop them. Fake tracks are thus tracks from secondary turns which are misinterpreted as separate particles and tracks that are reconstructed from hits of different turns. These fake tracks are problematic for the $\mu \rightarrow eX$ search as multiple fake tracks are reconstructed with nearly the same momentum which could emerge as a peak on the momentum spectrum. As can be seen from figure 5.8, a cut on the dip angle λ is an efficient way to reduce the number of fake tracks caused by *looping* particles. In addition, a λ cut reduces electron track contamination. An electron can only be misinterpreted as a positron when it is reconstructed in the wrong direction of rotation. This is only possible for recurling particles that cross not only the outer but also the vertex layers at least twice and thus have small $|\lambda|$.


On reconstruction level, the variable λ_{01} is available which is the dip angle between hits in the two vertex layers. This is a very good approximation of the λ of the true track as λ would equal λ_{01} in the absence of multiple scattering in the innermost layer.

The λ_{01} distributions for SM processes and $\mu \rightarrow eX$ events as well as the effect of cuts on λ_{01} on efficiency and purity are shown in figures 5.9 and 5.10, respectively. For short and long track reconstruction, there is a sharp peak visible for small $|\lambda_{01}|$ which is caused by the fake tracks. The small peak in the true track distribution is due to multiple reconstructions of the same track. This happens if the particle crosses an overlap of two pixel layers and thus creates two hits in one layer. The drop at $\lambda_{01} = 0$ rad is a result of pixelation.

³ In principle, it would be possible to set individual χ^2 cuts on six-, and eight-hit tracks in the long track reconstruction. As the relative fraction of eight-hit tracks in the long track ensemble is however small, a global cut is chosen for the sake of simplicity.



(a) Transverse view.



(b) Longitudinal view.




(c) Dip angle λ .


Figure 5.8: An event with a particle performing many turns in the central detector shown in transverse and longitudinal view, and the definition of the dip angle λ . The black tracks are reconstructed tracks that can be matched to a true particle, the red tracks are fake tracks.




(a) Distribution in λ_{01} for reconstructed short tracks.



(b) Efficiency and purity for cuts on λ_{01} using reconstructed short tracks.




(c) Distribution in λ_{01} for reconstructed long tracks.




(d) Efficiency and purity for cuts on λ_{01} using reconstructed long tracks.


Figure 5.9: Selection cuts on λ_{01} for SM processes reconstructed as short and long tracks. On the left hand side, the λ_{01} distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of λ_{01} . The sample contains about $5 \cdot 10^8$ positrons from SM muon decays.




(a) Distribution in λ_{01} for reconstructed short tracks.



(b) Efficiency and purity for cuts on λ_{01} using reconstructed short tracks.



(c) Distribution in λ_{01} for reconstructed long tracks.



(d) Efficiency and purity for cuts on λ_{01} using reconstructed long tracks.

Figure 5.10: Selection cuts on λ_{01} for $\mu \rightarrow eX$ decays with a mass m_X of 60 MeV reconstructed as short and long tracks. On the left hand side, the λ_{01} distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of λ_{01} . The sample contains about 75 000 $\mu \rightarrow eX$ decays.

The true tracks exhibit a broad distribution over the complete accessible range of λ_{01} whereas basically no fake tracks have large $|\lambda_{01}|$. As a result, the purity steeply rises with tighter cuts on $|\lambda_{01}|$, but also the efficiency deteriorates.

In the case of long tracks, there is an additional dip in the λ_{01} distribution left and right from $\lambda_{01} = 0$ rad. This is due to tracks that hit the gap between the central station and recurl stations.

For both types of reconstructed tracks, the cut is set at $|\lambda_{01}| \geq 0.1$ rad.

COMBINING THE CUTS The cuts are summarized in table 5.2. Details on the efficiency relative to all reconstructed tracks and the purity are given in table 5.3 for the SM processes and for $\mu \rightarrow eX$ decays as an average over masses m_X between 0 MeV and 99 MeV. With the combined cuts, for both types of reconstructed tracks efficiencies of above 80 % are achieved with the exception of $\mu \rightarrow eX$ decays with very high m_X that yield slightly smaller efficiencies (see table C.1 in the appendix). The cut on λ_{01} affects tracks with small momenta more than tracks with large momenta due to the acceptance cut-off for small transverse momenta. Therefore, the efficiency for large m_X is decreased. The purity is significantly enhanced in the case of short tracks. For long track reconstruction, the initial purity is already at a high level and therefore improves only moderately by applying the cuts. Overall, purities above 95 % are achieved, again with the exception of high m_X .

Figure 5.11 shows the reconstructed momentum spectrum of the SM sample and a $\mu \rightarrow eX$ sample with m_X of 60 MeV with applied cuts. For short and long track reconstruction, the contribution of fake tracks is considerably reduced while a large fraction of the true tracks remains. The best reduction of fake tracks is achieved in the region of the momentum peak of the $\mu \rightarrow eX$ events. This is a result of the cut on λ_{01} . Most of the fake tracks are reconstructed from hits of different turns of a recurling particle. As the recurling particles loose energy in the detector material, tracks built from hits of later turns have a smaller reconstructed momentum. These tracks naturally have larger λ_{01} and are thus more likely to pass the λ_{01} cut. In the case of SM processes, the shape of the momentum spectrum of the dominating Michel decay $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$ with the characteristic edge at about 50 MeV can be clearly identified. The spectrum is cut at momenta smaller than 10 MeV because of the acceptance cut-off for low transverse momenta. The spectrum of reconstructed long tracks exhibits two other features related to the acceptance. The momentum spectrum is partially cropped for momenta larger than about 40 MeV. At these large momenta, the particles can have suffi-

Short track reconstruction	Long track reconstruction
$ z_0 \leq 50 \text{ mm}$	$ z_0 \leq 50 \text{ mm}$
$0 \leq \chi^2 \leq 15$	$0 \leq \chi^2 \leq 30$
$ \lambda_{01} \geq 0.1 \text{ rad}$	$ \lambda_{01} \geq 0.1 \text{ rad}$

Table 5.2: Data selection cuts as applied in the sensitivity studies to $\mu \rightarrow eX$ decays.

cient longitudinal momentum to overshoot the recurl stations. Thus, an inefficiency arises. The shoulder at about 20 MeV arises in a similar way. Here, sufficient longitudinal momentum can be obtained to reach the non-instrumented gap between the central detector station and the recurl stations.

It is particularly notable that the cuts do not create any structure for fake tracks that could be mistaken for a peak of $\mu \rightarrow eX$ events. The same holds for electron contamination and for the various contributions of SM muon decays (see figure 5.12). This means that the sensitivities to $\mu \rightarrow eX$ would not significantly deviate from the ones estimated in this study even if these effects are slightly over- or underestimated.

5.4 EXPECTED BRANCHING FRACTION LIMITS

The procedure described in section 5.2.1 is applied to $\mu \rightarrow eX$ events with m_X between 0 MeV and 99 MeV under the assumption that an alternative calibration is applied that does not rely on the edge of the Michel spectrum. The expected limits on the $\mu \rightarrow eX$ branching fraction are given for the 90 %, 95 %, and 99 % confidence level, as well as the 5σ level that is usually required to claim a discovery. Figure 5.13 shows the expected limits⁴ for short and long track reconstruction that can be reached with the phase I *Mu3e* experiment after a total run time of 300 days at a muon stopping rate of $10^8 \mu/s$.


SHORT TRACK RECONSTRUCTION In the case of reconstructed short tracks, the expected limits on the branching fraction of $\mu \rightarrow eX$ decays range between $1 \cdot 10^{-8}$ and $2 \cdot 10^{-8}$ at 90 % CL for $m_X < 90$ MeV. The steep rise for $m_X > 90$ MeV is due to the low efficiency to reconstruct tracks with low momenta. The number $N_{\text{rec}}^{\mu \rightarrow eX}$ of recon-

⁴ Here and in the following, only statistical uncertainties are quoted. Systematic uncertainties are studied separately in section 5.7.



In all the following plots of the expected branching fraction limits, the data points are simply connected via straight lines.

Reconstruction		No cuts	Cuts applied
SM background			
Short	ϵ	1	0.845
	p	0.789	0.955
	$\epsilon \cdot p$	0.789	0.806
Long	ϵ	1	0.824
	p	0.902	0.960
	$\epsilon \cdot p$	0.902	0.791
$\mu \rightarrow eX$ signal			
Short	ϵ	1	0.845
	p	0.802	0.961
	$\epsilon \cdot p$	0.802	0.812
Long	ϵ	1	0.814
	p	0.907	0.963
	$\epsilon \cdot p$	0.907	0.784

Table 5.3: Efficiency ϵ , purity p and $\epsilon \cdot p$ before and after applying cuts for SM processes and $\mu \rightarrow eX$ decays with m_X ranging from 0 MeV to 99 MeV. The statistical uncertainty is on the 10^{-4} level.



(a) SM processes reconstructed as short tracks.

(b) $\mu \rightarrow eX$ events reconstructed as short tracks.

(c) SM processes reconstructed as long tracks.


(d) $\mu \rightarrow eX$ events reconstructed as long tracks.

Figure 5.11: Distribution of the reconstructed momentum of positron candidates in SM processes and $\mu \rightarrow eX$ decays with a mass m_X of 60 MeV before and after applying cuts as outlined in table 5.2. Reconstructed true tracks are shown in blue, fake tracks in red and the total reconstructed tracks in black. A dashed line indicates the distribution before applying cuts, and a solid line after applying cuts. The sample contains about $5 \cdot 10^8$ positrons from SM muon decays and 75 000 $\mu \rightarrow eX$ decays, respectively.

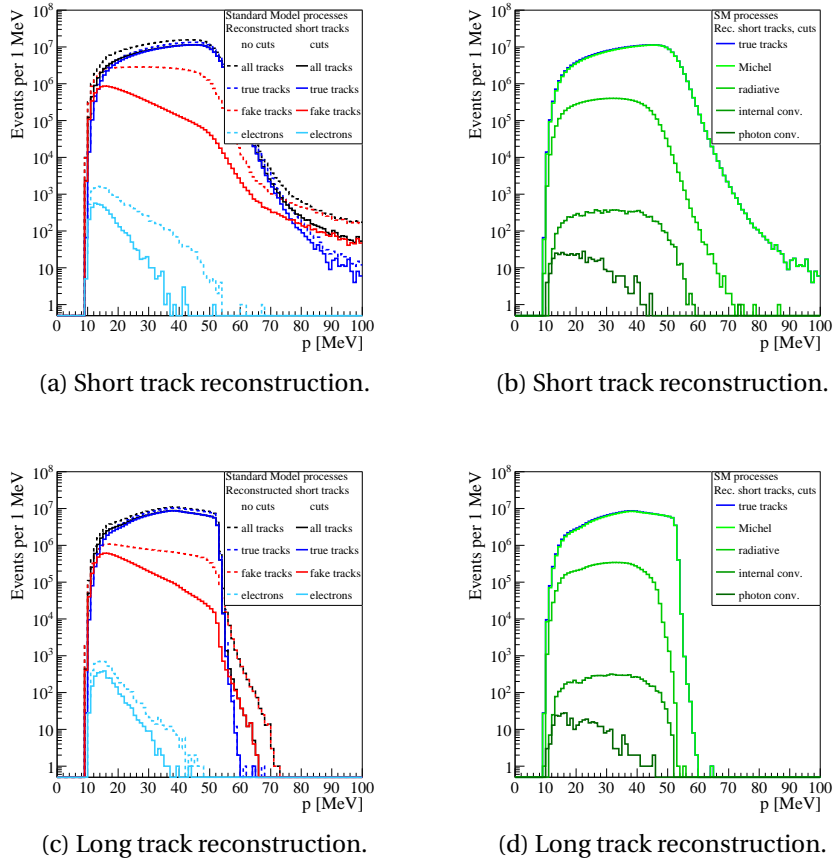



Figure 5.12: Distribution of the reconstructed momentum of positron candidates in SM processes with cuts applied as outlined in table 5.2. On the left-hand side, the contamination from electrons (light blue) is shown together with true (blue) and fake tracks (red). Dashed lines correspond to no cuts being applied, solid lines to applied cuts. On the right-hand side, the contributions of positrons from various SM muon decay modes are displayed. The distributions are given for true tracks with cuts applied. The sample contains about $5 \cdot 10^8$ positrons from SM muon decays.

structed $\mu \rightarrow eX$ events becomes small and thus the sensitivity rises according to equation 5.13. This effect is observed in all reconstruction modes (see figure 5.14b). Hence, in the following only masses smaller or equal to 90 MeV are considered.


In the range around $m_X \approx 40$ MeV, a hump becomes visible in the branching fraction limits. Translated to momentum, this corresponds to the region around the maximum of the Michel spectrum. As expected, a larger number of background events in the range of the signal peak results in weaker expected branching fraction limits. In the context of the toy Monte Carlo study, the width σ_{toyMC} of the signal fraction distribution becomes larger. Towards large m_X the limits improve because of the decline of the Michel spectrum at small momenta, and towards small m_X because of the Michel edge. On average for m_X between 0 MeV and 90 MeV, an expected limit on the branching fraction for $\mu \rightarrow eX$ decays of $1.48 \cdot 10^{-8}$ at 90 % CL can be achieved (see table 5.4).

LONG TRACK RECONSTRUCTION For reconstructed long tracks, limits on the branching fraction for $\mu \rightarrow eX$ around $1 \cdot 10^{-8}$ can be expected. The hump observed for short tracks is shifted towards smaller m_X as a result of the better momentum resolution. The reconstruction efficiency for $\mu \rightarrow eX$ events is shown in figure 5.14b. It increases towards 50 MeV, drops off at about 80 MeV before falling to zero after 90 MeV. The first rise is due to the geometric acceptance for recurling tracks. At higher momentum, tracks have a larger chance to overshoot the recurl stations without touching them, and thus the efficiency is reduced. The drop at 80 MeV is due to the non-instrumented gap between the central detector station and the recurl stations. Reconstructable low-momentum particles have mostly transverse momentum. Thus, the longitudinal momentum does not suffice to reach the gap. For higher momenta, tracks can also be reconstructed when a larger portion of the momentum is longitudinal. Therefore, the acceptance for recurling tracks drops as soon as enough longitudinal momentum can be afforded.

For masses between 0 MeV and 90 MeV, an average limit on the branching fraction for $\mu \rightarrow eX$ decays of $0.73 \cdot 10^{-8}$ at 90 % CL is expected. This is an improvement by about a factor of two compared to the sensitivity using short tracks due to the improved momentum resolution.




(a) Limits on the $\mu \rightarrow eX$ branching fraction for short track reconstruction.




(b) Limits on the $\mu \rightarrow eX$ branching fraction for long track reconstruction.

Figure 5.13: Limits on the $\mu \rightarrow eX$ branching fraction for 300 days of data taking at $10^8 \mu/s$ for reconstructed short and long tracks.



(a) Width σ_{toyMC} of the distribution of the signal fraction in toy Monte Carlo studies.



(b) Fraction of reconstructed tracks per $\mu \rightarrow eX$ decay.

Figure 5.14: Width σ_{toyMC} of the distribution of the signal fraction in toy Monte Carlo studies and fraction of reconstructed tracks per $\mu \rightarrow eX$ decay for various masses m_X .

Reconstruction	CL	$\mathcal{B} [10^{-8}]$
Short	90 %CL	1.48
	95 %CL	1.90
	99 %CL	2.69
	5σ	5.79
Long	90 %CL	0.73
	95 %CL	0.94
	99 %CL	1.33
	5σ	2.86

Table 5.4: Average expected limits on the branching fraction for $\mu \rightarrow eX$ decays in 300 days of data taking at $10^8 \mu/s$ for $0 \text{ MeV} \leq m_X \leq 90 \text{ MeV}$. The relative statistical uncertainty is on the 10^{-3} level.




Figure 5.15: Differential decay rate of SM Michel decays of the muon in positron energy E_e and cosine of the emission angle ϑ with respect to the muon polarization as given in equation 5.18. The muons are assumed to be fully polarized.

5.5 ANALYSIS IN MOMENTUM AND POLAR ANGLE

SM muon decays exhibit an angular dependence with respect to the polarization of the muon, whereas the positrons of $\mu \rightarrow eX$ decays are isotropically emitted in scalar interactions. The distribution of positrons from the main Michel decay mode of the muon dominates the spectrum. At leading order, the differential decay rate for $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$ in the SM becomes [46]

$$\frac{d^2\Gamma(\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e)}{dx d\cos\vartheta} = \frac{m_\mu^5 G_F^2}{192\pi^3} x^2 [(3 - 2x) + P \cos\vartheta (2x - 1)] \quad (5.18)$$

assuming the polarization of the positrons is not measured. Herein, G_F denotes the Fermi coupling constant, P the polarization of the muon, and ϑ the angle between the positron momentum and the muon polarization vector. x is the reduced positron energy

$$x = \frac{2m_\mu E_e}{m_\mu^2 + m_e^2}. \quad (5.19)$$

The differential decay rate is shown in figure 5.15.

The sensitivity to $\mu \rightarrow eX$ can thus be improved by additionally taking $\cos\vartheta$ into account. Furthermore, $\mu \rightarrow eX$ decays with pseudo-scalar couplings can be investigated.

In the following, θ denotes the polar angle i. e. the angle between the positron momentum and the beam axis. The muons are assumed to be polarized parallel to the beam axis with the spins pointing upstream, therefore $\theta = -\vartheta$.

5.5.1 Event Distributions

In this study, a fully polarized muon beam is assumed. The polarization is pointing upstream as a result of the production of μ^+ from the pion decay $\pi^+ \rightarrow \mu^+ \nu_\mu$. As before, it is assumed that the Michel edge is not used in the momentum calibration.

The sensitivities to $\mu \rightarrow eX$ decays are obtained in a similar fashion as described in section 5.2.1, with the only difference of using two-dimensional input histograms instead of one-dimensional ones. The histograms are filled with the reconstructed momentum and the cosine of the polar angle θ . As a large number of bins increases the run time of the toy Monte Carlo studies, the binning in momentum is adapted wherever applicable. The reconstructed short tracks have a larger momentum resolution than the long tracks allowing for a coarser binning of 1 MeV. For long tracks, the momentum is binned

in 100 keV bins. As shown in section 5.7, this choice for the momentum bin size hardly affects the expected branching fraction limits. For $\cos\theta$, a bin size of 0.1 is chosen for both reconstruction modes. In addition, the number of toy Monte Carlo samples is reduced to 1000 but the number of 10^9 events per sample remains unchanged. As before, smoothing is applied to the SM histograms.

Figure 5.16 shows the input histograms for SM processes and $\mu \rightarrow eX$ decays with a mass of 60 MeV for the short track reconstruction. In the case of $\mu \rightarrow eX$ decays, scalar and pseudo-scalar interactions are distinguished by the coupling parameter h as used in equation 5.17. In the following, scalar interactions ($h = 0$) as well as pseudo-scalar interactions with couplings $h = +1$ and $h = -1$ are studied.

In the case of SM processes, the events accumulate at momenta of about 50 MeV—the edge of the Michel spectrum—and large negative $\cos\theta$. The left border of the distribution is determined by the minimum transverse momentum of about 10 MeV that is required to reconstruct a track. The $\mu \rightarrow eX$ events are mostly contained in a small band in momentum. As expected, the events are evenly distributed in this band for scalar interactions, whereas positive $\cos\theta$ are preferred in the case of $h = +1$, and negative $\cos\theta$ in the case of $h = -1$.

The input histograms for long track reconstruction are shown in figure 5.17. The geometric acceptance for long track reconstruction becomes apparent in the distribution of SM muon decays. In addition to the minimum transverse momentum threshold, there are no entries at high momenta and large $|\cos\theta|$ because these particles reach further than the recurl stations. The gap between the central station and the recurl stations cast a shadow on the remaining distribution. Due to the excellent momentum resolution, the band in the $\mu \rightarrow eX$ distributions is very narrow.

In all histograms, the selection cut on λ_{01} is clearly visible. A few entries remain around $\cos\theta = 0$ as the bin width is larger than the λ_{01} cut, and because of smoothing in the case of SM processes.

5.5.2 *Expected Branching Fraction Limits*

In figure 5.18, the expected $\mu \rightarrow eX$ branching fraction limits at 90% CL as derived with the two-dimensional method are shown for the investigated couplings. As before, the limits on the branching fraction for $\mu \rightarrow eX$ decays are given for a total of 300 days of data taking at a muon stopping rate of $10^8 \mu/s$. The average expected limits are summarized in table 5.5.





Figure 5.16: Histograms of the reconstructed momentum and $\cos\theta$ for SM processes and $\mu \rightarrow eX$ decays with $m_X = 60 \text{ MeV}$ and various couplings h using reconstructed short tracks. About $5 \cdot 10^8$ SM events and 75 000 $\mu \rightarrow eX$ decays are generated in the target region.



(a) SM processes.




(b) $\mu \rightarrow eX$ signal events with $m_X = 60 \text{ MeV}$ and $h = 0$.(c) $\mu \rightarrow eX$ signal events with $m_X = 60 \text{ MeV}$ and $h = +1$.(d) $\mu \rightarrow eX$ signal events with $m_X = 60 \text{ MeV}$ and $h = -1$.

Figure 5.17: Histograms of the reconstructed momentum and $\cos\theta$ for SM processes and $\mu \rightarrow eX$ decays with $m_X = 60 \text{ MeV}$ and various couplings h using reconstructed long tracks. About $5 \cdot 10^8$ SM events and 75 000 $\mu \rightarrow eX$ decays are generated in the target region.

COUPLING $h = 0$ The limit curves for $h = 0$ have a similar shape and magnitude as the curves obtained with the one-dimensional method (see section 5.4). Because of the smaller sample size and, in the case of short track reconstruction, also coarser momentum bins, the curves for the two-dimensional histograms are however not as smooth and the uncertainties are larger but still acceptable.

For a mass range of $0 \text{ MeV} \leq m_X \leq 90 \text{ MeV}$, the expected average limit at 90 % CL is $1.33 \cdot 10^{-8}$ for reconstructed short tracks and $0.65 \cdot 10^{-8}$ for reconstructed long tracks, again about a factor of two smaller like in the one-dimensional method.


COUPLING $h = +1$ For a coupling $h = +1$, notably smaller limits are achieved than for scalar couplings. This results from the fact, that for $h = +1$ a $\cos \theta > 0$ is preferred—a region that is disfavoured by the SM muon decays.

The short track reconstruction yields an expected average limit at 90 % CL of $1.04 \cdot 10^{-8}$ for masses $0 \text{ MeV} \leq m_X \leq 90 \text{ MeV}$. For reconstructed long tracks, the limit is $0.55 \cdot 10^{-8}$ on average. The limits are 15 % to 22 % smaller than for $h = 0$.


COUPLING $h = -1$ The branching fraction limits in the case of $h = -1$ are comparable to the case of $h = 0$. At smaller positron momenta and thus larger m_X , the SM events tend to be more evenly distributed in $\cos \theta$ similar to the $\mu \rightarrow eX$ events with $h = 0$. For larger momenta, the SM decays prefer $\cos \theta < 0$ just like $\mu \rightarrow eX$ events with $h = -1$ (see figure 5.15). This is also observed in the sensitivity. For small m_X , smaller expected limits are obtained for scalar couplings, and for large m_X the limits for $h = -1$ are slightly better.

On average for masses of $0 \text{ MeV} \leq m_X \leq 90 \text{ MeV}$, the expected branching fraction limit at 90 % CL for $\mu \rightarrow eX$ decays with $h = -1$ becomes $1.36 \cdot 10^{-8}$ for short track reconstruction and $0.65 \cdot 10^{-8}$ for long track reconstruction.

COMPARING THE ONE- AND TWO-DIMENSIONAL METHOD In figure 5.19, the expected 90 % CL limits on the branching fraction for $\mu \rightarrow eX$ decays with scalar coupling derived with the one- and two-dimensional method are compared. It can be noticed, that the two-dimensional method yields smaller limits over the whole range in m_X . The amelioration is larger in regions where the asymmetry between $\cos \theta > 0$ and $\cos \theta < 0$ for SM processes is more pronounced,




(a) Limits on the $\mu \rightarrow eX$ branching fraction at 90 % CL for short track reconstruction.




(b) Limits on the $\mu \rightarrow eX$ branching fraction at 90 % CL for long track reconstruction.

Figure 5.18: Comparison of the expected limits on the $\mu \rightarrow eX$ branching fraction at 90 % CL for different couplings h obtained with the two-dimensional method assuming 300 days of data taking at $10^8 \mu/s$ using reconstructed short and long tracks.



(a) Limits on the $\mu \rightarrow eX$ branching fraction for short track reconstruction.



(b) Limits on the $\mu \rightarrow eX$ branching fraction for long track reconstruction.

Figure 5.19: Comparison of the expected branching fraction limits for scalar $\mu \rightarrow eX$ decays at 90% CL using either the one- or two-dimensional method. A data taking period of 300 days at a rate of $10^8 \mu/s$ is assumed.

Coupling	Reconstruction	$\mathcal{B} [10^{-8}]$	$\mathcal{B} [10^{-8}]$
		One-Dimensional	Two-Dimensional
$h = 0$	Short	1.48	1.33
	Long	0.73	0.65
$h = +1$	Short		1.04
	Long		0.55
$h = -1$	Short		1.36
	Long		0.65

Table 5.5: Average expected limits on the branching fraction for $\mu \rightarrow eX$ decays at 90 % CL derived using the two-dimensional method. For a coupling of $h = 0$, the limits of the one-dimensional method are given for comparison. 300 days of data taking at $10^8 \mu/s$ are assumed. The relative statistical uncertainty is on the 10^{-3} level.

i. e. at smaller masses. The limits are improved by 10 % to 12 % (see table 5.5).

5.6 MOMENTUM CALIBRATION

A proper calibration of the momentum scale is important to $\mu \rightarrow eX$ measurements. The current concept foresees to use the characteristic edge in the momentum spectrum of positrons from Michel decays for the calibration of online and offline reconstructed tracks [136]. This would however not allow to test $\mu \rightarrow eX$ decays with small masses m_X as these decays would affect the position and shape of the Michel edge. In the following, a procedure is discussed in which the momentum is calibrated alongside $\mu \rightarrow eX$ searches.

In principle, the same momentum spectrum can be used for both purposes, momentum calibration and peak search, as long as the signal region is left out during the calibration fit. This means, a calibration has to be performed for every m_X individually. In figure 5.20, the procedure is shown exemplarily for reconstructed short tracks and two different masses. In this toy Monte Carlo study, a sample of 10^9 events is drawn from the background p.d.f. In a first step, a mod-

ified background p.d.f. is fitted to the events (calibration fit⁵). The modified p.d.f. can be stretched or compressed with a free scaling parameter x_{scale} with respect to the unmodified p.d.f., i. e.

$$f_b^{\text{mod}}(x_{\text{scale}} \cdot p) \propto f_b(p) . \quad (5.20)$$

For this purpose, the momentum histogram of SM processes, that is the basis of both background p.d.f.s, is not only smoothed as described in section 5.2.1, but also an interpolation of second order is applied to even out the bin edges.

Events in a region of two root mean square width below and above the mean momentum of the positron of the investigated $\mu \rightarrow eX$ decay are left out during the calibration fit⁶.

Subsequently, a fit of the signal-plus-background p.d.f. is performed on all events. Signal and background p.d.f. are modified with the scaling parameter being fixed to the result of the previous fit. This procedure is repeated for a total of 1000 samples. The branching fraction limits are derived from the width of the signal fraction as before.

As shown in figure 5.20, the mean fitted scaling parameter is compatible with unity as expected. At a value of one, the initial and modified p.d.f. are the same. The actual spread of the fitted scaling parameter in this study is not very meaningful as it depends on the number of toy Monte Carlo events which is orders of magnitude below the expected number of observed muon decays. Thus, only qualitative conclusions can be drawn. Of course, the width of the scaling parameter distribution also depends on the size of the left-out momentum interval. In the definition chosen here, this interval is determined by the momentum resolution.

SHORT TRACK RECONSTRUCTION In figure 5.21, the expected $\mu \rightarrow eX$ branching fraction limits at 90 % CL are shown assuming 300 days of data taking at 10^8 muons/s. In a mass range from about 40 MeV to 90 MeV, the limits are marginally larger than the limits obtained in section 5.4 despite the additional scaling parameter. In this range, the momentum range that is left out for the calibration fit is

5 In contrast to the fits performed before that rely on the maximum likelihood method, here the fit is performed using the method of least squares. The maximum likelihood is not applicable as the modification of the histogram can result in a p.d.f. that is zero at the edges of the momentum range.

6 In fact, also structures in the momentum spectrum should be left out in the calibration fit that are due to the acceptance of the detector. This includes the exact position of the low-momentum threshold and the kink at low momenta caused by the incapability to measure transverse momenta below 10 MeV, as well as the region of inefficiency to reconstruct long tracks at high momenta.

located on the rising edge of the spectrum. As there are no sudden changes in the slope, the shape of the spectrum can be precisely fitted even with the left-out range. Therefore, the distribution of the fitted scaling parameter is narrow (see figure 5.20).

For smaller masses, the left-out range extends to the Michel edge. Omitting this characteristic feature of the spectrum increases the uncertainty by which the scaling parameter is determined. Consequently, the distribution of the fitted signal fraction broadens and larger expected branching fraction limits are derived. Also, the fluctuations of the limits are stronger. The mass range at which this turnover happens is a result of the choice of the left out fit range.

For large m_X , the left-out range spans over the low momentum threshold of the track reconstruction. At this point, the fit is performed on a single range and not on two separate ranges as in the case of smaller masses. Thus, the uncertainty of the fitted scaling parameter increases as well.


LONG TRACK RECONSTRUCTION For reconstructed long tracks, the expected limits exhibit a similar behaviour. As the momentum resolution is smaller, the Michel edge is steeper and the left-out range smaller. Therefore, the limits deviate from the ones derived in section 5.4 only up to masses of about 20 MeV.

The expected limits obtained with the simultaneous calibration fit are regarded to be a conservative estimate. The number of events is limited in the toy Monte Carlo study. With a larger number of events in the later experiment, the precision of the calibration fit is expected to improve. In figure 5.21, the expected limits from section 5.4 are shown in addition. These limits can be regarded as the best possible limits given an alternative calibration method can be found.


5.6.1 *Validation of the Method*

Uncertainties in the magnetic field or alignment of the detector result in an uncertainty in the measurement of the radius and thus momentum of the track. The calibration fit is required to function reliably under these conditions. This is investigated in the following.


For this purpose, the p.d.f. which is used to generate the events is modified. The initial histogram as used in section 5.4 is stretched or compressed with a scaling factor of $1 \pm \epsilon$. For ϵ , a value of $1 \cdot 10^{-4}$ is chosen which corresponds to the specified long term stability of




(a) Fit of the background p.d.f. to a single sample. A signal region corresponding to $\mu \rightarrow eX$ events with $m_X = 10\text{ MeV}$ is left out.



(b) Distribution of the fitted scaling parameter x_{scale} for $\mu \rightarrow eX$ searches with $m_X = 10\text{ MeV}$.



(c) Fit of the background p.d.f. to a single sample. A signal region corresponding to $\mu \rightarrow eX$ events with $m_X = 60\text{ MeV}$ is left out.



(d) Distribution of the fitted scaling parameter x_{scale} for $\mu \rightarrow eX$ searches with $m_X = 60\text{ MeV}$.

Figure 5.20: Simultaneous determination of the momentum calibration for reconstructed short tracks. On the left-hand side, a fit of the background p.d.f. (green line) is shown to a single sample (black dots) of 10^9 events derived from the background p.d.f. (blue line). On the right-hand side, the distribution of the fitted scaling parameter is shown for 1000 fits. Similar plots for long track reconstruction can be found in section C.3 in the appendix.




Figure 5.21: Expected 90 % confidence limits on the $\mu \rightarrow eX$ branching fraction for a simultaneously extracted momentum calibration assuming 300 days of data taking at $10^8 \mu/s$.


the magnetic field, as well as a value of $1 \cdot 10^{-3}$ in order to study even larger errors of the calibration. An interpolation of second order is applied to the histograms before modification in order to even out the bin edges. From this modified background p.d.f., the events for the toy Monte Carlo study are sampled.

Then, a calibration fit with x_{scale} as free fit parameter and a subsequent signal-plus-background fit is performed as described before in section 5.6. Again, events in a region of two root mean square width below and above the mean momentum of the positron of the investigated $\mu \rightarrow eX$ decay are left out in the calibration fit.


In figure 5.22, the expected branching fraction limits at 90 % CL are shown for the full range of masses m_X . The initial rescaling has no perceivable influence on the derived expected branching fraction limits. The limits are only influenced by the fluctuations of the scaling parameter. The method is thus regarded to be robust against the expected uncertainties of the magnetic field and alignment.

5.7 MOMENTUM BIN SIZE

The choice of the bin size for the histogram of the reconstructed momentum affects the sensitivity as soon as it is of comparable size as the momentum resolution. The bin size for the online histogram of the reconstructed momentum is not yet specified. Besides the momentum, also the polar and azimuthal angle need to be stored. This allows for a position dependent calibration. In addition, the polarization direction is rotated by a few degrees with respect to the



(a) Limits on the $\mu \rightarrow eX$ branching fraction for short track reconstruction.



(b) Limits on the $\mu \rightarrow eX$ branching fraction for long track reconstruction.

Figure 5.22: Expected branching fraction limits at 90 % CL on the $\mu \rightarrow eX$ decays assuming 300 days of data taking at $10^8 \mu/s$ that are expected when the calibration is simultaneously extracted from the momentum spectrum for various rescaled spectra. The limits for the case that the calibration is determined externally (ext. calibration) are given for reference.

beam axis [202]. Additional information about the azimuthal angle is thus required for the analysis of pseudo-scalar couplings.

Considerations about the technical realization and limitations of the online reconstruction on GPUs can be found in [157]. In general, a large number of bins is beneficial for the performance of the online reconstruction on the GPU as it decreases the probability that two threads need to access the same bin at the same time which causes latency. The number of bins is thus only limited by the available memory. In the order of 10^5 bins can be realized which translates to a few 10 to about 100 bins per variable. Thus, the maximum momentum bin size that does not affect the sensitivity of $\mu \rightarrow eX$ searches needs to be identified.


In the following, a set of fine to coarse bin sizes ranging from 50 keV to 2 MeV, and additionally 5 MeV and 10 MeV for short tracks, is studied. In figure 5.23, the SM histogram is shown in two exemplary bin sizes.

For coarse binnings, it can be observed that the $\mu \rightarrow eX$ events are either distributed over two bins or are mostly contained in a single bin. This is an artefact of the arbitrary choice of the bin edges and affects the sensitivity limits. In order to derive comparable limits for all masses and bin sizes, the bin edges in this study are chosen in such a way that the maximum of the $\mu \rightarrow eX$ peak is always located in the centre of one bin.


The limits on the $\mu \rightarrow eX$ branching fraction are derived in the same way as in section 5.4. Figure 5.24 shows the expected limits at 90 % CL for individual masses m_X for the three reconstruction modes. Although the bin edges are adapted, the limit curves for large bin sizes have some wiggles. These are caused by a varying ratio of the heights of the central and neighbouring bins.

In figure 5.25 and table 5.6 average limits for masses between 0 MeV and 90 MeV are presented. The results are given for 300 days of data taking at $10^8 \mu/s$.


SHORT TRACK RECONSTRUCTION The momentum resolution of short tracks ranges from 0.5 MeV for tracks with low momenta up to 3 MeV for high momenta (see figure 2.18b). This is reflected in the expected branching fraction limits. For bin sizes up to 2 MeV, the limits are basically unaffected, but for larger bin sizes a degradation becomes apparent. On average for masses m_X between 0 MeV and 90 MeV, the expected limits are $1.7 \cdot 10^{-8}$ for 5 MeV bins and $2.1 \cdot 10^{-8}$ for 10 MeV bins, whilst for smaller bin sizes limits of $1.5 \cdot 10^{-8}$ are expected.




(a) Background p.d.f. with 50 keV bin size for reconstructed short tracks.



(b) Background p.d.f. with 2 MeV bin size for reconstructed short tracks.



(c) Background p.d.f. with 50 keV bin size for reconstructed long tracks.



(d) Background p.d.f. with 2 MeV bin size for reconstructed long tracks.

Figure 5.23: Momentum spectrum of SM muon decays (background p.d.f.) with bin sizes of 50 keV and 2 MeV.


LONG TRACK RECONSTRUCTION Long tracks are reconstructed with a momentum resolution ranging from 0.1 MeV to 0.45 MeV, with the best resolution being achieved for 20 MeV tracks (see figure 2.18d). Therefore, the $\mu \rightarrow eX$ analysis is more sensitive to the choice of the bin size in the case of long tracks.

Up to 0.5 MeV, the bin size barely influences the expected branching fraction limit. But with larger bin sizes the limits deviate increasingly. For bin sizes smaller or equal to 0.5 MeV, the expected limits remain below $0.8 \cdot 10^{-8}$ on average for masses between 0 MeV and 90 MeV. At a bin size of 2 MeV, a limit of $1.1 \cdot 10^{-8}$ is expected.


5.8 COMPARISON TO EXISTING EXPERIMENTAL LIMITS

The latest and most stringent results on $\mu \rightarrow eX$ decays with massive bosons X are obtained by the *TWIST* experiment (see section 1.3.1.2). On average, $\mu \rightarrow eX$ decays are excluded with an upper 90 % confidence limit of $9 \cdot 10^{-6}$ [108].

In figure 5.26, the expected $\mu \rightarrow eX$ branching fraction limits at 90 % CL are put into comparison with the upper limits of the *TWIST* experiment. Again, 300 days of data taking at $10^8 \mu/s$ are assumed for



(a) Limits on the $\mu \rightarrow eX$ branching fraction for short track reconstruction.



(b) Limits on the $\mu \rightarrow eX$ branching fraction for long track reconstruction.

Figure 5.24: Expected branching fraction limits at 90% CL on the $\mu \rightarrow eX$ decays for various bin sizes assuming 300 days of data taking at $10^8 \mu/s$.

Bin size [MeV]	$\mathcal{B} [10^{-8}]$	$\mathcal{B} [10^{-8}]$
	Short tracks	Long tracks
0.05	1.48	0.73
0.10	1.49	0.73
0.20	1.49	0.74
0.50	1.49	0.78
1.00	1.50	0.89
2.00	1.52	1.14
5.00	1.68	—
10.00	2.09	—

Table 5.6: Average expected branching fraction limits at 90 % CL on $\mu \rightarrow eX$ decays with masses ranging from 0 MeV to 90 MeV for various bin sizes assuming 300 days of data taking at $10^8 \mu/s$.




Figure 5.25: Average branching fraction limits at 90 % CL on $\mu \rightarrow eX$ decays with masses ranging from 0 MeV to 90 MeV for various bin sizes.





Figure 5.26: Expected 90 % confidence limits on the $\mu \rightarrow eX$ branching fraction for a simultaneously extracted momentum calibration assuming 300 days of data taking at $10^8 \mu/s$ in comparison to the results by the *TWIST* experiment [108]. *TWIST* results by courtesy of R. Bayes.


the phase I *Mu3e* experiment. Due to the significantly larger number of muon decays, *Mu3e* is highly competitive to *TWIST*. Using short track reconstruction, *Mu3e* is sensitive to branching fractions about a factor of 600 smaller than investigated by *TWIST*. Using long track reconstruction, the factor becomes even larger than a thousand.

Given the Michel edge is used in momentum calibration, the sensitivity is reduced for small masses m_X . The expected limits obtained in section 5.6 are regarded to be conservative estimates as the precision of the calibration fit is expected to improve with a larger number of events. Even with the reduced sensitivity at small masses, an improvement of two orders of magnitude can be expected by the *Mu3e* experiment compared to *TWIST*.


TWIST has also investigated pseudo-scalar couplings. A comparison to the expected limits for *Mu3e* is shown in figure 5.27. *TWIST* has set average 90 % CL upper limits on $\mu \rightarrow eX$ decays of $9 \cdot 10^{-6}$ for $h = 0$, $6 \cdot 10^{-6}$ for $h = +1$, and $10 \cdot 10^{-6}$ for $h = -1$ [108]. In the case of *Mu3e*, 90 % CL branching fraction limits of $1.3 \cdot 10^{-8}$ ($6.5 \cdot 10^{-9}$) for $h = 0$, $1.0 \cdot 10^{-8}$ ($5.5 \cdot 10^{-9}$) for $h = +1$, and $1.4 \cdot 10^{-8}$ ($6.5 \cdot 10^{-9}$) for $h = -1$ are expected using short (long) track reconstruction and two-dimensional histograms. As in the case of the one-dimensional method, the expected sensitivities of *Mu3e* are about a factor of 600 smaller for short track reconstruction, and more than a factor of 1000 for long tracks.



(a) Limits on the $\mu \rightarrow eX$ branching fraction for $h = 0$.



(b) Limits on the $\mu \rightarrow eX$ branching fraction for $h = +1$.



(c) Limits on the $\mu \rightarrow eX$ branching fraction for $h = -1$.

Figure 5.27: Comparison of the expected limits on the $\mu \rightarrow eX$ branching fraction at 90 % CL for different couplings h assuming 300 days of data taking at $10^8 \mu/s$ with the results of the *TWIST* experiment [108]. *TWIST* results by courtesy of R. Bayes.

5.9 SUMMARY ON LEPTON FLAVOUR VIOLATING TWO-BODY DECAYS

SUMMARY In this chapter, a strategy is presented to search for lepton flavour violating two-body decays of the muon $\mu \rightarrow eX$ by searching for an excess on top of the momentum spectrum of positrons from anti-muon decays.

The accessible mass range is determined by the chosen method of momentum calibration. If the characteristic end-point of the spectrum of Michel decays $\mu \rightarrow e\nu\nu$ is used, small masses m_X have to be excluded up to about 25 MeV. Masses larger than 95 MeV cannot be studied as positrons with transverse momentum below 10 MeV are outside of the experimental acceptance. In fact, the efficiency for $m_X > 90$ MeV steeply declines.

The expected limits on the $\mu \rightarrow eX$ branching fraction at 90% CL range around 10^{-8} assuming a data taking period of 300 days at a muon stopping rate of $10^8 \mu/s$ in phase I. This is a result of the good momentum resolution and high acceptance as well as the large number of muon stops. *Mu3e* has thus the potential to improve the existing upper limits of the *TWIST* experiment by two to three orders of magnitude.

Along with the sensitivity studies, possible improvements to the procedure are evaluated.

The input data to the $\mu \rightarrow eX$ studies are histograms of the reconstructed momentum in the online reconstruction. In the baseline design of *Mu3e*, tracks are reconstructed as short tracks. With an extension to the data acquisition, also long tracks could be reconstructed. The improved momentum resolution would allow to test branching fractions that are up to a factor of two smaller.

In addition to the momentum, also the polar angle of the positron can be stored. As the positron in $\mu \rightarrow eX$ with scalar X is emitted isotropically in contrast to positrons from $\mu \rightarrow e\nu\nu$, the sensitivity improves by about 10% in this approach. Furthermore, possible pseudo-scalar couplings can be tested. The best expected limits are achieved when the positron is emitted in forward direction, i. e. opposite to the dominant direction of positrons from Michel decays.

If the shape of the momentum spectrum of SM muon decays is the only possibility to calibrate the momentum measurement, it is still possible to extract the calibration simultaneously to $\mu \rightarrow eX$ searches and thus measure also at small m_X . Here, a signal region is left out in the calibration fit. As soon as the left-out region extends

to the end-point, the uncertainty of the calibration increases. Therefore, the sensitivity at small m_X worsens.

Uncertainties in the magnetic field or alignment translate into uncertainties of the momentum measurement. It is observed that a relative uncertainty of up to $1 \cdot 10^{-3}$ in the momentum measurement does not pose a difficulty to the simultaneous calibration fit. The results obtained with a slightly mis-measured momentum are consistent to the results without this modification.

Furthermore, the influence of the bin size of the reconstructed momentum on the sensitivity is investigated. The available memory on the filter farm GPUs puts limitations on the bin size in particular if also the polar and azimuthal angle are stored for calibration purposes. The bin size can deteriorate the precision of $\mu \rightarrow eX$ searches as soon as it is comparable to the momentum resolution. For short tracks, deviations are observed for bin sizes larger than 2 MeV. For long tracks, the limits are already affected at a bin size of 0.5 MeV.

DISCUSSION AND OUTLOOK As the detector simulation incorporates some idealizations and simplifications, using fixed momentum spectrum shapes from simulation for the signal and background distributions is not suited for the later data analysis. Consequently, the analysis strategy has to be adapted. The search will not be conducted on the whole momentum spectrum but on a series of smaller search windows. The range of the search window should be chosen such that the background distribution is smooth within that range. Signal and background are then parametrized by a function. A polynomial for the background and a Crystal Ball function for the signal might be a good choice. In this way, the shape of the background function can adapt to some extent to the measured data and thus overcome deficiencies of the simulation. Furthermore, the mass can be treated as a continuous variable. Difficulties might arise at the Michel edge. The steep decline in background would favour a small search window. But at the same time, the search window needs to be wide enough to cover the full width of the $\mu \rightarrow eX$ signal with some safety margin.

If the reconstruction of long tracks on the filter farm is realized, the acceptance loss for long tracks compared to short tracks can be overcome. The acceptance loss stems from high momentum particles that overshoot the recurl stations as well as particles that cross the non-instrumented gaps. In these cases, a short track can be kept. The short and long track sample can then be combined. A weight

proportional to the inverse of the momentum resolution should be applied in order to account for the significantly better momentum resolution of long tracks.

The momentum resolution of reconstructed long tracks is best when the particle performs about a half turn between the two crossings of the outer layers. The momentum at which the resolution is optimum can be shifted with the magnetic field. For the nominal magnetic field of 1 T, this is around 20 MeV. With a higher magnetic field, this point is shifted to higher momenta. Thus, operating at the maximum specified magnetic field of 2 T can be used to improve the sensitivity at low m_X when the momentum calibration is simultaneously performed to the peak search. In the *Mu3e* experiment, a change in the magnetic field is in conflict with the acceptance for $\mu \rightarrow eee$ searches. The potential *Mu3e-Gamma* experiment in contrast is foreseen to operate at 2 T and can thus complement the $\mu \rightarrow eX$ searches at small m_X .

In phase II, the muon stopping rate will be increased to $2 \cdot 10^9 \mu/s$ and accordingly a larger total number of observed muons is expected. Thus, even smaller branching fractions are in principle accessible. However, the increase in the muon rate entails an increased probability for fake tracks. This can be compensated by tighter cuts—at the expense of efficiency as usual. Given long tracks can be reconstructed online, especially limits for lower masses in the range of 20 MeV to 30 MeV will profit from the increased acceptance for recurling tracks due to the installation of additional recurl stations.

The physics' motivation for $\mu \rightarrow eX$ searches is the familon (see section 1.3). Apart from the channel $\mu \rightarrow eX$, the familon can also be investigated in other muon decay channels.

As an example, an additional photon can be emitted in $\mu \rightarrow e\gamma X$ [105]. This can however only be distinguished from SM background if the photon is detected as well. In *Mu3e*, the photon conversion probability is very low due to the overall small material amount. The potential *Mu3e-Gamma* experiment features a photon conversion layer and can notably increase the sensitivity to $\mu \rightarrow e\gamma X$. Preliminary studies on the reconstruction of vertices from electron-positron pairs and the *Mu3e-Gamma* upgrade have been conducted and are ongoing [203–205].

The familon can also decay into an electron-positron pair allowing to test electron-electron coupling in addition to electron-muon coupling. A prompt decay is investigated within $\mu \rightarrow eee$ searches [206,

207]. For familons with a certain lifetime, again a displaced vertex fit for electron-positron pairs can be employed.

Furthermore, two virtual familons can be studied in $\mu \rightarrow eeeee$ decays [206, 207]. This is particularly interesting as it would allow to constrain all couplings of the familon involving electrons. This channel however suffers from a low efficiency as the decay particles tend to have low momentum. Thus, it is easy have at least one particle below the detection threshold. The magnetic field could be lowered for such searches. In this way, the transverse radius increases and thus the reconstruction efficiency for low momentum particles improves.

SEARCHES FOR ELECTRON-POSITRON RESONANCES

In this chapter various strategies for searches of dark photons emitted in muon decays with the *Mu3e* experiment are discussed. Without modifications to the data acquisition system, promptly decaying dark photons can be investigated by a resonance search in the electron-positron system. The sensitivity of this channel is investigated. With a large number of $\mu \rightarrow eee\nu$ decays that will be observed, *Mu3e* has the opportunity to improve on current limits on the kinetic mixing parameter in the accessible mass range.

6.1 SEARCHES FOR DARK PHOTONS

The search for a resonance of electron-positron pairs is motivated by dark photons A' and similar new light vector bosons (see section 1.4) that are possibly emitted in muon decays of the type $\mu \rightarrow e\nu A'$ and decay into an electron-positron pair $A' \rightarrow ee$.

There are three scenarios in which dark photons can potentially be probed with the *Mu3e* experiment.


If the dark photon has a very short life time, it immediately decays after production $\mu \rightarrow e\nu(A' \rightarrow ee)$. The signature is similar to radiative muon decays with internal conversion $\mu \rightarrow eee\nu$, but there is a resonance in the invariant mass spectrum of the electron-positron system at the mass $m_{A'}$ of the dark photon. Such prompt decays are studied further in the following.

For dark photons with a certain life time, a search for an electron-positron pair from a displaced vertex can be conducted. Previous studies on the reconstruction of displaced vertices in *Mu3e* have shown that the spatial resolution of the vertex position is of the order of mm [203, 204]. Given that the A' decay has to happen inside the volume enclosed by the innermost vertex pixel layer in order to be reconstructed by the standard reconstruction¹, a large portion of the accessible decay lengths cannot be resolved and can thus also be covered by searches for prompt decays. Apart from the target re-


¹ Offline reconstruction of long 5- and 7-hit tracks without a hit in the innermost vertex layer was recently implemented [204], but as these tracks are long tracks they are not selected by the online reconstruction. Reconstruction algorithms for tracks that only have hits in the outer layers are currently being developed [205].

gion, displaced vertices can also be searched for between the vertex layers and the fibre detector as well as shortly outside the outer layers. Regions in direct vicinity to the pixel layers and the fibre detector should be left out to avoid background from Bhabha scattering. These searches require modifications to the online event selection and reconstruction of tracks and vertices which are in principle feasible. Decay lengths of several mm up to a few cm can thus be investigated and new regions in the dark photon parameter space come into reach [208]. With the potential *Mu3e-Gamma* upgrade [177, 178], decay lengths in the range of a few 10 cm can be studied. The dark photon can also be stable, long-lived or decay into invisible particles so that no electron-positron pair can be detected in the experiment. The momentum of the remaining positron depends in principle on the mass $m_{A'}$, but is broadly distributed as expected for a four-body decay (see figure 6.1). Similar distributions can also be created by a small misalignment of the detector (see figure 6.1c for an example). Thus, muon decays with invisible dark photons are indistinguishable from detector effects and such a search is not feasible with the *Mu3e* detector.


As the final state of muon decays with prompt dark photon decay $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ contains three electrons, this channel can be investigated with the same data that is collected for $\mu \rightarrow eee$ searches. Thus, the full precision of offline track and vertex reconstruction is available. The dark photon can be identified by a resonance at the mass $m_{A'}$ in the invariant mass spectrum of electron-positron pairs. The background sources are essentially the same as for $\mu \rightarrow eee$ searches. There is an irreducible background from radiative decays with internal conversion $\mu \rightarrow eee\nu\nu$. As it has the same final state as $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$, the prompt dark photon searches cannot be performed background-free. For the internal conversion decays, the distribution of the invariant mass $m_{e^-e^+}$ is maximum at zero. Further background stems from accidental combinations of two positron and one electron candidate from different processes. It is expected that the dominant contribution is produced by two Michel positrons from $\mu^+ \rightarrow e^+\bar{\nu}_\mu\nu_e$ of which one positron undergoes Bhabha scattering in the detector material and transfers enough momentum to the electron for it to be visible in the detector. These cases will be referred to as *Bhabha events* in the following. Other possible sources of accidental background are radiative muon decays with conversion of the decay photon and combinations of three uncorrelated tracks, potentially with positron tracks being misidentified as electrons. The occurrence of the former is found to be



(a) True momentum spectra.



(b) Reconstructed momentum spectra.



(c) Reconstructed momentum spectrum with and without a shift of the vertex layers by ± 0.2 mm in beam direction (courtesy of U. Hartenstein [209]).

Figure 6.1: Momentum spectra of the positron in invisible dark photon decays $\mu \rightarrow e\nu\nu A'$, and the effect of misalignment on the reconstructed momentum spectrum of SM muon decays.

negligibly small. Out of a sample of $4.3 \cdot 10^8$ generated events not a single event is left after reconstruction and cuts. The events are produced with a minimum photon energy of 20 MeV so that both conversion products are potentially visible. The branching fraction for this selection is $3.4 \cdot 10^{-3}$. Thus, the occurrence is expected to be smaller than 10^{-11} per muon decay or even less and therefore more than an order of magnitude smaller as for the already rare Bhabha events which is estimated in the following to be around 10^{-10} .

The background of three uncorrelated tracks is expected to be even further suppressed by timing and vertex cuts. Thus, apart from Bhabha events, no further accidental background processes are considered. In fact, the omitted background processes are expected to have a smooth m_{ee} distribution so that a small contribution would anyhow have hardly an impact on the sensitivity of a m_{ee} resonance search.

6.2 ESTIMATING THE SENSITIVITY

Similarly to the $\mu \rightarrow eX$ searches in chapter 5, a bump search on a background spectrum is performed for the dark photon searches. In this case, spectra of the reconstructed invariant mass of electron-positron pairs are studied. For a large number of muon decays, the number of background events in a certain m_{ee} interval is expected to be Poisson distributed. Assuming no event is observed during the run time of the experiment, the smallest branching fraction that $Mu3e$ is sensitive to becomes (see also equation 5.7)

$$\mathcal{B}_{\mu \rightarrow e\nu\nu(A' \rightarrow ee)}^{\text{limit}} \propto \sqrt{\left. \frac{d\mathcal{R}_{\text{BG}}(m_{ee})}{dm_{ee}} \right|_{m_{ee}=m_{A'}}} \times \frac{\sqrt{\Delta m_{ee} \epsilon_{\text{BG}}(m_{A'})}}{\epsilon_{A'}(m_{A'})} \frac{1}{\sqrt{N_{\mu}}}. \quad (6.1)$$

Herein, $\mathcal{R}_{\text{BG}}(m_{ee})$ denotes the background rate at a given invariant electron-positron mass m_{ee} . This includes internal conversion and Bhabha events as well as background from mis-reconstruction. Δm_{ee} is the resolution by which m_{ee} can be measured, $\epsilon_{\text{BG}}(m_{A'})$ and $\epsilon_{A'}(m_{A'})$ are the efficiencies for background and signal events for m_{ee} around $m_{A'}$, and N_{μ} the number of observed muon decays. The efficiencies are meant to include the detector acceptance as well as the reconstruction and selection efficiency and depend in general on the investigated mass.

As expected, the best sensitivity is achieved in m_{ee} regions with a

low background or a low background reconstruction efficiency. As the number of background events enters however only with the square root, small contributions as from Bhabha events do not have a noticeable impact on the sensitivity. The inverse of the signal efficiency enters directly so that signal efficiency should be given preference over background rejection in order to yield an optimum sensitivity. Finally, the expected minimum limit scales with the square root of the number of muons and can be improved by longer run times up to a point where systematic uncertainties dominate. Neither the background distribution, because of the accidental background contribution, nor the efficiencies can be determined analytically. Therefore, the sensitivity of the *Mu3e* experiment to $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays is estimated by the means of toy Monte Carlo studies using simulated m_{ee} spectra of signal and background.

6.2.1 Toy Monte Carlo Approach

Histograms of the reconstructed invariant mass m_{ee} of electron-positron candidates from three-track vertex fits are used as input to the toy Monte Carlo study (see section 6.3 in the following for more details). For the background, internal conversion and Bhabha events are combined. All histograms are smoothed by taking the median in a series over three, five, and three bins [200].

In figure 6.2, a few steps of the toy Monte Carlo study are shown for an exemplary mass $m_{A'}$ of 45 MeV. Figure 6.2a shows the histogram of signal events used to derive the signal p.d.f. $f_s(m_{ee})dm_{ee}$. Figure 6.2b shows the background p.d.f. $f_b(m_{ee})dm_{ee}$ as derived from the background histograms. f_b is composed of the p.d.f.'s for internal conversion events f_{IC} and for Bhabha events f_{Bhabha}

$$f_b = (1 - f) \cdot f_{IC} + f \cdot f_{Bhabha} . \quad (6.2)$$

The fraction f of Bhabha events is inferred from the occurrence of internal conversion and Bhabha events per muon decay and their respective efficiencies ϵ . The efficiency is here taken as the ratio of the sum of weights² w of reconstructed and of generated events. In

² The events are weighted with the differential decay rate of the process. The weight corresponds thus to the number of events.

addition, the efficiency t due to the timing cuts needs to be considered.

$$\epsilon_{\text{IC/Bhabha}} = t_{\text{IC/Bhabha}} \frac{w_{\text{rec}}^{\text{IC/Bhabha}}}{w_{\text{gen}}^{\text{IC/Bhabha}}} \quad (6.3)$$

The branching fraction b_{IC} of radiative muon decays with internal conversion is $3.4 \cdot 10^{-5}$ [41]. The fraction of Bhabha events per muon decay b_{Bhabha} in the *Mu3e* experiment, requiring electron and positron to be potentially visible with a minimum energy of 10 MeV, is found to be $7.8 \cdot 10^{-5}$ [136]. Thus, the Bhabha contribution to the background p.d.f. becomes

$$f_{\text{gen}} = \frac{b_{\text{IC}}}{b_{\text{Bhabha}}} \quad (6.4)$$

$$f_{\text{rec}} = f_{\text{gen}} \frac{\epsilon_{\text{IC}}}{\epsilon_{\text{Bhabha}}} \quad (6.5)$$

$$f = \frac{1}{f_{\text{rec}} + 1}, \quad (6.6)$$

using the ratio of internal conversion events to Bhabha events at generator (f_{gen}) and reconstruction level (f_{rec}).

From the background p.d.f., N_{toyMC} events are generated in the toy Monte Carlo study. A signal-plus-background p.d.f.

$$f_{s+b} = (1 - f_{\text{sig}}) \cdot f_b + f_{\text{sig}} \cdot f_s \quad (6.7)$$

is fitted to the generated events with the signal fraction f_{sig} as free fit parameter using the maximum likelihood method on binned data. An exemplary fit is shown in figure 6.2c. This is repeated for a large number of samples. The width σ_{toyMC} of the distribution of the fitted signal fractions is taken to derive the sensitivity to the branching fraction for $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$. σ_{toyMC} is determined by fitting a Gaussian distribution with mean zero to the distribution of fitted f_{sig} (see figure 6.2d). The respective quantile $n_\sigma \sigma_{\text{toyMC}}$ gives the confidence level (CL, see table 5.1). Thus, in N_{toyMC} reconstructed background events, up to a total of $n_\sigma \sigma_{\text{toyMC}} N_{\text{toyMC}}$ reconstructed dark photon events can be excluded at the corresponding CL.

N_{toyMC} reconstructed background events correspond to

$$N'_\mu = N_{\text{toyMC}} \frac{1}{\epsilon_{\text{BG}}} \quad (6.8)$$

muon decays. The effective background efficiency ϵ_{BG} also incorporates the occurrence of the background processes.

$$\epsilon_{\text{BG}} = b_{\text{IC}}\epsilon_{\text{IC}} + b_{\text{Bhabha}}\epsilon_{\text{Bhabha}} \quad (6.9)$$

Translating N'_μ to the desired number of muon decays N_μ , involves scaling the width σ_{toyMC} . As shown before (see section C.2 in the appendix), σ_{toyMC} scales with the inverse of the square root of the number of events.

$$\sigma_{\text{toyMC}} \rightarrow \sigma_{\text{toyMC}} \sqrt{\frac{N'_\mu}{N_\mu}} = \sigma_{\text{toyMC}} \sqrt{\frac{N_{\text{toyMC}}}{N_\mu \epsilon_{\text{BG}}}} \quad (6.10)$$

Thus, up to a total of

$$n_\sigma \sigma_{\text{toyMC}} \sqrt{\frac{N_{\text{toyMC}}}{N_\mu \epsilon_{\text{BG}}}} N_\mu \epsilon_{\text{BG}} = n_\sigma \sigma_{\text{toyMC}} \sqrt{N_{\text{toyMC}} N_\mu \epsilon_{\text{BG}}} \quad (6.11)$$

reconstructed dark photon events can be excluded at the specified confidence level. With the dark photon reconstruction efficiency (again taking the timing efficiency into account)

$$\epsilon_{A'} = t_{A'} \frac{w_{\text{rec}}^{A'}}{w_{\text{gen}}^{A'}}, \quad (6.12)$$

this corresponds to

$$n_\sigma \sigma_{\text{toyMC}} \sqrt{N_{\text{toyMC}} N_\mu \epsilon_{\text{BG}}} \frac{1}{\epsilon_{A'}} \quad (6.13)$$

actual $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ events.

The expected limit on the branching fraction that $Mu3e$ is sensitive to thus becomes

$$\mathcal{B} \leq n_\sigma \sigma_{\text{toyMC}} \sqrt{N_{\text{toyMC}} N_\mu \epsilon_{\text{BG}}} \frac{1}{\epsilon_{A'}} \frac{1}{N_\mu} \quad (6.14)$$

$$\leq n_\sigma \sigma_{\text{toyMC}} \sqrt{\frac{N_{\text{toyMC}} \epsilon_{\text{BG}}}{N_\mu}} \frac{1}{\epsilon_{A'}} \quad (6.15)$$

for N_μ observed muon decays. Equation 6.15 is very similar to equation 6.1. The signal and background efficiencies $\epsilon_{\text{BG}}(m_{A'})$ and $\epsilon_{A'}(m_{A'})$ at the dark photon mass $m_{A'}$ are replaced by the overall efficiencies as the toy Monte Carlo study is here performed on the full-range invariant mass spectrum. Signal and background distri-

butions as well as detector effects are encoded in the width σ_{toyMC} .

As in the case of the toy Monte Carlo studies of $\mu \rightarrow eX$ searches, it has to be emphasized that the toy Monte Carlo approach as applied here incorporates some idealizations and simplifications.

The matrix element of the dominating radiative muon decay with internal conversion is implemented at leading-order in the simulation. Recent next-to-leading calculations suggest corrections in the order of 1 % [194, 195].

The prediction of accidental background is even less precise and is limited to the dominating contribution of Bhabha events. The rate and spatial distribution of Bhabha scattering events depends for instance on the material distribution in the detector as well as on the actual muon beam profile. The detector simulation incorporates the current best knowledge about the geometry but of course employs some simplifications, for instance on cabling. Furthermore, the material distribution in the real experiment will never be known with the same accuracy as in the simulation.

Moreover, the position of all active detector elements is known to full precision in the simulation. In the later experiment, there will be uncertainties due to misalignment. Furthermore, there are possibilities to miss a hit which are not considered in the simulation for instance because of small spacings between the individual pixel sensors or the limitations of the data acquisition. In addition, the simulated pixel sensors are idealized with an efficiency of 100 % and no noise hits or hit clusters. Although current prototypes are very close to these assumptions with efficiencies $\geq 99\%$ at low noise rates [148], the reconstruction efficiencies are thus slightly over- and the fake rate slightly underestimated.

In the current pre-construction stage of the experiment, the outlined toy Monte Carlo study is nevertheless an appropriate tool to estimate the sensitivities and to study the suitability of the reconstruction for dark photon searches. With a finalized detector geometry, the sensitivity studies have to be repeated for a more realistic estimate. A few considerations on improvements and a possible concept for the later analysis are presented in the outlook (see section 6.8).

In the following studies, 10^4 samples with $N_{\text{toyMC}} = 10^9$ events each are generated³. 300 days of data taking at $10^8 \mu/\text{s}$ are assumed which corresponds to a total of $N_\mu = 2.592 \cdot 10^{15}$ muon decays.

6.3 SIMULATION OF MUON DECAYS WITH PROMPT DARK PHOTON DECAY

6.3.1 *Event Generator*


The dark photon model is implemented using the FEYNRULES 2.3.23 framework [210]. FEYNRULES is a software tool to derive Feynman rules from a Lagrangian and is commonly used in particle physics phenomenology. For dark photons, the existing SM Lagrangian is extended by the dark photon Lagrangian as given by equations 1.24 and 1.25. The output in UNIVERSAL FEYNRULES OUTPUT [211] format is fed into the MADGRAPH5_aMC@NLO 2.4.3 event generator [212]. The $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e (A' \rightarrow e^- e^+)$ events are produced for fully polarized muons. A dedicated physics process is implemented in the *Mu3e* software framework that interprets the LES HOUCHES EVENT FILES [213, 214] generated with MADGRAPH. As the dark photon is emitted in muon decays and is required to decay into an electron-positron pair in the investigated channel, it can have a mass of

$$2m_e = 1.02 \text{ MeV} \leq m_{A'} \leq m_\mu - m_e = 105.15 \text{ MeV}, \quad (6.16)$$


neglecting neutrino masses. The efficiency for high $m_{A'}$ vanishes however as the additional positron from the muon decay has a low momentum and is thus outside of the detector acceptance. This could potentially be overcome by dedicated two-prong searches. About 700000 signal events are generated and simulated per mass. The masses range from 5 MeV to 105 MeV in steps of 5 MeV. In addition, events with a small mass of 2 MeV and of 7 MeV are generated motivated by the maxima in the background distributions.

As background sources, radiative muon decays with internal conversion $\mu \rightarrow e e \nu \nu$ and Bhabha events are taken into account. Both processes are generated as *special decays* with at most one event of the specified process per reconstruction frame of 50 ns (see section 2.3 for details). For Bhabha events, a normal Michel decay with


³ As shown in table 6.3, the number of background events varies around $3 \cdot 10^8$ for $2.592 \cdot 10^{15}$ observed muon decays. N_{toyMC} is nevertheless chosen to be 10^9 in order to be able to scale the results to any desired number of muon decays.




(a) Histogram of the reconstructed m_{ee} from $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ signal events with $m_{A'} = 45\text{MeV}$. About 700 000 decays in the target region are generated.



(b) A sample of 10^9 events generated from f_b . The contributions of internal conversion (green line) and Bhabha events (red line) are shown.



(c) A single sample of events derived from f_b . The solid blue line is a fit of f_{s+b} to the events. The fitted amount of signal (solid red line) and background (dashed blue line) are also shown. The solid and dashed blue line are indistinguishable on the logarithmic scale.



(d) Histogram of the fitted signal fraction for a study of 10^4 samples. A Gaussian distribution is fitted to the results (solid blue line).

Figure 6.2: Exemplary toy Monte Carlo study for reconstructed vertices using all tracks. The invariant mass m_{ee} of both electron-positron pairs is considered. An exemplary mass of $m_{A'} = 45\text{MeV}$ is used for the $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ sample.

immediate Bhabha scattering at the decay vertex is generated simultaneously to normal muon decays at a stopping rate of $10^8 \mu/s$. The generated electron and positron are required to be potentially visible, i. e. they both have a minimum energy of 10 MeV. For internal conversion events, no restrictions on the phase space are applied. The reconstruction efficiency for the background events is very low with about $4 \cdot 10^{-3}$ for internal conversion events and about $1 \cdot 10^{-4}$ for Bhabha events. A large number of events needs to be generated in order to be able to build a suitable background p.d.f. from the remaining reconstructed events. About $7 \cdot 10^8$ weighted internal conversion and $5 \cdot 10^7$ weighted Bhabha events are generated. Both samples are generated with weights⁴ according to their respective differential decay rate in order to perform the simulation within a reasonable time. The weight thus corresponds to the number of events.

6.3.2 Reconstruction

The resonance search in electron-positron pairs is performed on three particle vertex candidates—the same data that is used in the search for $\mu \rightarrow eee$. Thus, no modification to the data acquisition is needed. The full precision of the offline reconstruction can be exploited, in particular the reconstruction of recurling tracks.

The resonance search is performed on spectra of the invariant mass m_{ee} of an electron and a positron track candidate. In the following analysis, two cases will be differentiated. Either, no selection is applied based on whether the tracks are reconstructed as short or long tracks. This will be called *all tracks*. Or, it is required that the two tracks for the electron-positron pair candidate are reconstructed as recurlers. This is called *two long tracks*. The third track does not contribute in the resonance search. Therefore, no selection on the reconstruction of this track is made as it would unnecessarily

⁴ The events are generated with an even phase space distribution and subsequently weighted with the corresponding differential decay rate. Alternatively, the accept-reject method [168] can be used to produce an unweighted sample that is distributed according to the decay rate. This requires longer computation times because a larger amount of random numbers needs to be generated.

In the simulation, the information about the weight is lost, once the particle crosses a frame boundary. In these cases, the weight is set to the mean weight of the whole sample.

For internal conversion events, large weights are observed which disturb the smoothness of the later p.d.f. Therefore, weights larger than $1 \cdot 10^7$ are omitted when building the p.d.f. The mean weight for internal conversion events is about 5000. Only a fraction of about $5 \cdot 10^{-5}$ of the events are rejected with this selection. The events are not omitted when the weights are summed up for the calculation of the reconstruction efficiencies.

decrease the efficiency.

With two positrons and one electron, there are two possibilities to define the electron-positron pair. In the following, the spectra of the invariant electron-positron pair mass m_{ee} either contain the combination that yields the higher m_{ee} (*high m_{ee}*) or lower m_{ee} (*low m_{ee}*), or both combinations are considered (*both m_{ee}*).

Prior to the offline reconstruction, the events have to pass the online selection which includes online track and vertex finding. Currently, estimates of the online selection efficiency exist only for signal events. As this efficiency is likely to depend for instance on the invariant mass of the three electron system, these studies have to be repeated also for other physics processes which is beyond the scope of this work. Thus, the inefficiencies from the online selection cannot be considered in the following.

6.3.3 Selection and Cuts

As before, the branching fraction limits that can be expected for *Mu3e* will be quoted for a certain number of muon stops. Thus, only events with a true decay vertex in the target region are considered excluding in flight decays in the beam pipe. For Bhabha events, the true decay vertex is ill-defined. As Bhabha events are implemented such that the Bhabha scattering happens on the target anyhow, this selection can be omitted here.

As in the case of $\mu \rightarrow eee$ searches, a few selections are applied during the vertex fit (see chapter 4). Short tracks are required to have a maximum χ^2_{track} value of 32, and long tracks of 48. Events are rejected when two tracks have a momentum difference of less than 1 MeV and the opening angles indicate at the same time that these are two tracks of the same particle. Thus, for two positron track candidates the opening angle in this case has to be at least 0.14 rad, and for an electron and positron track candidate it has to be less than 3.00 rad.

A further selection is based on the χ^2 of the vertex fit. As all events are generated as special decays, only one event per reconstruction frame of 50 ns is expected. Some events are however reconstructed multiple times, for instance when a particle has mostly transverse momentum and performs many closely spaced loops. Then, many vertices are reconstructed by iterating over all the turns of the recurling particle. Multiple reconstructions are removed by keeping only the event with the smallest χ^2 in the event frame. Such a selection can also be applied in the later analysis as the chance to have two

events with similar kinematics in a certain time slice and volume is vanishingly small.

Cuts are applied on reconstructed variables from the vertex fit. As the two neutrinos in $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ carry an a priori unknown amount of the total energy, kinematic cuts like on the total momentum and invariant mass of the three electrons as applied in $\mu \rightarrow eee$ searches are not applicable. The remaining cut variables are the χ^2 value of the vertex fit and the distance of the reconstructed vertex to the target d .

Furthermore, cuts on the relative timing between the three tracks can be placed. Here, the same timing cuts as in $\mu \rightarrow eee$ searches are assumed. Details on the timing suppression can be found in [136, 154]. $\mu \rightarrow eee\nu\nu$ and $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ events are not strongly affected by the timing cuts as the three electrons emerge from a common vertex. An efficiency of 98% is expected, the same as for $\mu \rightarrow eee$ events. In Bhabha events, only the electron-positron pair from Bhabha scattering is correlated and thus a suppression by a factor of 100 can be achieved with timing cuts.

χ^2 OF THE VERTEX FIT The goodness of the vertex fit is estimated by the resulting χ^2 value. It is expected, that random combinations of three tracks yield larger values of χ^2 than decays that really have three electrons in the final state.

In figure 6.3, the distribution of χ^2 is shown for reconstructed $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ events with $m_{A'} = 45$ MeV as well as for internal conversion background and background from Bhabha events. Vertices built from all kinds of reconstructed tracks are considered.

The events are distributed as expected for a χ^2 distribution in the case of signal events and background from internal conversion. Bhabha events have a flat χ^2 distribution as there is no common vertex to be reconstructed.

In figure 6.4, signal efficiency and background rejection are shown in dependence on the χ^2 cut. The signal efficiency rises steeply and becomes then almost constant with increasing cuts on χ^2 . Likewise, the background rejection for internal conversion decays has a steep decline at tight cuts and nearly vanishes for looser cuts. The rejection of Bhabha events is strong and decreases linearly with the χ^2 cut as expected from the flat χ^2 distribution. Suppression factors of more than two orders of magnitude can be achieved while keeping a high signal efficiency.

As the χ^2 cut is not very effective in rejecting the dominant internal

conversion background, preference is given to signal efficiency and a loose cut of $\chi^2 \leq 30$ is chosen.

DISTANCE TO THE TARGET A $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ signal event is expected to originate from a stopped muon on the target, and thus the distance between the reconstructed vertex and the surface of the target is a suitable variable to suppress accidental background.

The distance to the target d is defined as the minimum distance between the reconstructed vertex position and the target surface. The sign of d indicating whether the reconstructed vertex is located inside or outside of the target is of no greater relevance here. Therefore, the cut is placed on $|d|$.


In figure 6.5, the distribution in $|d|$ for $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ signal events and background from internal conversion decays and Bhabha events are shown. Again, all kinds of reconstructed tracks are taken into account in the vertex fit.

For all three processes, the distribution is highest at small $|d|$ and decreases towards larger target distances. For signal and internal conversion decays, this is expected from the common vertex. And also in Bhabha events, at least the electron-positron pair from Bhabha scattering originates from one point in space. Therefore, the $|d|$ distribution of Bhabha events also decreases but with a smaller slope. The tail towards larger $|d|$ in signal and internal conversion events stems from decays in flight. In the case of Bhabha events, the third uncorrelated particle can have any distance to the Bhabha scattering pair so that the tail towards larger $|d|$ becomes significantly longer.


Figure 6.6 shows the signal efficiency and background rejection for cuts on the target distance. The signal efficiency increases steeply at small $|d|$ cuts and reaches a plateau within a few mm. In a similar way, the background rejection of internal conversion background decreases rapidly. The rejection of Bhabha events declines with a smaller slope than in the case of internal conversion decays.

As in the case of the χ^2 cut, no strong rejection for the internal conversion background can be achieved as internal conversion decays and signal decays behave very similarly. In order to maintain a good signal efficiency, a cut on the target distance is chosen at $|d| \leq 3$ mm for all dark photon masses.


COMBINING THE CUTS In figure 6.7, the distribution of the reconstructed invariant mass of electron-positron pair candidates is shown for signal decays and background before and after applying cuts. In addition, the reconstruction efficiencies are summarized in table 6.1.



(a) Distribution in χ^2 for prompt dark photon decays with $m_{A'} = 45 \text{ MeV}$.




(b) Distribution in χ^2 for background from internal conversion decays.

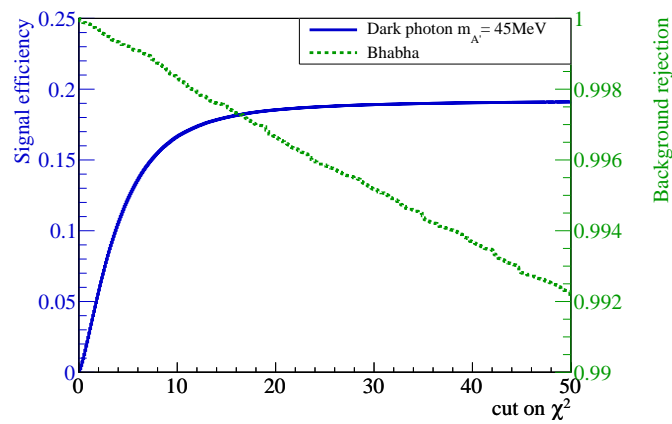


(c) Distribution in χ^2 for background from Bhabha events. Note that the range of the x-axis is twice as large as in the other two figures.

Figure 6.3: Distribution of the χ^2 value of the vertex fit for muon decays with prompt dark photon decay and background from internal conversion and Bhabha events. A dark photon mass of 45 MeV is chosen.




(a) Signal events and background from internal conversion.




(b) Signal events and background from Bhabha events.


Figure 6.4: Signal efficiency of muon decays with prompt dark photon decay and rejection of background from internal conversion and Bhabha events for cuts on the χ^2 of the vertex fit. A dark photon mass of 45 MeV is chosen. Note that the axis range of the background rejection is chosen differently for the two types of background.



(a) Distribution in target distance $|d|$ for prompt dark photon decays with $m_{A'} = 45 \text{ MeV}$.



(b) Distribution in target distance $|d|$ for background from internal conversion decays.



(c) Distribution in target distance $|d|$ for background from Bhabha events. Note that the range of the x-axis is changed compared to the other two figures.

Figure 6.5: Distribution of the distance to the target $|d|$ of the fitted vertex for muon decays with prompt dark photon decay and background from internal conversion and Bhabha events. A dark photon mass of 45 MeV is chosen.




Figure 6.6: Signal efficiency of muon decays with prompt dark photon decay and rejection of background from internal conversion and Bhabha events for cuts on the distance to the target $|d|$ of the fitted vertex. A dark photon mass of 45 MeV is chosen.

The m_{ee} spectra of the $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ signal decays are basically unaffected by the cuts. The resulting signal efficiencies relative to the number of generated events vary between a few percent up to about 20 % at $m_{A'}$ around 50 MeV including an expected inefficiency from timing cuts of 98 %. The cuts reduce the signal efficiencies by about 6 % for most of the considered dark photon masses.

Background from internal conversion events is likewise nearly unaffected by the cuts. The reconstruction efficiency is 0.46 %.

For Bhabha events, a remarkable reduction is achieved as both cuts are effective measures against accidental background. The reconstruction efficiency for Bhabha events is reduced by more than four orders of magnitude. This is also due to the timing cuts which are expected to reduce Bhabha events by a factor of 100. The resulting reconstruction efficiency is $2.6 \cdot 10^{-6}$.

The contribution of Bhabha events to the background is about three orders of magnitude smaller than the contribution of internal conversion events. The contribution f of Bhabha events as used throughout section 6.2.1 varies around $1.2 \cdot 10^{-3}$ depending on the choice of the reconstructed electron-positron pair (see table 6.2).

The efficiencies and the expected number of events for $\mu \rightarrow ee\nu\nu$ and Bhabha background are given in table 6.3. The vertex reconstruction efficiency especially for Bhabha events is very small so that only a few thousand of the simulated events are left over after reconstruction. This determines the choice of the mini-

mum bin width of the m_{ee} spectra for the sensitivity studies. The bin width is set to 1 MeV. For finer bin sizes, no suitable background p.d.f. can be built. As tracks and vertices will be reconstructed offline in the later experiment, considerably smaller bin widths can be chosen in the real data analysis which in turn is expected to improve the sensitivity. Spectra of accidental background can be easily obtained from data by inverting the timing cuts for the coincidence of the three particles.

6.4 EXPECTED SENSITIVITY TO THE BRANCHING FRACTION

In the following, invariant masses from an electron-positron pair candidate derived from long recurling tracks and from all kinds of reconstructed tracks are discussed. Also, for the m_{ee} spectra either the pair with the higher or lower m_{ee} as well as both pairs are considered.

In figure 6.8, m_{ee} spectra for signal decays with dark photons of different masses are shown. The same spectra for background events are shown in figure 6.9. There are two contributions to the dark photon spectra. The electron-positron pair from the dark photon decay exhibits as a sharp peak. The additional distribution stems from the second possible electron-positron combination, i. e. the electron from the dark photon decay is combined with the positron from the muon decay. The latter have a large range of possible m_{ee} and thus yield a broad distribution. Depending on the dark photon mass, it is more likely to cover the dark photon pair by the electron-positron combination with the higher or lower m_{ee} .



Both background spectra have a maximum at small m_{ee} . In the case of internal conversion decays, an m_{ee} of zero is most probable because of the conversion of the virtual photon. Due to the finite m_{ee} resolution, the maximum in the reconstructed m_{ee} distribution is shifted to a few MeV. Also in Bhabha scattering an invariant electron-positron mass of zero is most probable as here small opening angles are preferred. The maximum of the reconstructed m_{ee} is located at slightly larger values than for internal conversion events as a certain momentum transfer from the positron to the electron is required to actually detect the event. Both spectra have an additional broad distribution resulting from the second electron-positron combination. Towards large reconstructed m_{ee} , the background spectra fade. Only a few weighted events are left per bin. The background prediction is no longer considered reliable, and therefore only dark photon masses up to 80 MeV are investigated in the following. This is a limi-

$m_{A'}$ [MeV]	$\epsilon_{\text{No cuts}}$ [%]	ϵ_{Cuts} [%]	$\frac{\epsilon_{\text{No cuts}} - \epsilon_{\text{Cuts}}}{\epsilon_{\text{No cuts}}}$ [%]
2	0.369	0.331	10.3
5	0.559	0.517	7.5
7	0.728	0.677	7.0
10	0.957	0.905	5.4
15	1.53	1.44	5.9
20	2.45	2.32	5.3
25	4.51	4.27	5.3
30	8.37	7.93	5.3
35	13.1	12.4	5.3
40	17.5	16.5	5.7
45	20.4	19.2	5.9
50	21.2	20.0	5.7
55	20.3	19.2	5.4
60	18.0	16.9	6.1
65	14.8	13.9	6.1
70	11.0	10.4	5.5
75	6.98	6.54	6.3
80	3.55	3.31	6.8
Internal conversion	0.510	0.459	10.0
Bhabha events	15	0.000 26	~ 100


Table 6.1: Efficiency of reconstructed $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ events and background from $\mu \rightarrow eee\nu\nu$ and Bhabha events before and after applying cuts. The relative statistical uncertainties are on the 10^{-4} level except for Bhabha events ($\mathcal{O}(10^{-2})$). Dark photon masses larger than 80 MeV are not listed as they will not be considered in the later analysis. Inefficiencies due to timing cuts are also included in ϵ_{Cuts} .

m_{ee}	f
all tracks	$1.3 \cdot 10^{-3}$
at least 2 long tracks	$1.2 \cdot 10^{-3}$
2 long tracks for the higher m_{ee} pair	$1.2 \cdot 10^{-3}$
2 long tracks for the lower m_{ee} pair	$1.2 \cdot 10^{-3}$

Table 6.2: Contribution f of Bhabha events to the background p.d.f.

(a) $m_{A'} = 45 \text{ MeV}$.

(b) Internal conversion events.



(c) Bhabha events.

Figure 6.7: Reconstructed invariant mass for muon decays with prompt dark photon decay and background from internal conversion and Bhabha events before and after applying cuts. The invariant mass spectra of the electron-positron pair are shown, taking both combinations into account. All reconstructed tracks are considered. A dark photon mass of 45 MeV is chosen.

m_{ee}	ϵ	N
Internal conversion events		
all tracks	$4.6 \cdot 10^{-3}$	$4.0 \cdot 10^8$
at least 2 long tracks	$3.4 \cdot 10^{-3}$	$3.0 \cdot 10^8$
2 long tracks for the higher m_{ee} pair	$3.4 \cdot 10^{-3}$	$3.0 \cdot 10^8$
2 long tracks for the lower m_{ee} pair	$3.4 \cdot 10^{-3}$	$3.0 \cdot 10^8$
Bhabha events		
all tracks	$2.6 \cdot 10^{-6}$	$5.2 \cdot 10^5$
at least 2 long tracks	$1.8 \cdot 10^{-6}$	$3.6 \cdot 10^5$
2 long tracks for the higher m_{ee} pair	$1.7 \cdot 10^{-6}$	$3.4 \cdot 10^5$
2 long tracks for the lower m_{ee} pair	$1.8 \cdot 10^{-6}$	$3.7 \cdot 10^5$

Table 6.3: Efficiency and expected number of events from internal conversion and Bhabha events assuming 300 days of data taking at $10^8 \mu/s$.


tation of the sample size used in this study. The later experiment is of course also sensitive to larger $m_{A'}$. However, as shown in the following, the sensitivity of *Mu3e* at large $m_{A'}$ is no longer competitive with already existing constraints.

The procedure described in the previous section 6.2 is used to derive expected upper limits on the branching fraction of prompt dark photon decays $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ with either or both electron-positron combinations. Dark photon masses between 2 MeV and 80 MeV are investigated. The expected upper limits⁵ at 90 % CL are shown in figure 6.10. The sensitivity ranges from branching fractions of a few 10^{-9} at small dark photon masses down to a few 10^{-12} at high $m_{A'}$. The reconstruction efficiency for prompt dark photon decays and the width σ_{toyMC} of the distribution of fitted signal fractions in the toy Monte Carlo study are displayed in figure 6.11.


BOTH ELECTRON-POSITRON COMBINATIONS At first, the m_{ee} spectra in which both electron-positron combinations enter are discussed.

The shape of the sensitivity curve is determined on the one hand by the width σ_{toyMC} which depends on the square root of the number of background events in the range of the peak of the potential dark


⁵ Here and in the following plots, the data points are simply connected via straight lines.




(a) Both combinations of electron and positron candidates using all tracks.




(b) Both combinations of electron and positron candidates using two long tracks.




(c) Electron-positron pair with higher m_{ee} using all tracks.



(d) Electron-positron pair with higher m_{ee} using two long tracks.




(e) Electron-positron pair with lower m_{ee} using all tracks.




(f) Electron-positron pair with lower m_{ee} using two long tracks.


Figure 6.8: Spectra of the reconstructed invariant mass m_{ee} for prompt dark photon decays. The spectra are drawn normalized.




(a) Both combinations of electron and positron candidates using all tracks.




(b) Both combinations of electron and positron candidates using two long tracks.




(c) Electron-positron pair with higher m_{ee} using all tracks.



(d) Electron-positron pair with higher m_{ee} using two long tracks.




(e) Electron-positron pair with lower m_{ee} using all tracks.




(f) Electron-positron pair with lower m_{ee} using two long tracks.

Figure 6.9: Spectra of the reconstructed invariant mass m_{ee} for internal conversion and Bhabha events. 300 days of data taking at $10^8 \mu/s$ are assumed.




(a) Comparison of all combinations of electron and positron candidates.




(b) Electron-positron pair with the best expected sensitivity.

Figure 6.10: Limits on the $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ branching fraction at 90% CL that *Mu3e* is sensitive to in 300 days of data taking at $10^8 \mu/s$. All possible combinations of electron and positron candidates are compared to choosing the optimum combination.



(a) Width σ_{toyMC} of the distribution of the signal fraction in toy Monte Carlo studies.



(b) Fraction of reconstructed events per $\mu \rightarrow e \nu \nu (A' \rightarrow ee)$ decay.

Figure 6.11: Width σ_{toyMC} of the distribution of the signal fraction in toy Monte Carlo studies and fraction of reconstructed events per $\mu \rightarrow e \nu \nu (A' \rightarrow ee)$ decay.

photon events and thus on the background distribution and m_{ee} resolution. Indeed, the distribution of σ_{toyMC} has a similar shape as the background m_{ee} spectrum. Therefore, also the decrease in sensitivity at around 5 MeV is not visible. On the other hand, the reconstruction efficiency for $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ events influences the sensitivity. The efficiency increases up to a mass of about 50 MeV and then decreases again. Thus, the flat region in σ_{toyMC} at medium $m_{A'}$ is compensated by the increase in dark photon efficiency, and the steep decline of σ_{toyMC} at high $m_{A'}$ due to the vanishing background is attenuated by the decrease of the efficiency.

At $m_{A'} = 5$ MeV, the sensitivity is slightly weaker than what would be expected from a smooth curve. The effect is easier identified in the sensitivity estimates for the kinetic mixing parameter ε (see section 6.5 below). This is a result of the maximum of the spectrum of internal conversion events at around 5 MeV. An effect by the maximum of Bhabha events at about 7 MeV is not observed as the maximum is broader than the dark photon resonance, and because the distribution of the second electron-positron pair candidate in Bhabha events does not match the m_{ee} distribution of dark photons. Moreover, Bhabha events are strongly suppressed compared to internal conversion events and thus barely contribute.

HIGH INVARIANT MASS COMBINATION The expected sensitivity to the $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ branching fraction using only the electron-positron pair with the higher m_{ee} has a similar $m_{A'}$ dependence as the previously discussed combination, except for small masses $m_{A'}$. Here, the distribution of background events falls off as the events are more likely to be covered by the small m_{ee} combination which results in a better sensitivity at small $m_{A'}$.

LOW INVARIANT MASS COMBINATION As expected from the distribution of background, the expected minimum branching fraction the experiment is sensitive to falls off steeply towards higher $m_{A'}$ in the case of the low m_{ee} combination. From about 40 MeV onwards however, the curve flattens out. This results from the m_{ee} distribution of dark photon events. The higher $m_{A'}$, the more likely the actual electron-positron pair from the dark photon results in the higher m_{ee} . The distribution of the lower m_{ee} thus contains more and more of the second electron-positron combination. It broadens with increasing $m_{A'}$ and the resonance peak is less and less pronounced. At the highest studied $m_{A'}$, the sensitivity improves again due to the vanishing background.

As before, a slightly weaker sensitivity is observed for $m_{A'} = 5 \text{ MeV}$ due to the maximum of the internal conversion distribution.

COMPARISON OF ELECTRON-POSITRON COMBINATIONS Figure 6.10a puts the expected sensitivity to prompt dark photon decays into comparison for the high m_{ee} , low m_{ee} , and both electron-positron combinations. Simply adding both electron-positron pair combinations to the spectra does not have an advantage over the other two combinations for any $m_{A'}$. This is expected as not only the electron-positron combination belonging to the dark photon decay but also the second possible combination is always taken into account. Thus, the m_{ee} distributions exhibit a broad base in addition to the resonance. Furthermore, the overall background is twice as large and has in addition a broader distribution than for the other two e^+e^- -combinations.

At most low masses, the low m_{ee} combination yields the best sensitivity. As expected, the dark photon resonance is best pronounced in the low m_{ee} combination. In addition, the distribution of the second m_{ee} combination is narrower than for the high m_{ee} combination. In a similar way, the sensitivity is improved when using the high m_{ee} combination for larger $m_{A'}$. The transition takes place at around 45 MeV. At 2 MeV, the high m_{ee} combination yields a better sensitivity. This results purely from the background distribution that decreases for small m_{ee} in this combination. The dark photon resonance is however more clearly visible in the low m_{ee} distribution. Therefore, the low m_{ee} combination is chosen for $m_{A'} = 2 \text{ MeV}$ anyway.

In figure 6.10b, the smallest upper limits on the branching fraction of prompt dark photon decays at 90 % CL are shown that are expected for phase I of the *Mu3e* experiment given the most suitable m_{ee} combination is chosen. The experiment is sensitive to branching fractions larger or equal to $5 \cdot 10^{-9}$ at small $m_{A'}$ and down to $3 \cdot 10^{-12}$ at high $m_{A'}$.

COMPARING ALL TRACK VERTICES AND VERTICES WITH TWO LONG TRACKS The choice of the reconstruction of tracks affects the m_{ee} resolution and the signal reconstruction efficiency.


The invariant mass obtained from vertex fits with two long tracks for the electron-positron pair candidate is expected to yield an improved resolution compared to vertex fits with short tracks. Figure 6.12 shows the m_{ee} resolution as root mean square of the difference of reconstructed and generated m_{ee} when putting no requirements on the type of track reconstruction and when requesting at least one,

two, or three recurling tracks for the vertex fit. The m_{ee} resolution varies between 0.5 MeV and 2 MeV. Fits with three recurling tracks overall show better m_{ee} resolutions, in particular towards large invariant masses. At intermediate masses, the benefit of long tracks is however not very strong. The difference between the other three combinations is not that pronounced. The resolutions differ up to a factor of roughly two in the relevant mass range whereas the momentum resolution of the single tracks improves up to about a factor of ten when a long track is reconstructed instead of a short one (see figure 2.18). This results from the fact, that the mass resolution is not only dependent on the momentum resolution of the track fit but also on the angular resolution which does not profit from reconstructing long tracks. In addition, also many of the events in the all track ensemble are at least partially built from reconstructed long tracks.


The mass resolution is beneficial in reducing background. However, requiring long tracks comes at the expense of efficiency (see figure 6.11). The reconstruction efficiency for prompt dark photon decays is highly depended on the mass $m_{A'}$. Most events are reconstructed for medium $m_{A'}$ where all of the three charged particles have a good chance to have transverse momenta above 10 MeV and thus being reconstructable. Once, the momenta are more asymmetrically distributed—thus at both small and high $m_{A'}$ —the efficiency decreases. The best reconstruction efficiency is achieved for $m_{A'}$ around 50 MeV with 20 % for all tracks and 14 % for two long tracks relative to the number of generated signal events. The efficiency for long tracks is lower because the particles can hit the non-instrumented gap between central and recurl station or can overshoot the recurl stations and thus do not cause additional hits in the outer pixel layers.

The expected sensitivity limits on the branching fraction obtained with all and two long tracks do not significantly differ in shape but have different amplitudes (see figure 6.10). Using all tracks yields a higher reconstruction efficiency and thus improves upon requiring two long tracks over the whole $m_{A'}$ mass range despite the better mass resolution.

ESTIMATES FOR MU3E IN PHASE II In the following, also estimates for phase II of the *Mu3e* experiment are given. These are meant to give a rough estimate on how much the sensitivity to dark photons will improve with the upgraded experiment. A few simplifications are made. For instance, the same detector geometry and beam profile is used as in phase I as especially the latter is not yet known for phase II.




(a) m_{ee} resolution requiring different minimal numbers of long tracks in the vertex fit averaged over the internal conversion sample.




(b) m_{ee} resolution depending on m_{ee} .

Figure 6.12: Resolution of m_{ee} estimated as root mean square from reconstructed internal conversion events. A minimum of no, one, two, or three recurling tracks is requested in the vertex fit. In addition, the momentum resolution is shown in dependence of m_{ee} in intervals of 10 MeV.



(a) Sensitivity in branching fraction as expected in phase I.



(b) Sensitivity in branching fraction as expected in phase II.

Figure 6.13: Smallest expected upper limits on the branching fraction of $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ that the phase I and phase II *Mu3e* experiment is sensitive to. All tracks are considered in the vertex fit and the optimum m_{ee} combination is chosen.

With an expected increase of the muon stopping rate to $2 \cdot 10^9 \mu/s$, especially the accidental background increases. As the rate of Bhabha events is expected to scale linearly with the muon stopping rate, this type of background is scaled by a factor of twenty. This is a conservative approach as in phase II the vertex resolution improves due to a smaller target and changes in the beam profile, and consequently a better background rejection can be achieved. Other types of accidental background are not considered as no suitable simulation sample exists. Their contribution to the background might however increase to a size comparable to the contribution of Bhabha events because the rate of accidental background of three uncorrelated tracks scales with the muon beam rate squared.

It is further assumed that the timing and mass resolution are unchanged, thus the efficiencies for internal conversion events and dark photon events are unaffected. This neglects the fact, that the efficiency of reconstructing long tracks increases in phase II and as a result the two long track approach might become favourable.

In total, it is expected that the sensitivity is rather underestimated in this approach.

A total of $5.5 \cdot 10^{16}$ observed muon decays is assumed corresponding to a bit more than 300 days of data taking at $2 \cdot 10^9 \mu/s$. Figure 6.13 shows the expected upper limits on the signal branching fraction that *Mu3e* is sensitive to in each phase. For phase II, an improvement by a factor of five is expected for the more than 20 times larger number of muon decays.

6.5 EXPECTED SENSITIVITY TO THE KINETIC MIXING PARAMETER

The kinetic mixing parameter ε of the dark photon with the photon is directly related to the branching fraction of the $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays. In the work by Echenard, Essig and Zhong [118], $\mu \rightarrow e\nu\nu A'$ processes are studied with MADGRAPH5 and an empirical relationship between the branching fraction and the parameters ε and $m_{A'}$ is found:

$$\mathcal{B}(\varepsilon, m_{A'}) = \frac{1}{3 \cdot 10^{-19}} \left(\frac{\varepsilon}{0.1} \right)^2 \exp \left(\sum_{i=0}^5 a_i \left(\frac{m_{A'}}{\text{GeV}} \right)^i \right). \quad (6.17)$$

The coefficients a_i are determined in a fit.

In the accessible mass range of the *Mu3e* experiment, the dominant decay mode of the dark photon to SM particles is to an electron-positron pair. The branching fraction of $A' \rightarrow ee$ is therefore as-

sumed to be one⁶. Thus, equation 6.17 can also be applied to $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays. The kinetic mixing parameter can be estimated by

$$\varepsilon^2 = 3 \cdot 10^{-21} \exp\left(-\sum_{i=0}^5 a_i \left(\frac{m_{A'}}{\text{GeV}}\right)^i\right) \mathcal{B}. \quad (6.18)$$

The sensitivity to ε^2 is estimated with the expected sensitivity limits on the signal branching fraction from the previous section 6.4 using all types of reconstructed tracks in the vertex fit and the optimum combination for m_{ee} . The results for phase I of the *Mu3e* experiment are shown in figure 6.14a. For small masses, the experiment is sensitive to ε^2 around $3 \cdot 10^{-7}$. Again, a weaker sensitivity is observed at $m_{A'}$ around 5 MeV. From a dark photon mass of about 30 MeV onwards, the sensitivity worsens. At $m_{A'} = 80$ MeV, a ε^2 of about $3 \cdot 10^{-4}$ is expected to be in the reach of the experiment.

In phase II, *Mu3e* is sensitive to ε^2 in the range from $\mathcal{O}(10^{-7})$ to $\mathcal{O}(10^{-4})$ for the considered dark photon masses and thus to about a factor five smaller values than in phase I as expected from the increase in muon decays (see figure 6.14b).


6.6 SYSTEMATIC UNCERTAINTIES

A few sources of systematic uncertainties to the $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ sensitivities are presented here.


The largest uncertainty concerns the estimated amount of accidental background. This type of background is dominated by Bhabha events. In contrast for example to the branching fraction of internal conversion or other SM muon decays, no precise calculation or measurement for the occurrence of the accidental backgrounds exists. The occurrence has to be determined by means of detector simulation and thus bears a great potential for uncertainties. The amount of Bhabha events for instance depends greatly on the material distribution in the detector and the beam profile. In order to estimate the influence of the uncertainty in accidental background, the contribution by Bhabha events is altered.

Another systematic uncertainty might arise from the chosen bin size. The bin size of the m_{ee} spectra in the study presented here is limited by the method itself. The sample of reconstructed Bhabha events is small due to the low reconstruction efficiency for this type of back-

⁶ If the dark photon decays in addition to new light and dark particles, $\mathcal{B}(A' \rightarrow ee)$ becomes smaller than one, and the expected sensitivity to ε^2 weakens with a factor of $\mathcal{B}(A' \rightarrow ee)^{-1}$.



(a) Sensitivity to ϵ^2 as expected in phase I.



(b) Sensitivity to ϵ^2 as expected in phase II.


Figure 6.14: Smallest expected upper limits on the kinetic mixing parameter ϵ^2 in $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays that the phase I and phase II *Mu3e* experiment is sensitive to. All tracks are considered in the vertex fit and the optimum m_{ee} combination is chosen.

ground, and thus a too fine bin size would not result in a reliable p.d.f. Unfortunately, a bin size comparable to the size of the m_{ee} resolution has to be chosen. For the reconstruction in the later experiment, there are no restrictions on the choice of bin size. The influence of the bin size on the sensitivity is assessed by conducting a toy Monte Carlo study only on internal conversion background. The larger sample size allows to test a bin size of 0.5 MeV in addition to 1 MeV as used before.


Furthermore, the influence of cuts on the three-particle mass m_{eee} on the sensitivity to prompt dark photon decays is studied. A selection in m_{eee} might be applied on the filter farm level if the data rate turns out to be too high to be processed by the data collection server. Such a selection is however disfavoured as it affects not only $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ searches but also measurements of the radiative decays with internal conversion $\mu \rightarrow eee\nu\nu$ which also will be conducted with *Mu3e*.

6.6.1 Fraction of Bhabha Events

The fraction of Bhabha events per muon decay is varied in order to assess by how much the sensitivity changes if the accidental background is under- or overestimated. The sensitivity to the branching fraction of $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays and to the ε parameter are estimated as described previously (see section 6.4). The Bhabha fraction is additionally set to half and twice the nominal value of $b_{\text{Bhabha}} = 7.8 \cdot 10^{-5}$ as well as the Bhabha contribution to the background is completely taken out. The results are presented in figure 6.15. The expected sensitivity limits are shown for both branching fraction and the kinetic mixing parameter. For the vertex fit, all kinds of reconstructed tracks are utilized. Only the electron-positron pair candidate that yields the best sensitivity is considered. As the Bhabha contribution to the background is about three orders of magnitude smaller than the contribution by the dominant internal conversion events, the sensitivity is only marginally affected, if at all. No deviation between the various curves is observed that exceeds the statistical uncertainty. Accidental background from Bhabha events with an occurrence of $\mathcal{O}(10^{-10})$ per muon decay has thus no perceivable influence on the sensitivity. Likewise, accidental background from even rarer processes can be regarded as negligible. Problems would only arise if these processes exhibit peaks in the m_{ee} spectrum different from the maximum of the internal conversion background. Such a behaviour is not expected for any of the accidental combinations.




(a) Sensitivity in branching fraction.




(b) Sensitivity to the kinetic mixing parameter.

Figure 6.15: Smallest expected upper limits on the branching fraction of $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays and on the kinetic mixing parameter ϵ^2 at 90% CL that the $Mu3e$ experiment is sensitive to in 300 days of data taking at $10^8 \mu/s$. The contribution of Bhabha events is varied. The sensitivity is estimated using the optimum m_{ee} combination and vertices reconstructed from all tracks.



(a) Sensitivity in branching fraction.



(b) Sensitivity to the kinetic mixing parameter.

Figure 6.16: Smallest expected upper limits on the branching fraction of $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays and on the kinetic mixing parameter ϵ^2 at 90% CL that the *Mu3e* experiment is sensitive to in 300 days of data taking at $10^8 \mu/s$. An m_{ee} bin size of 0.5 MeV and 1 MeV is used. Background from Bhabha events is neglected. The sensitivity is estimated using the optimum m_{ee} combination and vertices reconstructed from all tracks.

6.6.2 *Bin Size*

The bin size of the m_{ee} spectra in the presented study is not optimal for the resonance search and is simply constrained by the limited sample size of reconstructed background events. The reconstruction efficiency for Bhabha events is particularly small. In order to build a reliable background p.d.f., a rather coarse bin size of 1 MeV is chosen. The sample of reconstructed internal conversion events is however larger than the one of Bhabha events and thus allows to test also a smaller bin size of 0.5 MeV. The influence of the bin size is therefore assessed using a background p.d.f. that only contains internal conversion events neglecting the anyhow minuscule Bhabha contribution.

In figure 6.16, estimates on the sensitivity to the branching fraction and kinetic mixing parameter are shown for bin sizes of 0.5 MeV and 1 MeV. Vertices are found using all kinds of reconstructed tracks. Again, the sensitivity is derived choosing the optimum electron-positron pair candidate.

With the finer bin size, the expected sensitivity limits improve over the whole mass range by 10 % to 40 %. Thus, all previously derived sensitivity estimates can be regarded to be in a realistic range although the bin size is not optimal, and are likely to improve in the later experiment. Though, a significant further improvement by choosing even smaller bin sizes is not expected as the investigated bin sizes are already in the same range as the m_{ee} resolution which varies around 1 MeV (see figure 6.12).

6.6.3 *Restrictions on the Three-Particle Mass*

Although currently not foreseen, a selection cut on the three-particle mass of the online reconstruction might be applied in order to reduce the output data rate. The data rate is dominated by internal conversion events. The corresponding m_{eee} spectrum is shown in figure 6.17b. The strongest suppressions are achieved for cuts up to about 60 MeV. Such a cut is not expected to harm $\mu \rightarrow eee$ searches. Searches for dark photon resonances from prompt $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays can however be severely affected. The effect of m_{eee} cuts on searches for prompt dark photon decays is estimated for cuts between 10 MeV and 90 MeV in steps of 10 MeV. The sensitivity is derived using all kinds of reconstructed tracks and by choosing the optimum m_{ee} combination.

The cuts alter not only the shape of the m_{ee} spectra but in particular the reconstruction efficiency. This is shown in figure 6.18 for dark

photon events and in figure 6.19 for internal conversion and Bhabha background. The m_{eee} spectra of all processes are shown in figure 6.17. As expected, increasing cuts on m_{eee} considerably reduce the reconstruction efficiency of dark photon events. Larger $m_{A'}$ are less affected than small ones. Likewise, the reconstruction efficiency of background events decreases corresponding to their respective m_{eee} distribution.

As the reconstruction efficiency vanishes for most processes at tight m_{eee} cuts, the procedure of estimating the sensitivity as used throughout this chapter finds its limits in some regions of the parameter space. In order to employ suitable p.d.f.s for all processes, it is required that the sum of weights after the m_{eee} cut equals at least one hundred times the average weight before cuts for each process. The resulting sensitivity to the signal branching fraction and to ϵ^2 are shown in figure 6.20. Starting from $m_{eee} > 70$ MeV, small values of $m_{A'}$ are getting out of the reach of the method.


At small $m_{A'}$, the sensitivity curves fan out with tighter m_{eee} cuts resulting in a weaker expected sensitivity. This is expected from the reconstruction efficiency of dark photon events. With increasing $m_{A'}$, the curves merge and are at about the same level. The curves for tighter cuts join at higher masses than the curves for looser cuts. As seen before, dark photon events with high $m_{A'}$ are simply less affected by the m_{eee} cuts. Additionally, the background decreases with tighter m_{eee} cuts. In fact, the sensitivity curves hint at an inversion. At large $m_{A'}$, tighter cuts result in a slightly improved expected sensitivity. An m_{eee} cut should thus be considered for analysis of prompt dark photon decays with larger $m_{A'}$.

The deviation from the sensitivity curve without m_{eee} cuts increases with tighter cuts. This is also observed in the reconstruction efficiency of dark photon events. The tail of the m_{eee} spectrum towards large three-particle masses in dark photon events is kinematically constrained to the muon mass⁷. Thus, the closer the m_{eee} cut is to the kinematic cutoff, the more severely the sensitivity is affected.


6.7 PROJECTIONS OF DARK PHOTON SEARCHES WITH THE MU3E EXPERIMENT

COMPARISON TO PREVIOUS STUDIES In the work by Echenard, Essig and Zhong [118], a first study of the sensitivity of the *Mu3e* experiment to prompt dark photon decays is conducted. The study is performed using a FASTSIM [215] simulation of a simplified *Mu3e*


⁷ Slightly larger reconstructed m_{eee} than the muon rest mass are possible because of the mass resolution.



(a) Dark photon events.




(b) Internal conversion events.



(c) Bhabha events.

Figure 6.17: Spectrum of the reconstructed three-particle mass of dark photon events as well as background events from internal conversion and Bhabha scattering using all kinds of reconstructed tracks.



(a) Both combinations of electron-positron candidates.





(b) Electron-positron pair with higher m_{ee} .(c) Electron-positron pair with lower m_{ee} .

Figure 6.18: Fraction of reconstructed events per generated $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decay for various cuts on the reconstructed three-particle mass. All kinds of reconstructed tracks are used in the vertex fit.




(a) Internal conversion events.




(b) Bhabha events.

Figure 6.19: Fraction of reconstructed events per generated internal conversion or Bhabha event for various cuts on the reconstructed three-particle mass. All kinds of reconstructed tracks are used in the vertex fit.



(a) Sensitivity in branching fraction.



(b) Sensitivity to the kinetic mixing parameter.

Figure 6.20: Smallest expected limits on the branching fraction of $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays and on the kinetic mixing parameter ϵ^2 at 90 % CL that the *Mu3e* experiment is sensitive to in 300 days of data taking at $10^8 \mu/s$. Various cuts on the three-particle mass m_{eee} are applied. The sensitivity is estimated using the optimum m_{ee} combination and vertices reconstructed from all tracks.

detector geometry. The geometry differs for instance in the composition of the pixel modules and the shape and material of the target. Furthermore, the fibre and tile detectors are not included. The total length of the detector is set to 180 cm in both phases of the experiment.

As potential sources of irreducible background, internal conversion events and radiative muon decays with photon conversion in the target are identified. The latter is estimated to be negligibly small in this study and is therefore left out. For accidental backgrounds, various combinations of multiple muon decays are considered some of which require a positron track to be misinterpreted as an electron track. Bhabha events are not examined as these would be beyond the scope of the simplified simulation. In the study presented here, Bhabha events are treated as the dominant source of accidental background and contributions from other processes are expected to be stronger suppressed and are thus neglected.


For phase I, a total number of observed muon decays of $1 \cdot 10^{15}$ is assumed in [118], and $5.5 \cdot 10^{16}$ in phase II. The phase I results of this study are translated accordingly.

In figure 6.21, the m_{ee} distributions of background events are shown in comparison. In the study by Echenard et al. [118], the spectra contain both combinations of electron and positron. For this study, either the combination with the higher or lower m_{ee} is considered. Therefore, the background is lower in this study. The contribution of Bhabha events that is considered here is about two orders of magnitude larger than what is estimated for accidental background in [118].


Both studies differ the most in mass resolution and signal reconstruction efficiency. In [118], a mass resolution of 0.2 MeV to 0.3 MeV is estimated, notably deviating from the resolutions of about 1 MeV estimated here. A comparison of the reconstruction efficiencies is shown in figure 6.21. The study by Echenard et al. [118] estimates considerably larger reconstruction efficiencies.

As a result, the expected sensitivity limits on prompt dark photon decays are weaker in the study presented here compared to the study by Echenard et al. [118]. The expected sensitivity limits at 95 % CL are shown in figure 6.22. The limits of all kinds of reconstructed tracks taking the optimum electron-positron pair candidate are used.


The study presented here extends to lower dark photon masses. This region in parameter space is however already excluded by existing experimental constraints.




(a) Background distribution estimated by [118] for $1 \cdot 10^{15}$ muon decays in phase I.




(b) Background distribution as estimated in this study for $1 \cdot 10^{15}$ muon decays in phase I.




(c) Background distribution estimated by [118] for $5.5 \cdot 10^{16}$ muon decays in phase II.



(d) Background distribution as estimated in this study for $5.5 \cdot 10^{16}$ muon decays in phase II.



(e) Signal reconstruction efficiency as estimated by [118].



(f) Signal reconstruction efficiency as estimated in this study.

Figure 6.21: Comparison of the m_{ee} spectra of background events as assumed in [118] (left-hand side) and this study (right-hand side) as well as of the reconstruction efficiency of prompt dark photon decays. $1 \cdot 10^{15}$ ($5.5 \cdot 10^{16}$) observed muon decays at a stopping rate of $10^8 \mu/s$ ($2 \cdot 10^9 \mu/s$) are assumed in phase I (II). The background distributions of this study are shown with bins of 0.2 MeV to be comparable to the distributions in [118]. The sensitivity is however estimated with a bin size of 1 MeV as before.

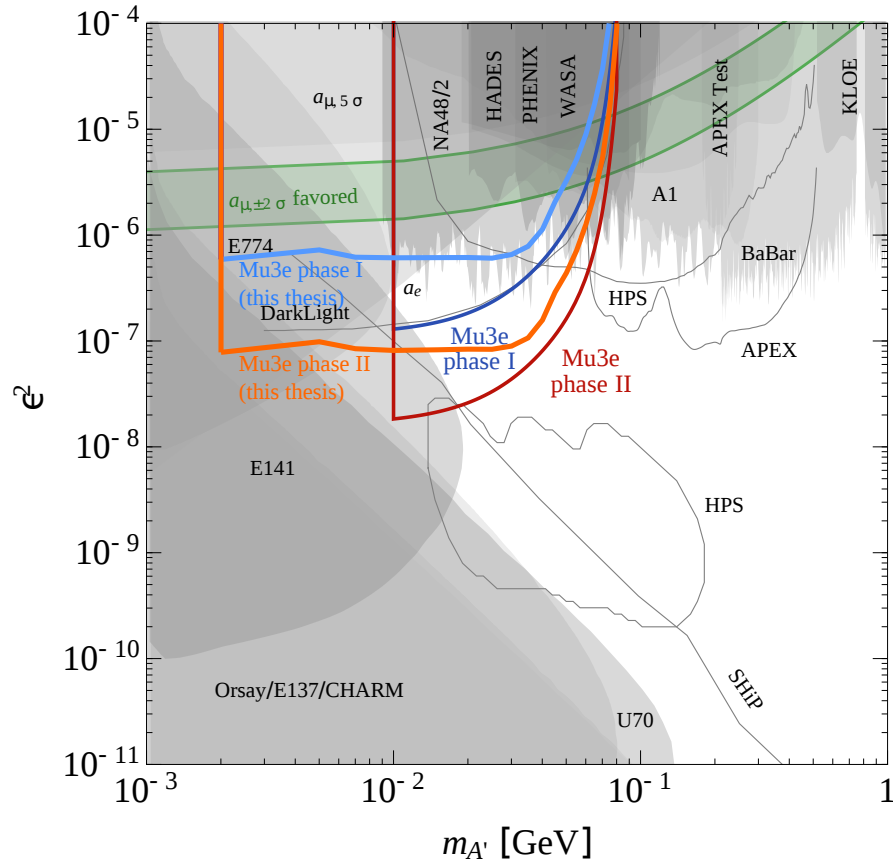


Figure 6.22: Expected sensitivity limits at 95 % CL on the kinetic mixing parameter ϵ^2 in $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays as estimated in [118] (blue and red line) and by this study (light blue and orange line). Results are shown for $1 \cdot 10^{15}$ muon decays in phase I and $5.5 \cdot 10^{16}$ in phase II. The plot is adapted from [118].




Figure 6.23: Expected sensitivity limits at 90 % CL on the kinetic mixing parameter ϵ^2 in $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays in phase I and phase II of the *Mu3e* experiment and current upper experimental limits at 90 % CL. All tracks are used in the vertex fit. The plot is adapted from [216].

COMPARISON TO CURRENT EXPERIMENTAL LIMITS ON VISIBLY DECAYING DARK PHOTONS In figure 6.23, the expected sensitivity limits of *Mu3e* on the kinetic mixing parameter ϵ^2 at 90 % CL are shown in comparison to current experimental upper limits. Details on some of the experiments are given in section 1.4.1. For phase I, the total number of observed muon decays is set to $2.592 \cdot 10^{15}$, and to $5.5 \cdot 10^{16}$ in phase II.

Already in phase I, the *Mu3e* experiment has the possibility to improve on current limits, in particular in a mass region from 10 MeV to 40 MeV. In phase II, the sensitivity extends further. In a mass region of 10 MeV to 60 MeV, *Mu3e* is competitive to other experiments with an expected sensitivity to ϵ^2 as small as $6 \cdot 10^{-8}$.

Typically, *Mu3e* is sensitive to a comparable mass range as experiments searching for dark photon production in pion decays, though for *Mu3e* a sensitivity to smaller kinetic mixing parameters is expected than the current best limits by the *NA48/2* experiment. This is due to the large number of expected observed muon decays. On the left-hand side, i. e. at lower $m_{A'}$, the sensitive parameter space of *Mu3e* is flanked by beam dump experiments (*E141*, *E137* and *E774*) as well as by constraints from measurements of the anomalous magnetic moment of the electron $(g-2)_e$.

6.8 SUMMARY ON ELECTRON-POSITRON RESONANCE SEARCHES

SUMMARY In this chapter, a strategy for electron-positron resonance searches in reconstructed three particle events is elaborated with regard to searches of promptly decaying dark photons emitted in muon decays. Dark photon events would appear as a resonance in the spectrum of the invariant mass m_{ee} of the electron-positron pair candidates. As the final state of the $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decay contains three charged particles, the search is conducted on the same data sample as $\mu \rightarrow eee$ searches. Thus, no modifications to the data acquisition are required. As background sources, the radiative muon decay with internal conversion $\mu \rightarrow eee\nu\nu$ and accidental combinations of three charged tracks are identified. In the case of accidental background, events in which a positron from $\mu \rightarrow e\nu\nu$ scatters off an electron via Bhabha scattering and in which the two are then combined with another random charged particle are found to dominate.

Toy Monte Carlo studies are performed for dark photon masses $m_{A'}$ between 2 MeV and 80 MeV. Optimum sensitivities are achieved when no selection on the type of reconstruction, i. e. short or long tracks, is applied. Furthermore, this study has shown that it is best to search for dark photon masses up to about 45 MeV in the m_{ee} spectrum of the electron-positron pair candidate with the lower m_{ee} , and for larger $m_{A'}$ in the higher m_{ee} pair candidate spectrum. A few systematic uncertainties are studied. The sensitivity is found to be robust against variations of the accidental background produced by Bhabha events to at least a factor of 100 %. This is expected as the Bhabha contribution to the background spectrum is about three orders of magnitude smaller than the dominating contribution from internal conversion events.

The bin size in this study is constrained by the limited size of the background samples. It is shown that the sensitivities improve to some extent with a finer bin size but the sensitivities estimated with the coarser bin size are still in a viable range.

Furthermore, it is demonstrated that potential selections on the three-particle mass m_{eee} at filter farm level can have a severe impact on the dark photon sensitivities, depending of course on the actual m_{eee} selection.

In the first phase of the *Mu3e* experiment, the expected sensitivity to $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ decays ranges from branching fractions as

small as $5 \cdot 10^{-9}$ at 90% CL for small $m_{A'}$ down to $3 \cdot 10^{-12}$ for the largest investigated $m_{A'}$ assuming 300 days of data taking at $10^8 \mu/s$. The sensitivity is highly dependent on the reconstruction efficiency for dark photons events which alters with $m_{A'}$ as expected, as well as on the background which is highest for low $m_{A'}$.

Under the assumption that dark photons in this mass range decay exclusively to electron-positron pairs, the sensitivity to the kinetic mixing parameter ε can be estimated. At low $m_{A'}$, the expected sensitivity limits on ε^2 are $3 \cdot 10^{-7}$ at 90% CL in phase I. The limits increase to $3 \cdot 10^{-4}$ at large $m_{A'}$.

Thus, the potential of the *Mu3e* experiment to investigate currently unexplored dark photon parameter space is demonstrated. Already in phase I, *Mu3e* can improve on existing limits by *NA48/2* and beam dump experiments. With phase II, even smaller values of ε come into reach.

DISCUSSION AND OUTLOOK The toy Monte Carlo study as performed here embodies some idealizations and simplifications and is only suited to give a preliminary estimate on the sensitivity in the current pre-construction phase of the *Mu3e* experiment.

One aspect to be further studied concerns the efficiency of the on-line selection of three-track candidates on the filter farm. So far, these studies are only conducted for $\mu \rightarrow eee$ signal events. As the kinematics of muon decays with prompt dark photon decay differs from $\mu \rightarrow eee$, this study has to be repeated for dark photon events and various dark photon masses.

The accuracy of the background prediction can be improved by taking next-to-leading order calculations of the matrix element of the dominating radiative decay with internal conversion into account [194, 195]. The corrections with respect to leading order calculations range around 1%. For a more precise estimate of the accidental background, other types of accidental background need to be simulated as well. The occurrence of these processes often depends on the material distribution in the detector—for instance in scenarios with Bhabha scattering or photon conversion—and is thus difficult to simulate reliably. But once the experiment is running, accidental background can easily be measured by inverting the timing cut for the coincidence of the three tracks.

In simulation, the position of the active detector elements and the target are known with full precision, but in the later experiment there will be uncertainties from the alignment. In addition, the beam profile and muon stopping distribution cannot be perfectly

modelled. Since these factors influence the reconstructed m_{ee} distribution, a fit with a fixed background shape as in this toy Monte Carlo study is not applicable. It needs to be replaced by a parametrized background function which allows to adapt to detector effects to some extent. Typically, the background is described by a polynomial. Likewise, the signal distribution can be parametrized, for instance by a double Gaussian distribution or Crystal Ball function. The search is then conducted on select invariant mass windows. The background should be smooth in each of the search windows and the window should be wide enough to cover the signal with some safety margin.

The sensitivity can be further enhanced by applying weights. As the sensitivity depends on the resolution of the invariant mass of the electron-positron pair, one possibility is to weight the events with the inverse of it. This circumvents the signal efficiency loss from choosing only events with two long tracks for the electron-positron pair.

Similarly, instead of choosing either the reconstructed electron-positron pair with the higher or lower invariant mass, both can be combined with a weight corresponding to the probability to cover the e^+e^- -pair from the dark photon decays with the respective combination. This would increase the signal efficiency without increasing the underlying background.

For phase II, the sensitivity is expected to improve by about a factor of five due to the larger number of muons to be observed. As also the muon beam rate is increased, especially the accidental background needs to be re-evaluated. Background from three uncorrelated tracks, e.g. two positrons from a Michel decay and an electron from a photon conversion or Bhabha scattering event, scales with the beam rate to the square and might be an important contribution to the background of prompt dark photon searches. The contribution of Bhabha events to the background also rises, in this case about linearly with the beam rate. The sensitivity will be further influenced by the detector upgrade. A smaller beam profile and target will improve the vertex position resolution. Installing additional recoil stations enhances the efficiency to reconstruct long tracks. These are many unknowns so that only rough estimates on the phase II sensitivity are given in this study.

Apart from prompt dark photon decays, one could think of two other channels to search for dark photons with $Mu3e$. On the one hand, the dark photon could decay preferentially or ex-

clusively to invisible particles with the only detectable particle being a single positron. The momentum distribution of this positron is unfortunately very broad and has no characteristic features. It is for instance easily mistaken for a slightly misaligned detector. Thus, searches for invisible dark photon decays in *Mu3e* are deemed unfeasible.

On the other hand, searches for dark photons with lifetime are very well feasible. As long as the dark photon decays to an electron-positron pair within the volume enclosed by the second vertex layer, the event can be reconstructed, at least in the current reconstruction. This requires already changes to the online event selection. With further modifications to the reconstruction, in principle also $A' \rightarrow ee$ events between the second vertex layer and the fibre detector or shortly outside the outer pixel layers can be detected. Such searches would allow to test even smaller values of ε .

The reach in ε for displaced vertex searches is in general restricted by statistics and background suppression towards small ε , and by the resolution of the vertex position towards larger ε . With a potential *Mu3e-Gamma* upgrade, the reach can be extended to longer decay lengths and thus smaller ε . The feasibility of such searches has to be determined in detailed studies.

SUMMARY, CONCLUSION, AND OUTLOOK

Searches for lepton flavour violation in the charged lepton sector in decays like $\mu \rightarrow eee$ are powerful tools to discover New Physics beyond the Standard Model. In the Standard Model, lepton flavour is conserved. And even in extensions incorporating neutrino masses and mixing, decays like $\mu \rightarrow eee$ are suppressed to unobservable levels. Consequently, an observation of $\mu \rightarrow eee$ would be a clear indication of the existence of New Physics.

The upcoming *Mu3e* experiment has an ambitious sensitivity goal for the $\mu^+ \rightarrow e^+e^-e^+$ decay of ultimately one signal event in $1 \cdot 10^{16}$ muon decays. In a first phase, a single event sensitivity of about $2 \cdot 10^{-15}$ is envisaged.

This requires for instance a data acquisition system capable of coping with high data rates. Parts of the data acquisition firmware on the front-end of the pixel detector are developed and tested in the context of this thesis.

Furthermore, sensitivity studies on the $\mu \rightarrow eee$ signal decay in the first phase of the experiment are performed. For the first time, these studies are expanded to $\mu \rightarrow eee$ decays mediated via various effective operators in order to account for the specific kinematics of the potential new interactions. It is shown that a background-free operation is feasible and that the envisaged sensitivity goal can be realized.

Although *Mu3e* is a specialized experiment, the physics programme is not limited to $\mu \rightarrow eee$ searches. With a precise measurement of low momentum electrons and a high acceptance detector combined with an immense number of expected muon decays, opportunities to investigate additional channels arise. In the presented work, the sensitivity to two other channels involving New Physics is estimated: lepton flavour violating two-body decays of the muon $\mu \rightarrow eX$, and muon decays involving promptly decaying dark photons $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$. In both cases, it is demonstrated that the *Mu3e* experiment can set competitive limits.

7.1 FIRMWARE ON THE PIXEL DETECTOR FRONT-END

The first implementation of a firmware-based hit-sorting algorithm for the pixel detector front-end of the *Mu3e* experiment is developed and evaluated. The pixel sensors feature an on-chip readout system but due to the readout sequence, the time order of the hits is not strictly conserved. The time order can be restored on the front-end FPGA by making use of the time stamp information that is part of the hit data. This is a crucial part of the pixel detector data acquisition system as it is essential for the subsequent event filtering to have time ordered hit information.

The developed hit sorter is capable to process a continuous stream of hit data and is robust against large variations of time stamps. It relies on a random-access memory used as a circular buffer to which the hits are written in correct time order and then read back. The suitability for the *Mu3e* experiment is evaluated with the MUPIX TELESCOPE equipped with either of the latest two *HV-MAPS* prototypes for *Mu3e*: the MUPIX7 or the MUPIX8 sensor. The telescope represents a small scale realization of the later pixel detector front-end and is thus suited for integration tests of the hit sorter and of the front-end firmware in general.

The first version of the hit sorter is implemented for the use with a MUPIX7 based telescope. The MUPIX7 is a small scale *HV-MAPS* prototype that features for the first time an on-chip readout control. This setup is extensively tested and operated at numerous test beams. The hit sorter has demonstrated an excellent efficiency of 99.995 % on average in a high beam rate measurement with an electron beam at MAMI at an average hit rate of up to 1.7 MHz per sensor.

With the next prototype, MUPIX8, the hit sorter is upgraded. The MUPIX8 is a large scale sensor and transmits data on up to three links. Therefore, the number of input channels to the hit sorter and the size of the sorting memory are increased. In fact, the number of input channels is made scalable so that it can be easily adapted to the needs of the front-end board.

In a first test at an electron beam at DESY with a total hit rate of (20-40) kHz, an average efficiency of 99.97 % is measured. The dominating contribution to the inefficiency is caused by incoming hits with deviant time stamp. The increase in inefficiency compared to the MUPIX7 implementation is consistent with the larger size of the MUPIX8 sensor.

As the memory has a fixed size, only a limited number of hits with the same time stamp value can be processed. Inefficiencies that

result from this limitation are found to be negligible at least at the investigated hit rates. A high rate test is not yet performed.

Furthermore, the requirements on the hit sorter in terms of bandwidth and efficiency are assessed by means of detector simulation. It is found that the latest implementation of the hit sorter can meet all these requirements with only minor modifications.

The developed hit sorter is successfully operated in the MUPIX TELESCOPE environment with actual pixel detector prototypes and the suitability of the hit sorting concept on the front-end FPGA for the phase I of the *Mu3e* experiment is proven. In a next step, the hit sorter is migrated to a front-end board prototype and combined with further firmware components, thus paving the way to the final pixel detector front-end.

The hit sorter concept is not limited to the use in the *Mu3e* experiment but can also be adopted in other experiments that make use of *HV-MAPS* with an on-chip readout system. One example is the upcoming *P2* experiment for precision measurements of the weak mixing angle at MESA [217]. Furthermore, the hit sorter might be applied in the luminosity detector of the upcoming *PANDA* experiment at FAIR [218].

7.2 SENSITIVITY STUDIES ON NEW PHYSICS

SIGNAL DECAY $\mu \rightarrow eee$ The sensitivity of the experiment to the signal decay $\mu \rightarrow eee$ is determined by the detector acceptance as well as by the capability to suppress background and thus by the measurement resolution.

In the presented work, the possibility to operate free of background in $\mu \rightarrow eee$ searches when exploiting the excellent momentum resolution of reconstructed recurling tracks is confirmed. In the case of phase space distributed signal events, an expected single event sensitivity of $2.3 \cdot 10^{-15}$ in 300 days of data taking at $10^8 \mu/s$ in the first phase of the experiment is accomplished. This poses an improvement in sensitivity by a factor of 200 compared to the preceding *SINDRUM* experiment.

As the acceptance and thus sensitivity greatly depends on the kinematics of the decay, the actual lepton flavour violating interaction needs to be accounted for. For this purpose, an effective field theory approach is established. The sensitivity is estimated for various effective operators. In the case of $\mu \rightarrow eee$ decays mediated via dipole operators, a decrease in acceptance and reconstruction efficiency is observed. This results in a weaker expected single event sensitivity

of $3.7 \cdot 10^{-15}$. Signal decays mediated by four-fermion operators in contrast yield a slightly enhanced acceptance. Consequently, an improved single event sensitivity of $2.0 \cdot 10^{-15}$ is estimated.

Furthermore, characteristic footprints in Dalitz plots of the invariant mass of the electron-positron combinations as well as the angular distributions are observed. These can be used to pin down the type of interaction and chiralities of the involved fermion fields in the case a sufficient number of events is observed.

The effective theory approach enables comparisons between experiments searching for lepton flavour violation in the decays $\mu \rightarrow eee$, $\mu \rightarrow e\gamma$ and in muon to electron conversion on nuclei $\mu N \rightarrow eN$. In the case of dipole operators, existing limits measured by the *MEG* experiment are already more stringent than limits expected from *Mu3e* in phase I. The prospects of the *MEG II* experiment also exceed the *Mu3e* experiment in phase II in terms of sensitivity to dipole operators. In the case of four-fermion operators of muon to electron transitions, the *Mu3e* experiment yields the best expected sensitivity compared to experiments of the other lepton flavour violation observables whereas the best sensitivity to four-fermion operators involving quark fields is expected for $\mu N \rightarrow eN$ searches. This comparison illustrates the power of the effective theory approach and the importance to perform lepton flavour violation searches in various channels. In the event a signal is observed in one or more experiments, a comparison based on effective operators is capable to reveal the type of the involved New Physics.

The presented effective theory approach is not only suited for the *Mu3e* experiment but also for example for the potential *Mu3e-Gamma* upgrade. Compared to *MEG* and *MEG II*, *Mu3e-Gamma* has a larger acceptance and thus an improved sensitivity to the chiralities of the involved operators.

LEPTON FLAVOUR VIOLATING TWO-BODY DECAYS Searches for lepton flavour violating two-body decays are motivated by familons, Nambu-Goldstone bosons arising from introducing a broken flavour symmetry. These new particles can be massive and can have scalar and pseudo-scalar couplings.

In the context of this thesis, a search strategy for $\mu \rightarrow eX$ decays in the *Mu3e* experiment is developed and the sensitivity is estimated for the first time.

The signal decay is identified by an excess on top of the momentum spectrum of positrons from Standard Model muon decays. The high muon decay rate in *Mu3e* does not allow to keep the data of every

event. Therefore, every particle is reconstructed online and only data of potential $\mu \rightarrow eee$ candidates is stored for offline analysis. In addition, histograms of the online reconstructed momentum and polar angle of all particles are stored. These histograms are the input to $\mu \rightarrow eX$ searches.

The sensitivity is estimated in a toy Monte Carlo study. In the baseline concept, only short tracks are reconstructed on the filter farm. In this case, upper branching fraction limits of $1.5 \cdot 10^{-8}$ at 90 % CL can be expected for familon masses m_X up to 90 MeV assuming 300 days of data taking at $10^8 \mu/s$ in the first phase of the experiment. Taking also the polar angle into account, a 10 % improvement in sensitivity can be achieved. In addition, $\mu \rightarrow eX$ decays with pseudo-scalar couplings can be investigated.

As a consequence of the constant improvement of the performance of GPUs, online reconstruction of long tracks might come into reach. In this case, the sensitivity is improved by a factor of two.

The characteristic edge in the momentum spectrum of positrons from the dominant Michel decay is used in momentum calibration if alternatives like calibration with Mott or Bhabha scattering cannot be employed. This turns $\mu \rightarrow eX$ searches at small masses m_X unfeasible. In this work, a modified strategy for the calibration is investigated. The calibration is performed with the shape of the momentum spectrum leaving out a momentum region around the potential $\mu \rightarrow eX$ signal. In this way, the whole mass range can be investigated with the drawback of a reduced sensitivity at small m_X .

It is thus shown that the *Mu3e* experiment has a high sensitivity to $\mu \rightarrow eX$ decays. It can be expected to improve the currently most stringent limits on massive X by the *TWIST* experiment by two to three orders of magnitude. At the moment, there is no other experiment in operation or planning with a comparable sensitivity.

Besides the $\mu \rightarrow eX$ channel, familons can also be investigated in $\mu \rightarrow eee$, $\mu \rightarrow eeeee$, and $\mu \rightarrow e\gamma X$. If the familon in the first channel decays immediately, these searches can be performed with the data of the standard $\mu \rightarrow eee$ searches. Searches for displaced $X \rightarrow ee$ vertices however require modifications to the data acquisition. The acceptance for $\mu \rightarrow eeeee$ events is low as the decay electrons have typically a low momentum, but can be enhanced by lowering the magnetic field. A search for $\mu \rightarrow e\gamma X$ is not feasible with *Mu3e* because of the extremely low conversion probability of the additional photon. These decays can be studied in the potential *Mu3e-Gamma* upgrade.

PROMPT DECAYS OF DARK PHOTONS Dark photons are proposed particles that kinematically mix with the photon of the Standard Model and can serve as a messenger to a dark sector. Dark photons can thus couple to charged leptons and decay for instance into an electron-positron pair. In the context of this thesis, the sensitivity of the *Mu3e* experiment to promptly decaying dark photons emitted in muon decays $\mu \rightarrow e\nu\nu(A' \rightarrow ee)$ is estimated for the first time with a realistic detector simulation.

These searches are performed with the same data as in the case of $\mu \rightarrow eee$ searches. Dark photons can be identified by a resonance in the invariant mass spectra of electron-positron pairs at a mass $m_{A'}$. In the presented study, dark photon masses between 2 MeV and 80 MeV are investigated. The background is dominated by the radiative muon decay with internal conversion $\mu \rightarrow eee\nu$ which has the same final state.

As in the case of $\mu \rightarrow eX$ searches, the sensitivity is estimated by means of toy Monte Carlo studies. The phase I *Mu3e* experiment is found to be sensitive to branching fractions between $5 \cdot 10^{-9}$ at 90 % CL at dark photon masses of a few MeV and $3 \cdot 10^{-12}$ at a dark photon mass of 80 MeV. The sensitivity can be expected to improve by a factor of five in the second phase of the experiment.

The presented sensitivity study has thus proven that the sensitivity of the *Mu3e* experiment is competitive compared to existing limits and that a currently unexplored dark photon parameter space can be tested. In particular, the sensitivity of *Mu3e* exceeds current bounds of the *NA48/2* experiment in which dark photons produced in neutral pion decays are investigated. The search for prompt dark photon decays requires no modifications to the data acquisition.

Moreover, dark photon decays at displaced vertices and thus smaller mixing parameters can be studied with *Mu3e*. The implementation of dark photons into the simulation and the toy Monte Carlo study performed in this thesis can serve as a basis for these studies. A search for events at displaced vertices requires modifications to the data acquisition and online reconstruction. Studies are ongoing.

In addition, there is the opportunity to search for dark photons at displaced vertices with the potential *Mu3e-Gamma* upgrade. Longer decay lengths and thus smaller mixing parameters come into reach.

The presented work contributes to the technical realization of the data acquisition of the *Mu3e* experiment. In addition, the feasibility of a background-free $\mu \rightarrow eee$ search in phase I with the envisaged signal sensitivity is confirmed. Furthermore, the range of observables is extended and competitive sensitivities to further channels involving New Physics beyond the Standard Model are predicted. All in all, the *Mu3e* experiment offers an exciting physics programme and a great impact on many models of physics beyond the Standard Model can be expected.

Part IV

APPENDIX

DATA ACQUISITION

A.1 MUIX DATA PROTOCOL

The MUIX sensors continuously stream zero-suppressed, 8 bit/10 bit encoded serial data over a fast 1.25 Gbit/s LVDS link. The data output format of the MUIX7 prototype is given in table A.1, the output format of MUIX8 in table A.2.

After power-up or reset, a synchronization pattern is transmitted that allows the following data acquisition chain to align to the word boundary. For this purpose, the comma word K.28.5 is sent repeatedly. After an adjustable number of synchronization words, the actual data output starts.

In the case of MUIX7, at first the binary counter is output, followed by some comma words. Then a sequence of hit data follows. The information of each hit is terminated by the comma word K.28.1. When all hits in the EoC logic are transferred, a new readout cycle starts again with the binary counter. In the case that no hit data is available, only binary counter information and comma words are sent. There is thus a constant data stream from the sensor. If always only one hit per readout cycle is transmitted, about 15 % of the band width is used for transmission of hit data. The remaining bandwidth is used for header information as well as for 8 bit/10 bit encoding. Hit rates of 7.8 MHz can be transferred in this case. At maximum, one hit of each column, thus 56 hits in total, is transmitted in one readout cycle. In this case, the bandwidth usage increases to 76 % and the hit rate to 30 MHz. This is the maximal hit output rate of the MUIX7.

At the beginning of a readout cycle on the MUIX8, K.28.5 comma words are sent. Then follows an identifier for the sub-matrix of the sensor indicated by the comma word K.28.0. Next, the binary counter is transmitted followed by a sequence of hit data. In MUIX8, the size of the time stamp is increased compared to MUIX7 and information about the analogue pulse is available in addition. Therefore, no comma word is transmitted in between hits. In the case of no hits, the matrix identifier and comma words are sent. The binary counter is unfortunately omitted due to a bug in the state ma-

Cycle	Data	Comment	
0	K.28.5	Synchronization pattern	
1	K.28.5		
...	K.28.5		
$N_{\text{res}} - 1$	K.28.5		
N_{res}	Gray encoded binary counter[7:0]	Header: time information	
$N_{\text{res}} + 1$	Binary counter[7:0]		
$N_{\text{res}} + 2$	Binary counter[15:8]		
$N_{\text{res}} + 3$	Binary counter[23:16]		
$N_{\text{res}} + 4$	K.28.5		
$N_{\text{res}} + 5$	K.28.5		
$N_{\text{res}} + 6$	K.28.5		
$N_{\text{res}} + 7$	K.28.5		
$N_{\text{res}} + 8$	K.28.5		
$N_{\text{res}} + 9$	K.28.5		
$N_{\text{res}} + 10$	K.28.5		
$N_{\text{res}} + 11$	K.28.5		
$N_{\text{res}} + 12$	Row		Data: hit information
$N_{\text{res}} + 13$	Column		
$N_{\text{res}} + 14$	Time stamp		
$N_{\text{res}} + 15$	K.28.1		
...	Row		
...	Column	Repeat for up to 56 hits then return to cycle N_{res}	
...	Time stamp		
...	K.28.1		

Table A.1: Data output protocol of the MUPix7 prototype [219]. The synchronization sequence is transmitted once after reset or power-up. The length N_{res} is adjustable.

Cycle	Data	Comment
0	K.28.5	Synchronization
1	K.28.5	pattern
...	K.28.5	
$N_{\text{res}} - 1$	K.28.5	
N_{res}	K.28.5	Header:
$N_{\text{res}} + 1$	K.28.5	
$N_{\text{res}} + 2$	K.28.5	
$N_{\text{res}} + 3$	K.28.5	
$N_{\text{res}} + 4$	K.28.5	
$N_{\text{res}} + 5$	K.28.5	
$N_{\text{res}} + 6$	K.28.5	
$N_{\text{res}} + 7$	K.28.5	
$N_{\text{res}} + 8$	K.28.5	
$N_{\text{res}} + 9$	K.28.5	
$N_{\text{res}} + 10$	K.28.5	
$N_{\text{res}} + 11$	K.28.5	
$N_{\text{res}} + 12$	K.28.0	matrix information
$N_{\text{res}} + 13$	Matrix identifier	
$N_{\text{res}} + 14$	K.28.0	
$N_{\text{res}} + 15$	Matrix identifier	
$N_{\text{res}} + 16$	Binary counter[23:16]	time information
$N_{\text{res}} + 17$	Binary counter[15:8]	
$N_{\text{res}} + 18$	Binary counter[7:0]	
$N_{\text{res}} + 19$	Gray encoded binary counter[7:0]	
$N_{\text{res}} + 20$	Time stamp[9:2]	Data:
$N_{\text{res}} + 21$	Time stamp[1:0] & Charge[5:0]	hit information
$N_{\text{res}} + 22$	Column	
$N_{\text{res}} + 23$	Row	
...	Time stamp[9:2]	Repeat for up to 48/32 hits
...	Time stamp[1:0] & Charge[5:0]	then return to cycle N_{res}
...	Column	
...	Row	

Table A.2: Data output protocol of the MUPIX8 prototype [220]. The synchronization sequence is transmitted once after reset or power-up. The length N_{res} is adjustable. The three matrices are identified by a fixed pattern, i. e. 0xAA, 0xBB and 0xCC. Matrix C has less columns (32 instead of 48) and therefore a maximum of 32 hits is transmitted in one readout cycle. Charge denotes the analogue pulse information.

chine. It is planned to output also the binary counter in the case of no hits in the following prototypes.

If only one hit is transferred per readout cycle, the bandwidth usage is 13 % and the hit output rate 5.2 MHz. Using the three output links in parallel increases the hit rate to 16 MHz. At maximum, 32 or 48 hits are sent out per readout cycle. In this case, the bandwidth usage becomes 71 % and the hit rate 28 MHz when operating at a single merged link. With three parallel links, the hit output rate increases to 84 MHz.

THE SIGNAL DECAY MEDIATED VIA DIFFERENT EFFECTIVE OPERATORS

B.1 EFFECTIVE OPERATORS

In the work by Kuno and Okada [46, 62], the operator basis for muon to electron transitions in effective theories is identified up to dimension six. The corresponding Lagrangian for $\mu \rightarrow eee$ decays is

$$\begin{aligned} \mathcal{L}_{\mu \rightarrow eee} = & -\frac{4G_F}{\sqrt{2}} \left[m_\mu A_R \bar{\mu}_R \sigma^{\mu\nu} e_L F_{\mu\nu} + m_\mu A_L \bar{\mu}_L \sigma^{\mu\nu} e_R F_{\mu\nu} \right. \\ & + g_1 (\bar{\mu}_R e_L) (\bar{e}_R e_L) + g_2 (\bar{\mu}_L e_R) (\bar{e}_L e_R) \\ & + g_3 (\bar{\mu}_R \gamma^\mu e_R) (\bar{e}_R \gamma_\mu e_R) + g_4 (\bar{\mu}_L \gamma^\mu e_L) (\bar{e}_L \gamma_\mu e_L) \\ & + g_5 (\bar{\mu}_R \gamma^\mu e_R) (\bar{e}_L \gamma_\mu e_L) + g_6 (\bar{\mu}_L \gamma^\mu e_L) (\bar{e}_R \gamma_\mu e_R) \\ & \left. + \text{h.c.} \right], \end{aligned} \quad (\text{B.1})$$

with the Fermi coupling constant G_F and the field strength tensor $F_{\mu\nu}$ for photons. μ and e are the fermion fields in their left- (index L) and right-handed (index R) projection.

The operator basis is the same as used in equation 4.1. The coefficients $A_{R,L}$ and g_i of equation B.1 can be translated to the coefficients C of equation 4.1 as follows [196]

$$\begin{aligned} -\frac{4G_F}{\sqrt{2}} m_\mu A_R &= \frac{1}{\Lambda^2} e m_\mu (C_D^R)^*, \\ -\frac{4G_F}{\sqrt{2}} m_\mu A_L &= \frac{1}{\Lambda^2} e m_\mu (C_D^L)^*, \\ -\frac{4G_F}{\sqrt{2}} g_1 &= \frac{1}{\Lambda^2} (C_S^{RR})^*, \\ -\frac{4G_F}{\sqrt{2}} g_2 &= \frac{1}{\Lambda^2} (C_S^{LL})^*, \\ -\frac{4G_F}{\sqrt{2}} g_3 &= \frac{1}{\Lambda^2} (C_V^{RR})^*, \\ -\frac{4G_F}{\sqrt{2}} g_4 &= \frac{1}{\Lambda^2} (C_V^{LL})^*, \\ -\frac{4G_F}{\sqrt{2}} g_5 &= \frac{1}{\Lambda^2} (C_V^{RL})^*, \\ -\frac{4G_F}{\sqrt{2}} g_6 &= \frac{1}{\Lambda^2} (C_V^{LR})^*, \end{aligned} \quad (\text{B.2})$$

with e being the electron charge.

For the given set of operators, a differential branching fraction for $\mu^+ \rightarrow e^+e^-e^+$ decays is derived in [46, 62]

$$\begin{aligned} \frac{d\mathcal{B}_{\mu^+ \rightarrow e^+e^-e^+}}{dx_1 dx_2 d(\cos \tilde{\vartheta}) d\tilde{\varphi}} &= \frac{3}{2\pi} \left[\right. \\ &C_1 \alpha_1(x_1, x_2) (1 + P_\mu \cos \tilde{\vartheta}) + C_2 \alpha_1(x_1, x_2) (1 - P_\mu \cos \tilde{\vartheta}) \\ &+ C_3 [\alpha_2(x_1, x_2) + P_\mu \beta_1(x_1, x_2) \cos \tilde{\vartheta} + P_\mu \gamma_1(x_1, x_2) \sin \tilde{\vartheta} \cos \tilde{\varphi}] \\ &+ C_4 [\alpha_2(x_1, x_2) - P_\mu \beta_1(x_1, x_2) \cos \tilde{\vartheta} - P_\mu \gamma_1(x_1, x_2) \sin \tilde{\vartheta} \cos \tilde{\varphi}] \\ &+ C_5 [\alpha_3(x_1, x_2) + P_\mu \beta_2(x_1, x_2) \cos \tilde{\vartheta} + P_\mu \gamma_2(x_1, x_2) \sin \tilde{\vartheta} \cos \tilde{\varphi}] \\ &+ C_6 [\alpha_3(x_1, x_2) - P_\mu \beta_2(x_1, x_2) \cos \tilde{\vartheta} - P_\mu \gamma_2(x_1, x_2) \sin \tilde{\vartheta} \cos \tilde{\varphi}] \\ &+ C_7 [\alpha_4(x_1, x_2) (1 - P_\mu \cos \tilde{\vartheta}) + P_\mu \gamma_3(x_1, x_2) \sin \tilde{\vartheta} \cos \tilde{\varphi}] \\ &+ C_8 [\alpha_4(x_1, x_2) (1 + P_\mu \cos \tilde{\vartheta}) - P_\mu \gamma_3(x_1, x_2) \sin \tilde{\vartheta} \cos \tilde{\varphi}] \\ &+ C_9 [\alpha_5(x_1, x_2) (1 + P_\mu \cos \tilde{\vartheta}) - P_\mu \gamma_4(x_1, x_2) \sin \tilde{\vartheta} \cos \tilde{\varphi}] \\ &+ C_{10} [\alpha_5(x_1, x_2) (1 - P_\mu \cos \tilde{\vartheta}) + P_\mu \gamma_4(x_1, x_2) \sin \tilde{\vartheta} \cos \tilde{\varphi}] \\ &+ C_{11} P_\mu \gamma_3(x_1, x_2) \sin \tilde{\vartheta} \sin \tilde{\varphi} - C_{12} P_\mu \gamma_4(x_1, x_2) \sin \tilde{\vartheta} \sin \tilde{\varphi} \\ &\left. + \text{h.c.} \right]. \end{aligned} \quad (\text{B.3})$$

Herein, $x_i \equiv 2 \frac{E_i}{m_\mu}$ ($i = 1, 2$) are the normalized energies E_i of the two positrons with $E_1 \geq E_2$, P_μ is the net polarization of the muon, and the angles $\tilde{\vartheta}$ and $\tilde{\varphi}$ denote the angle between the electron and the muon spin direction and the angle between the electron-muon spin plane and the decay plane of $\mu \rightarrow eee$, respectively (see figure 4.16). The coefficients C_i are combinations of the coefficients $A_{R,L}$ and g_i of equation B.1

$$\begin{aligned} C_1 &= \frac{|g_1|^2}{16} + |g_3|^2, & C_2 &= \frac{|g_2|^2}{16} + |g_4|^2, \\ C_3 &= |g_5|^2, & C_4 &= |g_6|^2, \\ C_5 &= |eA_R|^2, & C_6 &= |eA_L|^2, \\ C_7 &= \text{Re}(eA_R g_4^*), & C_8 &= \text{Re}(eA_L g_3^*), \\ C_9 &= \text{Re}(eA_R g_6^*), & C_{10} &= \text{Re}(eA_L g_5^*), \\ C_{11} &= \text{Im}(eA_R g_4^* + eA_L g_3^*), & C_{12} &= \text{Im}(eA_R g_6^* + eA_L g_5^*). \end{aligned} \quad (\text{B.4})$$

The kinematical functions $\alpha_i(x_1, x_2)$, $\beta_i(x_1, x_2)$ and $\gamma_i(x_1, x_2)$ are given as follows

$$\begin{aligned}
 \alpha_1(x_1, x_2) &= 8(2 - x_1 - x_2)(x_1 + x_2 - 1) , \\
 \alpha_2(x_1, x_2) &= 2[x_1(1 - x_1) + x_2(1 - x_2)] , \\
 \alpha_3(x_1, x_2) &= 8 \left[\frac{2x_2^2 - 2x_2 + 1}{1 - x_1} + \frac{2x_1^2 - 2x_1 + 1}{1 - x_2} \right] , \\
 \alpha_4(x_1, x_2) &= 32(x_1 + x_2 - 1) , \\
 \alpha_5(x_1, x_2) &= 8(2 - x_1 - x_2) , \\
 \beta_1(x_1, x_2) &= 2 \frac{(x_1 + x_2)(x_1^2 + x_2^2) - 3(x_1 + x_2)^2 + 6(x_1 + x_2) - 4}{(2 - x_1 - x_2)} , \\
 \beta_2(x_1, x_2) &= \frac{8}{(1 - x_1)(1 - x_2)(2 - x_1 - x_2)} \times \\
 &\quad [2(x_1 + x_2)(x_1^3 + x_2^3) - 4(x_1 + x_2)(2x_1^2 + x_1x_2 + 2x_2^2) \\
 &\quad + (19x_1^2 + 30x_1x_2 + 19x_2^2) - 12(2x_1 + 2x_2 - 1)] , \\
 \gamma_1(x_1, x_2) &= 4 \frac{\sqrt{(1 - x_1)(1 - x_2)(x_1 + x_2 - 1)}(x_1 - x_2)}{(2 - x_2 - x_1)} , \\
 \gamma_2(x_1, x_2) &= 32 \sqrt{\frac{(x_1 + x_2 - 1)}{(1 - x_1)(1 - x_2)} \frac{(x_1 + x_2 - 1)(x_2 - x_1)}{(2 - x_1 - x_2)}} , \\
 \gamma_3(x_1, x_2) &= 16 \sqrt{\frac{(x_1 + x_2 - 1)}{(1 - x_1)(1 - x_2)}} (x_1 + x_2 - 1)(x_2 - x_1) , \\
 \gamma_4(x_1, x_2) &= 8 \sqrt{\frac{(x_1 + x_2 - 1)}{(1 - x_1)(1 - x_2)}} (2 - x_1 - x_2)(x_2 - x_1) . \quad (\text{B.5})
 \end{aligned}$$

Integrating the differential branching fraction from equation B.3 yields the total branching fraction for $\mu^+ \rightarrow e^+e^-e^+$ decays [46, 62]

$$\begin{aligned}
 \mathcal{B}_{\mu^+ \rightarrow e^+e^-e^+} &= \int_{\frac{1}{2}}^1 dx_1 \int_{1-x_1}^{x_1} dx_2 \int_{-1}^1 d(\cos \tilde{\vartheta}) \int_0^\pi d\tilde{\varphi} \frac{d\mathcal{B}_{\mu^+ \rightarrow e^+e^-e^+}}{dx_1 dx_2 d(\cos \tilde{\vartheta}) d\tilde{\varphi}} \\
 &= 2(C_1 + C_2) + (C_3 + C_4) + 32 \left[\ln \left(\frac{m_\mu^2}{m_e^2} \right) - \frac{11}{4} \right] (C_5 + C_6) \\
 &\quad + 16(C_7 + C_8) + 8(C_9 + C_{10}) . \quad (\text{B.6})
 \end{aligned}$$

Equation B.3 is used in chapter 4.4 to assign a weight to a $\mu \rightarrow eee$ event corresponding to interactions mediated by an effective operator. With equation B.6, sensitivities to the various effective operators are derived.




Figure B.1: Schematic of a Dalitz plot. The figure is taken from [1].

B.2 DALITZ ANALYSIS

The phase space of a three-body decay is often displayed as a Dalitz plot. It was introduced by Dalitz for studies of kaon decays into three pions [221].

The boundaries of the phase space of a three body decay are limited by the masses M and m_i of the mother and daughter particles, respectively ($i \in \{1, 2, 3\}$), but the kinematics are not fully constrained. Within the decay plane in the centre of mass system of the mother particle, a three body decay has two degrees of freedom, for instance the opening angles or invariant masses m_{ij} of two pairs of the decay products. Therefore, three-body decays in high energy physics are typically displayed as scatter plots of m_{ij}^2 . An exemplary Dalitz plot is shown in figure B.1.


From a Dalitz plot, information about the nature of the interaction can be drawn. The differential decay rate is given by [1]

$$d\Gamma = \frac{1}{(2\pi)^3} \frac{1}{32M^3} \overline{|\mathcal{M}|^2} dm_{12}^2 dm_{23}^2, \quad (\text{B.7})$$


with the spin-averaged matrix element $\overline{|\mathcal{M}|^2}$. For a constant $\overline{|\mathcal{M}|^2}$, the events are evenly distributed over the Dalitz plot. A resonance in one of the pairs of daughter particles for example, would manifest as a band.

B.3 DALITZ PLOTS


In the following, Dalitz plots and angular distributions of $\mu \rightarrow eee$ signal events mediated via effective operators are shown. The operators have opposite chiralities as the examples given in section 4.4. The distributions for the dipole operator \mathcal{O}_D^L are shown in figure B.2, for the scalar operator \mathcal{O}_S^{LL} and the vector operator \mathcal{O}_V^{LL} in figure B.3, and for the vector operator \mathcal{O}_V^{LR} in figure B.4.




(a) Dalitz plots of generated $\mu \rightarrow eee$ signal events for operator \mathcal{O}_D^L .




(b) Angular distribution of generated $\mu \rightarrow eee$ signal events for operator \mathcal{O}_D^L .




(c) Dalitz plots of reconstructed $\mu \rightarrow eee$ signal events for operator \mathcal{O}_D^L (truth variables).



(d) Angular distribution of reconstructed $\mu \rightarrow eee$ signal events for operator \mathcal{O}_D^L (truth variables).




(e) Dalitz plots of reconstructed $\mu \rightarrow eee$ signal events for operator \mathcal{O}_D^L (reconstructed variables).




(f) Angular distribution of reconstructed $\mu \rightarrow eee$ signal events for operator \mathcal{O}_D^L (reconstructed variables).


Figure B.2: Dalitz plots of the truth and reconstructed invariant mass of the e^+e^- system and the truth and reconstructed angular distribution of $\mu \rightarrow eee$ signal events mediated via the dipole operator \mathcal{O}_D^L on generator level and after vertex reconstruction. Note the logarithmic scale.




(a) Dalitz plots of generated $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_S^{\text{LL}}$ or $\mathcal{O}_V^{\text{LL}}$.




(b) Angular distribution of generated $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_S^{\text{LL}}$ or $\mathcal{O}_V^{\text{LL}}$.




(c) Dalitz plots of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_S^{\text{LL}}$ or $\mathcal{O}_V^{\text{LL}}$ (truth variables).



(d) Angular distribution of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_S^{\text{LL}}$ or $\mathcal{O}_V^{\text{LL}}$ (truth variables).




(e) Dalitz plots of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_S^{\text{LL}}$ or $\mathcal{O}_V^{\text{LL}}$ (reconstructed variables).




(f) Angular distribution of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_S^{\text{LL}}$ or $\mathcal{O}_V^{\text{LL}}$ (reconstructed variables).


Figure B.3: Dalitz plots of the truth and reconstructed invariant mass of the e^+e^- system and the truth and reconstructed angular distribution of $\mu \rightarrow eee$ signal events mediated via the scalar 4-fermion operator $\mathcal{O}_S^{\text{LL}}$ or the vector 4-fermion operator $\mathcal{O}_V^{\text{LL}}$ on generator level and after vertex reconstruction.




(a) Dalitz plots of generated $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{LR}}$.




(b) Angular distribution of generated $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{LR}}$.




(c) Dalitz plots of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{LR}}$ (truth variables).



(d) Angular distribution of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{LR}}$ (truth variables).



(e) Dalitz plots of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{LR}}$ (reconstructed variables).



(f) Angular distribution of reconstructed $\mu \rightarrow eee$ signal events for operator $\mathcal{O}_V^{\text{LR}}$ (reconstructed variables).

Figure B.4: Dalitz plots of the truth and reconstructed invariant mass of the e^+e^- system and the truth and reconstructed angular distribution of $\mu \rightarrow eee$ signal events mediated via the vector 4-fermion operator $\mathcal{O}_V^{\text{LR}}$ on generator level and after vertex reconstruction.

LEPTON FLAVOUR VIOLATING TWO-BODY DECAYS OF THE MUON

This chapter contains additional plots and auxiliary explanations in the context of sensitivity studies of lepton flavour violating two-body decays of the muon that are left out in chapter 5.

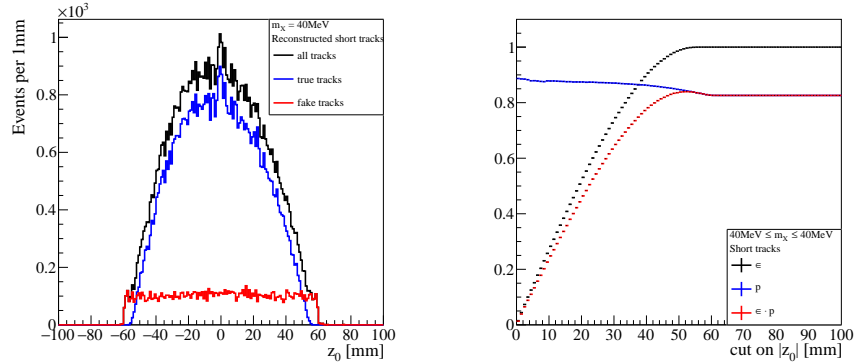
C.1 DATA SELECTION

In section 5.3.3.2, the effect of various cuts are discussed for $\mu \rightarrow eX$ decays with an exemplary mass m_X of 60 MeV. In this section, results for two additional masses are shown: For $m_X = 40$ MeV, the peak is located around the maximum of the momentum spectrum of the positrons and close to the Michel edge. For $m_X = 90$ MeV, the peak lies at the lower edge close to the momentum threshold defined by the geometric acceptance.

C.1.1 Cuts on the Decay Vertex Estimate z_0


Figures C.1 and C.2 show the effect of cuts on the decay vertex estimate z_0 for $\mu \rightarrow eX$ decays with masses m_X of 40 MeV and 90 MeV, respectively. The same distributions for m_X ranging from 0 MeV and 99 MeV are shown in figure C.3. A similar behaviour is observed as in section 5.3.3.2.

The distribution of z_0 is peaking at $z_0 = 0$ which is not expected from the true z position of the decay vertex z_{MC} . In figure C.4, the reconstructed z_0 is shown in comparison to z_{MC} . The shape of the z_{MC} distribution reflects the stopping distribution of the muons. The excess at $z_0 = 0$ arises from the reconstruction. Events with a true z_{MC} close to zero exhibit a long tail in z_0 towards both positive and negative z_0 . This is not observed for large $|z_{MC}|$. At $z_0 \approx 0$, the target and thus true decay vertices are close to the cylinder that defines the target region. The positron can be emitted in such a way that it crosses the target volume and is detected in the pixel layers on the opposite side. In this way, also z_0 is evaluated on the opposite side of the cylinder resulting in the observed tails. This poses however no problem as the estimated z_0 is still well within the cut range. The ef-



(a) Distribution in z_0 for reconstructed short tracks.


(b) Efficiency and purity for cuts on z_0 using reconstructed short tracks.




(c) Distribution in z_0 for reconstructed long tracks.

(d) Efficiency and purity for cuts on z_0 using reconstructed long tracks.


Figure C.1: Selection cuts on z_0 for $\mu \rightarrow eX$ decays with a mass m_X of 40 MeV reconstructed as short and long tracks. On the left hand side, the z_0 distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of z_0 . The sample contains about 75 000 $\mu \rightarrow eX$ decays.




(a) Distribution in z_0 for reconstructed short tracks.



(b) Efficiency and purity for cuts on z_0 using reconstructed short tracks.




(c) Distribution in z_0 for reconstructed long tracks.




(d) Efficiency and purity for cuts on z_0 using reconstructed long tracks.


Figure C.2: Selection cuts on z_0 for $\mu \rightarrow eX$ decays with a mass m_X of 90 MeV reconstructed as short and long tracks. On the left hand side, the z_0 distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of z_0 . The sample contains about 75 000 $\mu \rightarrow eX$ decays.




(a) Distribution in z_0 for reconstructed short tracks.



(b) Efficiency and purity for cuts on z_0 using reconstructed short tracks.



(c) Distribution in z_0 for reconstructed long tracks.



(d) Efficiency and purity for cuts on z_0 using reconstructed long tracks.

Figure C.3: Selection cuts on z_0 for $\mu \rightarrow eX$ decays with a mass m_X ranging from 0 MeV to 99 MeV reconstructed as short and long tracks. On the left hand side, the z_0 distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of z_0 . The sample contains about 75 000 $\mu \rightarrow eX$ decays per m_X .

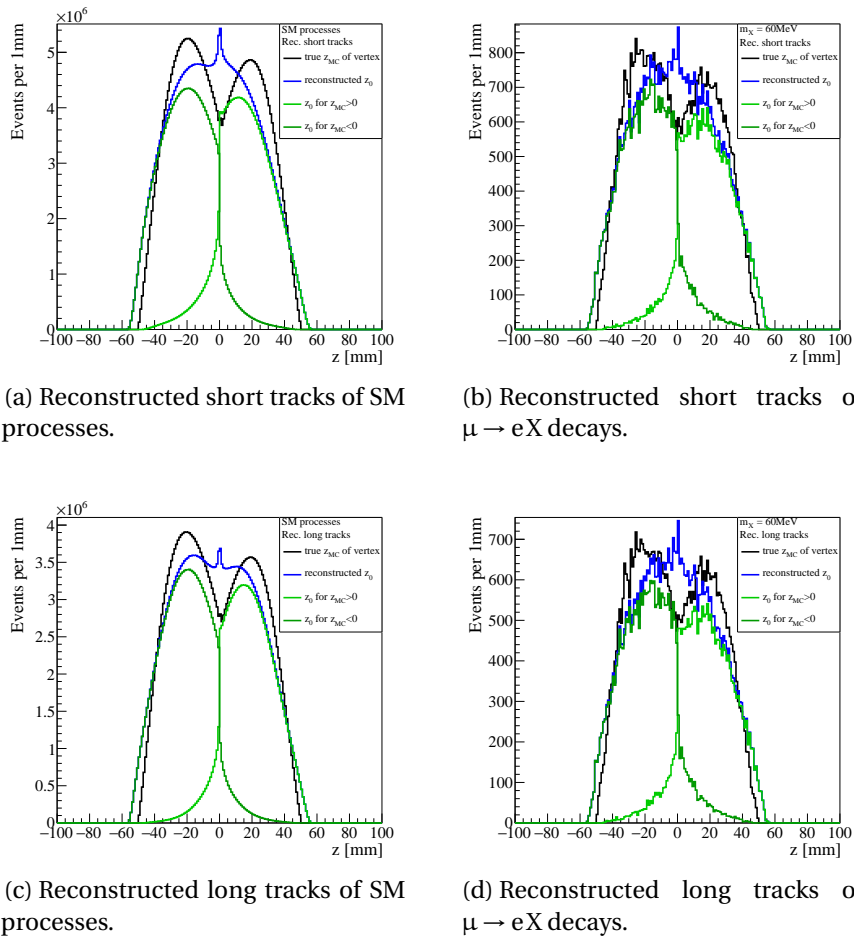


Figure C.4: True vertex position and reconstructed z_0 of single tracks for SM processes and $\mu \rightarrow eX$ decays with $m_X = 60$ MeV reconstructed as short and long tracks.

fect becomes negligible when the decay vertices are located around the tips of the target.

C.1.2 Cuts on the χ^2 of the Track Fit

The distribution of the χ^2 of the track fit is shown in figures C.5, C.6 and C.7 for masses of 40 MeV and 90 MeV, respectively, as well as for masses ranging from 0 MeV to 99 MeV. The purity of reconstructed recurling tracks is smaller in the case of 90 MeV than for smaller masses. For $m_X = 90$ MeV, the momentum of the decay positron is small. Only particles that have mostly transverse momentum are reconstructed which are prone to perform many loops in the central detector. Thus, the amount of fake tracks is enhanced.

Apart from this, the behaviour of the curves does not deviate from

the previous observations.

C.1.3 Cuts on the Dip Angle λ_{01}

The effect of cuts on λ_{01} on reconstructed $\mu \rightarrow eX$ events with masses of 40 MeV and 90 MeV, as well as masses in a range from 0 MeV to 99 MeV are depicted in figures C.8, C.9 and C.10. No significant deviations with respect to SM processes and $\mu \rightarrow eX$ decays with $m_X = 60$ MeV are observed. As the decay positrons for $m_X = 90$ MeV have a small momentum, only events with mostly transverse momentum are reconstructed and thus a smaller range of λ_{01} is covered.

C.1.4 Combined Cuts

In figure C.11, histograms of the reconstructed momentum of $\mu \rightarrow eX$ decays with masses of 40 MeV and 90 MeV, and for masses between 0 MeV to 99 MeV in steps of 1 MeV are shown after applying the combined cuts. In the latter, the various peaks can be distinguished especially for low m_X .

The efficiencies and purities for selected masses are summarized in table C.1.


C.2 EXPECTED BRANCHING FRACTION LIMITS

C.2.1 One-Dimensional Analysis


C.2.1.1 Mean and Width of the Signal Fraction

In the derivation of $\mu \rightarrow eX$ branching fraction limits as outlined in section 5.2.1, the mean μ_{toyMC} of the fit of a Gaussian distribution to the signal fraction is fixed to zero. As demonstrated in figure C.12, this is a valid assumption. On the left-hand side, the fraction of the mean μ_{toyMC} and width σ_{toyMC} of a Gaussian fit with μ_{toyMC} as a free fit parameter is shown for different numbers of toy Monte Carlo events N_{toyMC} . The deviation from zero is very small with less than 2% of σ_{toyMC} , and it is driven by the statistical uncertainty due to the limited number of events.


On the right-hand side of figure C.12, the evolution of $\ln(\sigma_{\text{toyMC}})$




(a) Distribution in χ^2 for reconstructed short tracks.



(b) Efficiency and purity for cuts on χ^2 using reconstructed short tracks.




(c) Distribution in χ^2 for reconstructed long tracks.




(d) Efficiency and purity for cuts on χ^2 using reconstructed long tracks.


Figure C.5: Selection cuts on the χ^2 of the track fit for $\mu \rightarrow eX$ decays with a mass m_X of 40 MeV reconstructed as short and long tracks. On the left hand side, the χ^2 distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of χ^2 . The sample contains about 75 000 $\mu \rightarrow eX$ decays.




(a) Distribution in χ^2 for reconstructed short tracks.



(b) Efficiency and purity for cuts on χ^2 using reconstructed short tracks.




(c) Distribution in χ^2 for reconstructed long tracks.




(d) Efficiency and purity for cuts on χ^2 using reconstructed long tracks.


Figure C.6: Selection cuts on the χ^2 of the track fit for $\mu \rightarrow eX$ decays with a mass m_X of 90 MeV reconstructed as short and long tracks. On the left hand side, the χ^2 distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of χ^2 . The sample contains about 75 000 $\mu \rightarrow eX$ decays.




(a) Distribution in χ^2 for reconstructed short tracks.



(b) Efficiency and purity for cuts on χ^2 using reconstructed short tracks.




(c) Distribution in χ^2 for reconstructed long tracks.




(d) Efficiency and purity for cuts on χ^2 using reconstructed long tracks.


Figure C.7: Selection cuts on the χ^2 of the track fit for $\mu \rightarrow eX$ decays with a mass m_X ranging from 0 MeV to 99 MeV reconstructed as short and long tracks. On the left hand side, the χ^2 distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of χ^2 . The sample contains about 75 000 $\mu \rightarrow eX$ decays per m_X .




(a) Distribution in λ_{01} for reconstructed short tracks.



(b) Efficiency and purity for cuts on λ_{01} using reconstructed short tracks.




(c) Distribution in λ_{01} for reconstructed long tracks.




(d) Efficiency and purity for cuts on λ_{01} using reconstructed long tracks.


Figure C.8: Selection cuts on λ_{01} for $\mu \rightarrow eX$ decays with a mass m_X of 40 MeV reconstructed as short and long tracks. On the left hand side, the λ_{01} distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of λ_{01} . The sample contains about 75 000 $\mu \rightarrow eX$ decays.




(a) Distribution in λ_{01} for reconstructed short tracks.



(b) Efficiency and purity for cuts on λ_{01} using reconstructed short tracks.




(c) Distribution in λ_{01} for reconstructed long tracks.




(d) Efficiency and purity for cuts on λ_{01} using reconstructed long tracks.


Figure C.9: Selection cuts on λ_{01} for $\mu \rightarrow eX$ decays with a mass m_X of 90 MeV reconstructed as short and long tracks. On the left hand side, the λ_{01} distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of λ_{01} . The sample contains about 75 000 $\mu \rightarrow eX$ decays.




(a) Distribution in λ_{01} for reconstructed short tracks.



(b) Efficiency and purity for cuts on λ_{01} using reconstructed short tracks.




(c) Distribution in λ_{01} for reconstructed long tracks.



(d) Efficiency and purity for cuts on λ_{01} using reconstructed long tracks.


Figure C.10: Selection cuts on λ_{01} for $\mu \rightarrow eX$ decays with a mass m_X ranging from 0 MeV to 99 MeV reconstructed short and long tracks. On the left hand side, the λ_{01} distribution is shown, and on the right hand side the yielded efficiency ϵ , purity p , and the product $\epsilon \cdot p$ for cuts of λ_{01} . The sample contains about 75 000 $\mu \rightarrow eX$ decays per m_X .



(a) $\mu \rightarrow eX$ events with $m_X = 40\text{MeV}$ reconstructed as short tracks.

(b) $\mu \rightarrow eX$ events with $m_X = 90\text{MeV}$ reconstructed as short tracks.

(c) $\mu \rightarrow eX$ events with m_X ranging from 0 MeV to 99 MeV reconstructed as short tracks.



(d) $\mu \rightarrow eX$ events with $m_X = 40\text{MeV}$ reconstructed as long tracks.

(e) $\mu \rightarrow eX$ events with $m_X = 90\text{MeV}$ reconstructed as long tracks.

(f) $\mu \rightarrow eX$ events with m_X ranging from 0 MeV to 99 MeV reconstructed as long tracks.

Figure C.11: Distribution of the reconstructed momentum of positron candidates in $\mu \rightarrow eX$ decays with a mass m_X of 40 MeV and 90 MeV, and m_X ranging from 0 MeV to 99 MeV before and after applying cuts as outlined in table 5.2. Reconstructed true tracks are shown in blue, fake tracks in red and the total reconstructed tracks in black. A dashed line indicates the distribution before applying cuts, and a solid line after applying cuts. The sample contains about 75 000 $\mu \rightarrow eX$ decays.

m_X	Reconstruction		No cuts	Cuts applied
40 MeV	Short	ϵ	1	0.849
		p	0.826	0.982
	Long	ϵ	1	0.812
		p	0.927	0.983
60 MeV	Short	ϵ	1	0.847
		p	0.794	0.968
	Long	ϵ	1	0.845
		p	0.926	0.978
90 MeV	Short	ϵ	1	0.768
		p	0.647	0.788
	Long	ϵ	1	0.766
		p	0.764	0.828

Table C.1: Efficiency ϵ and purity p before and after applying cuts for $\mu \rightarrow eX$ decays with $m_X = 40$ MeV, $m_X = 60$ MeV and $m_X = 90$ MeV. The statistical uncertainty is on the 10^{-3} level.

with the number of toy Monte Carlo events is shown. The mean is fixed to zero in this case. A linear fit


$$\ln(\sigma_{\text{toyMC}}) \propto m \ln(N_{\text{toyMC}}) \quad (\text{C.1})$$

with slope m is applied. For all three reconstruction modes, the fit yields a slope compatible to -0.5 . It is thus demonstrated, that σ_{toyMC} scales with the inverse of the square root of N_{toyMC}


$$\sigma_{\text{toyMC}} \propto \frac{1}{\sqrt{N_{\text{toyMC}}}}. \quad (\text{C.2})$$

C.2.1.2 Toy Monte Carlo Studies for Long Track Reconstruction


In figure C.13, an exemplary toy Monte Carlo study for reconstructed long tracks is shown. As recurving tracks have a smaller momentum resolution than short tracks, the distribution of the reconstructed momentum for $\mu \rightarrow eX$ events is narrower and the Michel edge of the SM processes steeper (see figure 5.3). This results in smaller fitted signal fractions.




(a) Fraction of μ_{toyMC} and σ_{toyMC} using reconstructed short tracks.



(b) $\ln(\sigma_{\text{toyMC}})$ using reconstructed short tracks.




(c) Fraction of μ_{toyMC} and σ_{toyMC} using reconstructed long tracks.




(d) $\ln(\sigma_{\text{toyMC}})$ using reconstructed long tracks.


Figure C.12: Scaling of the mean μ_{toyMC} and width σ_{toyMC} in Gaussian fits to the signal fraction for different numbers of toy Monte Carlo events. A linear fit with slope m is applied to σ_{toyMC} . The study is performed with $\mu \rightarrow eX$ events with a mass m_X of 60 MeV.




(a) Histogram of the reconstructed momentum from SM background processes. About $5 \cdot 10^8$ decays in the target region are generated.



(b) Histogram of the reconstructed momentum from $\mu \rightarrow eX$ signal events with $m_X = 60\text{MeV}$. About 75 000 decays in the target region are generated.




(c) A single sample of events derived from f_b . The solid blue line is a fit of f_{s+b} to the events. The fitted amount of signal (solid red line) and background (dashed blue line) are also shown. Please note, that the solid and dashed blue line are indistinguishable on the logarithmic scale.




(d) Histogram of the signal fraction f_{sig} for a complete study with fit of a Gaussian distribution (solid blue line).


Figure C.13: Exemplary toy Monte Carlo study for the long track reconstruction and $\mu \rightarrow eX$ events with $m_X = 60\text{MeV}$.




(a) Fit of the background p.d.f. to a single sample. A signal region corresponding to $\mu \rightarrow eX$ events with $m_X = 10 \text{ MeV}$ is left out.



(b) Distribution of the fitted scaling parameter x_{scale} for $\mu \rightarrow eX$ searches with $m_X = 10 \text{ MeV}$.



(c) Fit of the background p.d.f. to a single sample. A signal region corresponding to $\mu \rightarrow eX$ events with $m_X = 60 \text{ MeV}$ is left out.



(d) Distribution of the fitted scaling parameter x_{scale} for $\mu \rightarrow eX$ searches with $m_X = 60 \text{ MeV}$.

Figure C.14: Simultaneous determination of the momentum calibration for reconstructed long tracks. On the left-hand side, a fit of the background p.d.f. (green line) is shown to a single sample (black dots) of 10^9 events derived from the background p.d.f. (blue line). On the right-hand side, the distribution of the fitted scaling parameter is shown for 1000 fits.

C.3 MOMENTUM CALIBRATION

Figure C.14 shows the determination of the calibration of the momentum spectrum in the case of long track reconstruction. As observed for reconstructed short tracks in section 5.6, the distribution of the fit results for the scaling parameter broadens for smaller masses m_X , as in these cases the characteristic Michel edge is left out for the fit of the background p.d.f. In all cases, the average fitted scaling parameter is compatible with unity meaning the calibration fit is capable to reproduce the initial p.d.f. shape.

PUBLICATIONS

Some of the ideas presented in this thesis have appeared previously or will be published in the following journal articles and conference proceedings.

Technical design of the Phase I Mu3e Experiment

A. Blondel et al.
to be published

Searching for Lepton Flavour Violation with the Mu3e Experiment

A.K. Perrevoort (for the Mu3e collaboration)
arXiv: 1803.01581[physics.ins-det]
to be published in *PoS NuFact2017* (2018), 105.

Efficiency and Timing Performance of the MuPix7 High Voltage Monolithic Active Pixel Sensor

H. Augustin et al.
arXiv: 1802.09851[physics.ins-det]
to be published in *Nucl. Instrum. Meth. A*

Irradiation Study of a Fully Monolithic HV-CMOS Pixel Sensor Design in AMS 180 nm

H. Augustin et al.
(2017). arXiv: 1712.03921[physics.ins-det]

The MuPix Telescope: A Thin, High Rate Tracking Telescope

H. Augustin et al.
In: *JINST* 12.01 (2017), C01087. DOI: 10.1088/1748-0221/12/01/C01087. arXiv: 1611.03102[physics.ins-det]

MuPix7—A Fast Monolithic HV-CMOS Pixel Chip for Mu3e

H. Augustin et al.
In: *JINST* 11.11 (2017), C11029. DOI: 10.1088/1748-0221/11/11/C11029. arXiv: 1610.02210[physics.ins-det]

Status of the Mu3e Experiment at PSI

A.K. Perrevoort (for the Mu3e collaboration)
In: *EPJ Web Conf.* 118 (2016), 01028. DOI: 10.1051/epjconf/201611801028. arXiv: 1605.02906[physics.ins-det]

The MuPix System-on-Chip for the Mu3e Experiment

H. Augustin et al.

In: *Nucl. Instrum. Meth.* A845 (2017), 194-198. DOI: 10.1016/j.nima.2016.06.095. arXiv: 1603.08751[physics.ins-det]**Overview of HVCMOS Pixel Sensors**

I. Perić et al.

In: *JINST* 10.05 (2015), C05021. DOI: 10.1088/1748-0221/10/05/C05021.

The content of the following publications is related to previous research conducted during my Bachelor and Master thesis as well as during a research stay prior to the doctoral studies.

CASCADE - A Multi-Layer Boron-10 Neutron Detection System

M. Köhli, M. Klein, F. Allmendinger, A.K. Perrevoort, T. Schröder, N. Martin, C.J. Schmidt, U. Schmidt

In: *J. Phys. Conf. Ser.* 746.1 (2016), 012003. DOI: 10.1088/1742-6596/746/1/012003. arXiv: 1602.04064[physics.ins-det]**Clock Monitoring at the GET Electronics**

A.K. Perrevoort, T. Isobe, H. Baba, A. Taketani

In: *RIKEN Accelerator Progress Report* 47 (2014), 182.**Missing Mass Spectroscopy on Carbon Isotopes beyond Proton Drip-Line**

T. Miyazaki et al.

In: *RIKEN Accelerator Progress Report* 47 (2014), 23.**High-Voltage Pixel Detectors in Commercial CMOS Technologies for ATLAS, CLIC and Mu3e Experiments**

I. Perić et al.

In: *Nucl. Instrum. Meth.* A731 (2013), 131-136. DOI: 10.1016/j.nima.2013.05.006.**A Tracker for the Mu3e Experiment based on High-Voltage Monolithic Active Pixel Sensors**

N. Berger et al.

In: *Nucl. Instrum. Meth.* A732 (2013), 61-65. DOI: 10.1016/j.nima.2013.05.035. arXiv: 1309.7896[physics.ins-det]**A Novel Experiment Searching for the Lepton Flavor Violating Decay $\mu \rightarrow eee$**

D. Wiedner, H. Augustin, S. Bachmann, N. Berger, M. Kiehn, I. Perić, A.K. Perrevoort, A. Schöning

In: *PoS Vertex2012* (2013), 033.

BIBLIOGRAPHY

-
- [1] C. Patrignani et al. “Review of Particle Physics”. In: *Chin. Phys. C* 40.10 (2016), p. 100001. DOI: 10.1088/1674-1137/40/10/100001.
- [2] R. Aaij et al. “Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\Lambda_b^0 \rightarrow J/\psi K^- p$ Decays”. In: *Phys. Rev. Lett.* 115 (2015), p. 072001. DOI: 10.1103/PhysRevLett.115.072001. arXiv: 1507.03414 [hep-ex].
- [3] C. Burgard and D. Galbraith. *Standard Model of Physics*. CERN Webfest. 2012. URL: <http://davidgalbraith.org/portfolio/ux-standard-model-of-the-standard-model/>.
- [4] F. Halzen and A. D. Martin. *Quarks and Leptons. an introductory course in modern particle physics*. eng. [Chichester et al.]: Wiley, 1984, XVI, 396 S.
- [5] A. Hillairet et al. “Precision muon decay measurements and improved constraints on the weak interaction”. In: *Phys. Rev. D* 85 (2012), p. 092013. DOI: 10.1103/PhysRevD.85.092013. arXiv: 1112.3606 [hep-ex].
- [6] 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”. In: (2012). arXiv: 1207.7214 [hep-ex].
- [7] S. Chatrchyan et al. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Phys.Lett.B* (2012). arXiv: 1207.7235 [hep-ex].
- [8] F. Zwicky. “Die Rotverschiebung von extragalaktischen Nebeln”. In: *Helv. Phys. Acta* 6 (1933). [Gen. Rel. Grav.41,207(2009)], pp. 110–127. DOI: 10.1007/s10714-008-0707-4.
- [9] V. C. Rubin, N. Thonnard, and W. K. Ford Jr. “Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/”. In: *Astrophys. J.* 238 (1980), p. 471. DOI: 10.1086/158003.

- [10] X.-P. Wu, T. Chiueh, L.-Z. Fang, and Y.-J. Xue. “A comparison of different cluster mass estimates: consistency or discrepancy?” In: *Mon. Not. Roy. Astron. Soc.* 301 (1998), p. 861. DOI: 10.1046/j.1365-8711.1998.02055.x. arXiv: astro-ph/9808179 [astro-ph].
- [11] E. Corbelli and P. Salucci. “The Extended Rotation Curve and the Dark Matter Halo of M33”. In: *Mon. Not. Roy. Astron. Soc.* 311 (2000), pp. 441–447. DOI: 10.1046/j.1365-8711.2000.03075.x. arXiv: astro-ph/9909252 [astro-ph].
- [12] A. Refregier. “Weak gravitational lensing by large scale structure”. In: *Ann. Rev. Astron. Astrophys.* 41 (2003), pp. 645–668. DOI: 10.1146/annurev.astro.41.111302.102207. arXiv: astro-ph/0307212 [astro-ph].
- [13] G. Hinshaw et al. “Five-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Data Processing, Sky Maps, and Basic Results”. In: *Astrophys. J. Suppl.* 180 (2009), pp. 225–245. DOI: 10.1088/0067-0049/180/2/225. arXiv: 0803.0732 [astro-ph].
- [14] P. A. R. Ade et al. “Planck 2015 results. XIII. Cosmological parameters”. In: *Astron. Astrophys.* 594 (2016), A13. DOI: 10.1051/0004-6361/201525830. arXiv: 1502.01589 [astro-ph.CO].
- [15] P. Parihar, M. S. Vogeley, J. R. Gott, Y.-Y. Choi, J. Kim, S.-S. Kim, R. Speare, J. R. Brownstein, and J. Brinkmann. “A Topological Analysis for Large-Scale Structure, Studied Using the CMASS Sample of SDSS-III”. In: *Astrophys. J.* 796.2 (2014), p. 86. DOI: 10.1088/0004-637X/796/2/86.
- [16] A. G. Riess et al. “Observational evidence from supernovae for an accelerating universe and a cosmological constant”. In: *Astron. J.* 116 (1998), pp. 1009–1038. DOI: 10.1086/300499. arXiv: astro-ph/9805201 [astro-ph].
- [17] R. Adam et al. “Planck 2015 results. I. Overview of products and scientific results”. In: *Astron. Astrophys.* 594 (2016), A1. DOI: 10.1051/0004-6361/201527101. arXiv: 1502.01582 [astro-ph.CO].
- [18] A. D. Sakharov. “Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe”. In: *Pisma Zh. Eksp. Teor. Fiz.* 5 (1967). [Usp. Fiz. Nauk161,no.5,61(1991)], pp. 32–35. DOI: 10.1070/PU1991v034n05ABEH002497.

- [19] R. Davis Jr., D. S. Harmer, and K. C. Hoffman. “Search for neutrinos from the sun”. In: *Phys. Rev. Lett.* 20 (1968), pp. 1205–1209. DOI: 10.1103/PhysRevLett.20.1205.
- [20] Q. R. Ahmad et al. “Measurement of the charged current interactions produced by B-8 solar neutrinos at the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 87 (2001), p. 071301. eprint: nucl-ex/0106015.
- [21] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. eprint: hep-ex/9807003.
- [22] F.P. An et al. “Observation of electron-antineutrino disappearance at Daya Bay”. In: *Phys.Rev.Lett.* 108 (2012), p. 171803. DOI: 10.1103/PhysRevLett.108.171803. arXiv: 1203.1669 [hep-ex].
- [23] J.K. Ahn et al. “Observation of Reactor Electron Antineutrino Disappearance in the RENO Experiment”. In: *Phys.Rev.Lett.* 108 (2012), p. 191802. DOI: 10.1103/PhysRevLett.108.191802. arXiv: 1204.0626 [hep-ex].
- [24] K. Abe et al. “Evidence of Electron Neutrino Appearance in a Muon Neutrino Beam”. In: *Phys. Rev. D* 88.3 (2013), p. 032002. DOI: 10.1103/PhysRevD.88.032002. arXiv: 1304.0841 [hep-ex].
- [25] P. J. Mohr, D. B. Newell, and B. N. Taylor. “CODATA Recommended Values of the Fundamental Physical Constants: 2014”. In: *Rev. Mod. Phys.* 88.3 (2016), p. 035009. DOI: 10.1103/RevModPhys.88.035009. arXiv: 1507.07956 [physics.atom-ph].
- [26] J. C. Bernauer et al. “High-precision determination of the electric and magnetic form factors of the proton”. In: *Phys. Rev. Lett.* 105 (2010), p. 242001. DOI: 10.1103/PhysRevLett.105.242001. arXiv: 1007.5076 [nucl-ex].
- [27] R. Pohl et al. “The size of the proton”. In: *Nature* 466 (2010), pp. 213–216. DOI: 10.1038/nature09250.
- [28] R. Pohl et al. “Laser spectroscopy of muonic deuterium”. In: *Science* 353.6300 (2016), pp. 669–673. DOI: 10.1126/science.aaf2468.
- [29] T. Blum, A. Denig, I. Logashenko, E. de Rafael, B. Lee Roberts, T. Teubner, and G. Venanzoni. “The Muon (g-2) Theory Value: Present and Future”. In: (2013). arXiv: 1311.2198 [hep-ph].

- [30] J. Grange et al. “Muon (g-2) Technical Design Report”. In: (2015). arXiv: 1501.06858.
- [31] J. L. Holzbauer. “The Muon g-2 Experiment Overview and Status”. In: *International Workshop on Neutrinos from Accelerators (NuFact17)*. 2017. arXiv: 1712.05980 [hep-ex]. URL: <http://inspirehep.net/record/1643822/files/arXiv:1712.05980.pdf>.
- [32] R. Aaij et al. “Test of lepton universality using $B^+ \rightarrow K^+ \ell^+ \ell^-$ decays”. In: *Phys. Rev. Lett.* 113 (2014), p. 151601. DOI: 10.1103/PhysRevLett.113.151601. arXiv: 1406.6482 [hep-ex].
- [33] R. Aaij et al. “Measurement of the ratio of branching fractions $\mathcal{B}(\bar{B}^0 \rightarrow D^{*+} \tau^- \bar{\nu}_\tau) / \mathcal{B}(\bar{B}^0 \rightarrow D^{*+} \mu^- \bar{\nu}_\mu)$ ”. In: *Phys. Rev. Lett.* 115.11 (2015). [Erratum: *Phys. Rev. Lett.* 115, no. 15, 159901 (2015)], p. 111803. DOI: 10.1103/PhysRevLett.115.159901, 10.1103/PhysRevLett.115.111803. arXiv: 1506.08614 [hep-ex].
- [34] R. Aaij et al. “Angular analysis of the $B^0 \rightarrow K^{*0} \mu^+ \mu^-$ decay using 3 fb^{-1} of integrated luminosity”. In: *JHEP* 02 (2016), p. 104. DOI: 10.1007/JHEP02(2016)104. arXiv: 1512.04442 [hep-ex].
- [35] R. Aaij et al. “Test of lepton universality with $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ decays”. In: *JHEP* 08 (2017), p. 055. DOI: 10.1007/JHEP08(2017)055. arXiv: 1705.05802 [hep-ex].
- [36] A. de Gouvea and P. Vogel. “Lepton Flavor and Number Conservation, and Physics Beyond the Standard Model”. In: *Prog. Part. Nucl. Phys.* 71 (2013), pp. 75–92. DOI: 10.1016/j.pnpnp.2013.03.006. arXiv: 1303.4097 [hep-ph].
- [37] S. T. Petcov. “The Processes $\mu \rightarrow e \gamma$, $\mu \rightarrow e e \text{ anti-}e$, Neutrino \rightarrow Neutrino γ in the Weinberg-Salam Model with Neutrino Mixing”. In: *Sov. J. Nucl. Phys.* 25 (1977). [Erratum: *Yad. Fiz.* 25, 1336 (1977)], p. 340.
- [38] S. M. Bilenky, S. T. Petcov, and B. Pontecorvo. “Lepton Mixing, $\mu \rightarrow e + \gamma$ Decay and Neutrino Oscillations”. In: *Phys. Lett.* 67B (1977), p. 309. DOI: 10.1016/0370-2693(77)90379-3.
- [39] S. L. Glashow, J. Iliopoulos, and L. Maiani. “Weak Interactions with Lepton-Hadron Symmetry”. In: *Phys. Rev. D* 2 (1970), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285.

- [40] W. J. Marciano, T. Mori, and J. M. Roney. “Charged Lepton Flavor Violation Experiments”. In: *Ann.Rev.Nucl.Part.Sci.* 58 (2008), pp. 315–341. DOI: 10 . 1146 / annurev . nucl . 58 . 110707 . 171126.
- [41] C. Patrignani et al. “Review of Particle Physics”. In: *Chin. Phys.* C40.10 (2016), p. 100001. DOI: 10 . 1088 / 1674 - 1137 / 40 / 10 / 100001.
- [42] F. Borzumati and A. Masiero. “Large Muon and Electron Number Violations in Supergravity Theories”. In: *Phys. Rev. Lett.* 57 (1986), p. 961.
- [43] K. Huitu, J. Maalampi, M. Raidal, and A. Santamaria. “New constraints on R-parity violation from mu e conversion in nuclei”. In: *Phys. Lett.* B430 (1998), pp. 355–362. DOI: 10 . 1016 / S0370 - 2693 (98) 00344 - X. arXiv: hep - ph / 9712249 [hep-ph].
- [44] A. Faessler, T. S. Kosmas, S. Kovalenko, and J. D. Vergados. “Constraints on R-parity violating supersymmetry from muon electron nuclear conversion”. In: *Nucl. Phys.* B587 (2000), pp. 25–44. DOI: 10 . 1016 / S0550 - 3213 (00) 00446 - 6. arXiv: hep - ph / 9904335 [hep-ph].
- [45] F. R. Joaquim and A. Rossi. “Lepton flavour violation in the supersymmetric type-II seesaw mechanism”. In: *AIP Conf. Proc.* 1078 (2009), pp. 372–374. DOI: 10 . 1063 / 1 . 3051965. arXiv: 0809 . 2859 [hep-ph].
- [46] Y. Kuno and Y. Okada. “Muon decay and physics beyond the standard model”. In: *Rev. Mod. Phys.* 73 (2001), pp. 151–202. eprint: hep - ph / 9909265.
- [47] J. C. Pati and A. Salam. “Lepton Number as the Fourth Color”. In: *Phys. Rev.* D10 (1974), pp. 275–289.
- [48] R. N. Mohapatra and J. C. Pati. “Left-Right Gauge Symmetry and an Isoconjugate Model of CP Violation”. In: *Phys. Rev.* D11 (1975), pp. 566–571.
- [49] G. Senjanovic and R. N. Mohapatra. “Exact Left-Right Symmetry and Spontaneous Violation of Parity”. In: *Phys. Rev.* D12 (1975), p. 1502.
- [50] R. E. Marshak and R. N. Mohapatra. “Quark - Lepton Symmetry and B-L as the U(1) Generator of the Electroweak Symmetry Group”. In: *Phys. Lett.* 91B (1980), pp. 222–224. DOI: 10 . 1016 / 0370 - 2693 (80) 90436 - 0.

- [51] Aharon Davidson. “ $B - L$ as the fourth color within an $SU(2)_L \times U(1)_R \times U(1)$ model”. In: *Phys. Rev. D* 20 (3 1979), pp. 776–783. DOI: 10.1103/PhysRevD.20.776. URL: <https://link.aps.org/doi/10.1103/PhysRevD.20.776>.
- [52] U. Chattopadhyay and P. B. Pal. “Radiative Neutrino Decay in Left-right Models”. In: *Phys. Rev. D* 34 (1986), p. 3444. DOI: 10.1103/PhysRevD.34.3444.
- [53] H. M. Georgi, S. L. Glashow, and S. Nussinov. “Unconventional Model of Neutrino Masses”. In: *Nucl. Phys.* B193 (1981), pp. 297–316. DOI: 10.1016/0550-3213(81)90336-9.
- [54] R. N. Mohapatra and Palash B. Pal. *Massive neutrinos in physics and astrophysics*. 3. ed. World Scientific lecture notes in physics. Singapore [et al.]: World Scientific, 2004, XXII, 451 S.
- [55] E. J. Chun, K. Y. Lee, and S. C. Park. “Testing Higgs triplet model and neutrino mass patterns”. In: *Phys. Lett.* B566 (2003), pp. 142–151. DOI: 10.1016/S0370-2693(03)00770-6. arXiv: hep-ph/0304069 [hep-ph].
- [56] M. Kakizaki, Y. Ogura, and F. Shima. “Lepton flavor violation in the triplet Higgs model”. In: *Phys. Lett.* B566 (2003), pp. 210–216. DOI: 10.1016/S0370-2693(03)00833-5. arXiv: hep-ph/0304254 [hep-ph].
- [57] J. Bernabeu, E. Nardi, and D. Tommasini. “ $\mu - e$ conversion in nuclei and Z' physics”. In: *Nucl. Phys.* B409 (1993), pp. 69–86. eprint: hep-ph/9306251.
- [58] L. Randall and R. Sundrum. “A large mass hierarchy from a small extra dimension”. In: *Phys. Rev. Lett.* 83 (1999), p. 3370.
- [59] N. Arkani-Hamed and M. Schmaltz. “Hierarchies without symmetries from extra dimensions”. In: *Phys. Rev.* D61 (2000), p. 33005.
- [60] A. Crivellin, S. Davidson, G. M. Pruna, and A. Signer. “Complementarity in lepton-flavour violating muon decay experiments”. In: (2016). arXiv: 1611.03409 [hep-ph].
- [61] A. Crivellin, S. Davidson, G. M. Pruna, and A. Signer. “Renormalisation-group improved analysis of $\mu \rightarrow e$ processes in a systematic effective-field-theory approach”. In: *JHEP* 05 (2017), p. 117. DOI: 10.1007/JHEP05(2017)117. arXiv: 1702.03020 [hep-ph].

- [62] Y. Okada, K. Okumura, and Y. Shimizu. “ $\mu \rightarrow e\gamma$ and $\mu \rightarrow 3e$ processes with polarized muons and supersymmetric grand unified theories”. In: *Phys.Rev. D* 61 (2000), p. 094001. DOI: 10.1103/PhysRevD.61.094001. arXiv: hep-ph/9906446 [hep-ph].
- [63] A. M. Baldini et al. “Search for the lepton flavour violating decay $\mu^+ \rightarrow e^+\gamma$ with the full dataset of the MEG experiment”. In: *Eur. Phys. J. C* 76.8 (2016), p. 434. DOI: 10.1140/epjc/s10052-016-4271-x. arXiv: 1605.05081 [hep-ex].
- [64] U. Bellgardt et al. “Search for the Decay $\mu^+ \rightarrow e^+e^+e^-$ ”. In: *Nucl.Phys. B* 299 (1988), p. 1. DOI: 10.1016/0550-3213(88)90462-2.
- [65] W. H. Bertl et al. “A Search for muon to electron conversion in muonic gold”. In: *Eur.Phys.J. C* 47 (2006), pp. 337–346. DOI: 10.1140/epjc/s2006-02582-x.
- [66] B. Aubert et al. “Searches for Lepton Flavor Violation in the Decays $\tau^\pm \rightarrow e^\pm\gamma$ and $\tau^\pm \rightarrow \mu^\pm\gamma$ ”. In: *Phys.Rev.Lett.* 104 (2010), p. 021802. DOI: 10.1103/PhysRevLett.104.021802. arXiv: 0908.2381 [hep-ex].
- [67] Y. Yusa et al. “Search for neutrinoless decays $\tau \rightarrow 3l$ ”. In: *Phys. Lett. B* 589 (2004), pp. 103–110. DOI: 10.1016/j.physletb.2004.04.002. arXiv: hep-ex/0403039 [hep-ex].
- [68] CMS Collaboration. “Search for Lepton Flavor Violation in Z decays in pp collisions at $\sqrt{s}=8$ TeV”. In: (2015).
- [69] R. Akers et al. “A Search for lepton flavor violating Z^0 decays”. In: *Z.Phys. C* 67 (1995), pp. 555–564. DOI: 10.1007/BF01553981.
- [70] P. Abreu et al. “Search for lepton flavor number violating Z^0 decays”. In: *Z.Phys. C* 73 (1997), pp. 243–251. DOI: 10.1007/s002880050313.
- [71] V. Khachatryan et al. “Search for lepton flavour violating decays of the Higgs boson to $e\tau$ and $e\mu$ in proton–proton collisions at $\sqrt{s} = 8$ TeV”. In: *Phys. Lett. B* 763 (2016), pp. 472–500. DOI: 10.1016/j.physletb.2016.09.062. arXiv: 1607.03561 [hep-ex].
- [72] CMS Collaboration. “Search for lepton flavour violating decays of the Higgs boson to $\mu\tau$ and $e\tau$ in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: (2017).

- [73] A.M. Baldini, F. Cei, C. Cerri, S. Dussoni, L. Galli, et al. “MEG Upgrade Proposal”. In: *ArXiv e-prints* (Jan. 2013). arXiv: 1301.7225 [physics.ins-det].
- [74] A. Blondel et al. “Research Proposal for an Experiment to Search for the Decay $\mu \rightarrow eee$ ”. In: (2013). arXiv: 1301.6113 [physics.ins-det].
- [75] M. Aoki. “An experimental search for muon-electron conversion in nuclear field at sensitivity of 10^{-14} with a pulsed proton beam”. In: *AIP Conf.Proc.* 1441 (2012), pp. 599–601. DOI: 10.1063/1.3700628.
- [76] R. Abramishvili et al. “COMET Phase-I Technical Design Report”. In: (2016).
- [77] L. Bartoszek et al. “Mu2e Technical Design Report”. In: (2014). arXiv: 1501.05241.
- [78] R. J. Barlow. “The PRISM/PRIME project”. In: *Nucl. Phys. Proc. Suppl.* 218 (2011), pp. 44–49. DOI: 10.1016/j.nuclphysbps.2011.06.009.
- [79] J. Adam et al. “The MEG detector for $\mu^+ \rightarrow e^+\gamma$ decay search”. In: *Eur. Phys. J. C* 73.4 (2013), p. 2365. DOI: 10.1140/epjc/s10052-013-2365-2. arXiv: 1303.2348 [physics.ins-det].
- [80] W. H. Bertl et al. “Search for the Decay $\mu^+ \rightarrow e^+e^+e^-$ ”. In: *Nucl. Phys.* B260 (1985), pp. 1–31. DOI: 10.1016/0550-3213(85)90308-6.
- [81] L. Morescalchi. “The Mu2e Experiment at Fermilab”. In: *PoS DIS2016* (2016), p. 259. arXiv: 1609.02021 [hep-ex].
- [82] B. E. Krikler. “An Overview of the COMET Experiment and its Recent Progress”. In: *17th International Workshop on Neutrino Factories and Future Neutrino Facilities Search (NuFact15)*. 2015. arXiv: 1512.08564 [physics.ins-det]. URL: <http://inspirehep.net/record/1411643/files/arXiv:1512.08564.pdf>.
- [83] F. Wilczek. “Axions and Family Symmetry Breaking”. In: *Phys. Rev. Lett.* 49 (1982), pp. 1549–1552. DOI: 10.1103/PhysRevLett.49.1549.
- [84] D. B. Reiss. “Can the Family Group Be a Global Symmetry?”. In: *Phys. Lett.* 115B (1982), pp. 217–220. DOI: 10.1016/0370-2693(82)90647-5.

- [85] J. E. Kim. “Light Pseudoscalars, Particle Physics and Cosmology”. In: *Phys. Rept.* 150 (1987), pp. 1–177. DOI: 10.1016/0370-1573(87)90017-2.
- [86] B. Grinstein, J. Preskill, and M. B. Wise. “Neutrino Masses and Family Symmetry”. In: *Phys. Lett.* 159B (1985), pp. 57–61. DOI: 10.1016/0370-2693(85)90119-4.
- [87] A. Pilaftsis. “Astrophysical and terrestrial constraints on singlet Majoron models”. In: *Phys. Rev. D* 49 (1994), pp. 2398–2404. DOI: 10.1103/PhysRevD.49.2398. arXiv: hep-ph/9308258 [hep-ph].
- [88] J. L. Feng, T. Moroi, H. Murayama, and E. Schnapka. “Third generation familons, b factories, and neutrino cosmology”. In: *Phys. Rev. D* 57 (1998), pp. 5875–5892. DOI: 10.1103/PhysRevD.57.5875. arXiv: hep-ph/9709411 [hep-ph].
- [89] M. Hirsch, A. Vicente, J. Meyer, and W. Porod. “Majoron emission in muon and tau decays revisited”. In: *Phys. Rev. D* 79 (2009). [Erratum: *Phys. Rev. D* 79, 079901 (2009)], p. 055023. DOI: 10.1103/PhysRevD.79.055023. arXiv: 0902.0525 [hep-ph].
- [90] J. Jaeckel. “A Family of WISPy Dark Matter Candidates”. In: *Phys. Lett. B* 732 (2014), pp. 1–7. DOI: 10.1016/j.physletb.2014.03.005. arXiv: 1311.0880 [hep-ph].
- [91] A. Celis, J. Fuentes-Martin, and H. Serodio. “An invisible axion model with controlled FCNCs at tree level”. In: *Phys. Lett. B* 741 (2015), pp. 117–123. DOI: 10.1016/j.physletb.2014.12.028. arXiv: 1410.6217 [hep-ph].
- [92] A. Celis, J. Fuentes-Martín, and H. Serôdio. “A class of invisible axion models with FCNCs at tree level”. In: *JHEP* 12 (2014), p. 167. DOI: 10.1007/JHEP12(2014)167. arXiv: 1410.6218 [hep-ph].
- [93] Y. Ema, K. Hamaguchi, T. Moroi, and K. Nakayama. “Flaxion: a minimal extension to solve puzzles in the standard model”. In: *JHEP* 01 (2017), p. 096. DOI: 10.1007/JHEP01(2017)096. arXiv: 1612.05492 [hep-ph].
- [94] C. Garcia-Cely and J. Heeck. “Neutrino Lines from Majoron Dark Matter”. In: *JHEP* 05 (2017), p. 102. DOI: 10.1007/JHEP05(2017)102. arXiv: 1701.07209 [hep-ph].
- [95] J. Goldstone, A. Salam, and S. Weinberg. “Broken Symmetries”. In: *Phys. Rev.* 127 (1962), pp. 965–970. DOI: 10.1103/PhysRev.127.965.

- [96] S. Weinberg. “Approximate symmetries and pseudo-Goldstone bosons”. In: *Phys. Rev. Lett.* 29 (1972), pp. 1698–1701. DOI: 10.1103/PhysRevLett.29.1698.
- [97] R.D. Peccei and H.R. Quinn. “CP Conservation in the Presence of Instantons”. In: *Phys. Rev. Lett.* 38 (1977), pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440.
- [98] A. Ringwald. “Axions and Axion-Like Particles”. In: *Proceedings, 49th Rencontres de Moriond on Electroweak Interactions and Unified Theories*. 2014, pp. 223–230. arXiv: 1407.0546 [hep-ph]. URL: <http://inspirehep.net/record/1304441/files/arXiv:1407.0546.pdf>.
- [99] C. D. Carone. “Pseudo-Familon Dark Matter”. In: *Phys. Lett. B* 707 (2012), pp. 529–533. DOI: 10.1016/j.physletb.2012.01.010. arXiv: 1201.1205 [hep-ph].
- [100] A. S. Joshipura. “Generation Dependent Peccei-Quinn Symmetry and Neutrino Masses”. In: *Z. Phys.* C38 (1988), p. 479. DOI: 10.1007/BF01584399.
- [101] Z. G. Berezhiani and M. Yu. Khlopov. “Cosmology of Spontaneously Broken Gauge Family Symmetry”. In: *Z. Phys.* C49 (1991), pp. 73–78. DOI: 10.1007/BF01570798.
- [102] M. Bauer, T. Schell, and T. Plehn. “Hunting the Flavon”. In: *Phys. Rev. D* 94.5 (2016), p. 056003. DOI: 10.1103/PhysRevD.94.056003. arXiv: 1603.06950 [hep-ph].
- [103] K. Tsumura and L. Velasco-Sevilla. “Phenomenology of flavon fields at the LHC”. In: *Phys. Rev. D* 81 (2010), p. 036012. DOI: 10.1103/PhysRevD.81.036012. arXiv: 0911.2149 [hep-ph].
- [104] C.-Y. Chen, H. Davoudiasl, W. J. Marciano, and C. Zhang. “Implications of a light “dark Higgs” solution to the $g_{\mu}-2$ discrepancy”. In: *Phys. Rev. D* 93.3 (2016), p. 035006. DOI: 10.1103/PhysRevD.93.035006. arXiv: 1511.04715 [hep-ph].
- [105] J. T. Goldman et al. “Light Boson Emission in the Decay of the μ^+ ”. In: *Phys. Rev. D* 36 (1987), pp. 1543–1546. DOI: 10.1103/PhysRevD.36.1543.
- [106] A. Jodidio et al. “Search for Right-Handed Currents in Muon Decay”. In: *Phys. Rev. D* 34 (1986). [Erratum: *Phys. Rev. D* 37, 237(1988)], p. 1967. DOI: 10.1103/PhysRevD.34.1967, 10.1103/PhysRevD.37.237.

- [107] R. Eichler et al. “Limits for shortlived neutral particles emitted in μ^+ or π^+ decay”. In: *Phys.Lett.* B175 (1986), p. 101. DOI: 10.1016/0370-2693(86)90339-4.
- [108] R. Bayes et al. “Search for two body muon decay signals”. In: *Phys. Rev.* D91.5 (2015), p. 052020. DOI: 10.1103/PhysRevD.91.052020. arXiv: 1409.0638 [hep-ex].
- [109] R. S. Henderson et al. “Precision planar drift chambers and cradle for the TWIST muon decay spectrometer”. In: *Nucl. Instrum. Meth.* A548 (2005), pp. 306–335. DOI: 10.1016/j.nima.2005.02.043. arXiv: hep-ex/0409066 [hep-ex].
- [110] J. Alexander et al. “Dark Sectors 2016 Workshop: Community Report”. In: 2016. arXiv: 1608.08632 [hep-ph]. URL: <http://inspirehep.net/record/1484628/files/arXiv:1608.08632.pdf>.
- [111] R. Essig et al. “Working Group Report: New Light Weakly Coupled Particles”. In: *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013)*. 2013. arXiv: 1311.0029 [hep-ph]. URL: <http://inspirehep.net/record/1263039/files/arXiv:1311.0029.pdf>.
- [112] L. Ackerman, M. R. Buckley, S. M. Carroll, and M. Kamionkowski. “Dark Matter and Dark Radiation”. In: *Phys. Rev.* D79 (2009). [,277(2008)], p. 023519. DOI: 10.1103/PhysRevD.79.023519, 10.1142/9789814293792_0021. arXiv: 0810.5126 [hep-ph].
- [113] J.L. Hewett, H. Weerts, R. Brock, J.N. Butler, B.C.K. Casey, et al. “Fundamental Physics at the Intensity Frontier”. In: (2012). arXiv: 1205.2671 [hep-ex].
- [114] J. Jaeckel and A. Ringwald. “The Low-Energy Frontier of Particle Physics”. In: *Ann. Rev. Nucl. Part. Sci.* 60 (2010), pp. 405–437. DOI: 10.1146/annurev.nucl.012809.104433. arXiv: 1002.0329 [hep-ph].
- [115] B. Holdom. “Two U(1)’s and Epsilon Charge Shifts”. In: *Phys. Lett.* 166B (1986), pp. 196–198. DOI: 10.1016/0370-2693(86)91377-8.
- [116] P. Galison and A. Manohar. “Two Z’s or not two Z’s?” In: *Phys. Lett.* 136B (1984), pp. 279–283. DOI: 10.1016/0370-2693(84)91161-4.

- [117] K. R. Dienes, C. F. Kolda, and J. March-Russell. “Kinetic mixing and the supersymmetric gauge hierarchy”. In: *Nucl. Phys. B* 492 (1997), pp. 104–118. DOI: 10.1016/S0550-3213(97)80028-4, 10.1016/S0550-3213(97)00173-9. arXiv: hep-ph/9610479 [hep-ph].
- [118] B. Echenard, R. Essig, and Y.-M. Zhong. “Projections for Dark Photon Searches at Mu3e”. In: *JHEP* 01 (2015), p. 113. DOI: 10.1007/JHEP01(2015)113. arXiv: 1411.1770 [hep-ph].
- [119] M. Pospelov. “Secluded U(1) below the weak scale”. In: *Phys. Rev. D* 80 (2009), p. 095002. DOI: 10.1103/PhysRevD.80.095002. arXiv: 0811.1030 [hep-ph].
- [120] G. W. Bennett et al. “Final Report of the Muon E821 Anomalous Magnetic Moment Measurement at BNL”. In: *Phys. Rev. D* 73 (2006), p. 072003. DOI: 10.1103/PhysRevD.73.072003. arXiv: hep-ex/0602035 [hep-ex].
- [121] G. W. Bennett et al. “Measurement of the negative muon anomalous magnetic moment to 0.7 ppm”. In: *Phys. Rev. Lett.* 92 (2004), p. 161802. DOI: 10.1103/PhysRevLett.92.161802. arXiv: hep-ex/0401008 [hep-ex].
- [122] N. Arkani-Hamed, D. P. Finkbeiner, T. R. Slatyer, and N. Weiner. “A Theory of Dark Matter”. In: *Phys. Rev. D* 79 (2009), p. 015014. DOI: 10.1103/PhysRevD.79.015014. arXiv: 0810.0713 [hep-ph].
- [123] M. Pospelov. “Light and dark: A survey of new physics ideas in the 1-100 MeV window”. In: *AIP Conf. Proc.* 1563 (2013), pp. 51–56. DOI: 10.1063/1.4829375.
- [124] O. Adriani et al. “An anomalous positron abundance in cosmic rays with energies 1.5-100 GeV”. In: *Nature* 458 (2009), pp. 607–609. DOI: 10.1038/nature07942. arXiv: 0810.4995 [astro-ph].
- [125] O. Adriani et al. “Ten years of PAMELA in space”. In: *Riv. Nuovo Cim.* 40.10 (2017), p. 1. DOI: 10.1393/ncr/i2017-10140-x. arXiv: 1801.10310 [astro-ph.HE].
- [126] M. Ackermann et al. “Measurement of separate cosmic-ray electron and positron spectra with the Fermi Large Area Telescope”. In: *Phys. Rev. Lett.* 108 (2012), p. 011103. DOI: 10.1103/PhysRevLett.108.011103. arXiv: 1109.0521 [astro-ph.HE].

- [127] M. Aguilar et al. “Cosmic-ray positron fraction measurement from 1 to 30-GeV with AMS-01”. In: *Phys. Lett. B* 646 (2007), pp. 145–154. DOI: 10.1016/j.physletb.2007.01.024. arXiv: astro-ph/0703154 [ASTRO-PH].
- [128] M. Aguilar et al. “First Result from the Alpha Magnetic Spectrometer on the International Space Station: Precision Measurement of the Positron Fraction in Primary Cosmic Rays of 0.5–350 GeV”. In: *Phys. Rev. Lett.* 110 (2013), p. 141102. DOI: 10.1103/PhysRevLett.110.141102.
- [129] A. J. Krasznahorkay et al. “Observation of Anomalous Internal Pair Creation in Be8 : A Possible Indication of a Light, Neutral Boson”. In: *Phys. Rev. Lett.* 116.4 (2016), p. 042501. DOI: 10.1103/PhysRevLett.116.042501. arXiv: 1504.01527 [nucl-ex].
- [130] E. Goudzovski. “Neutral pion form factor measurement at NA62”. In: *PoS ICHEP2016* (2017), p. 642. arXiv: 1611.04972 [hep-ex].
- [131] J. Heeck and W. Rodejohann. “Gauged $L_\mu - L_\tau$ Symmetry at the Electroweak Scale”. In: *Phys. Rev. D* 84 (2011), p. 075007. DOI: 10.1103/PhysRevD.84.075007. arXiv: 1107.5238 [hep-ph].
- [132] W. Altmannshofer, S. Gori, M. Pospelov, and I. Yavin. “Quark flavor transitions in $L_\mu - L_\tau$ models”. In: *Phys. Rev. D* 89 (2014), p. 095033. DOI: 10.1103/PhysRevD.89.095033. arXiv: 1403.1269 [hep-ph].
- [133] M. Bauer. “Hidden Photons”. Talk at the Mu3e external collaboration meeting. 2017.
- [134] M. Bauer, P. Foldenauer, and J. Jaeckel. “Hunting All the Hidden Photons”. In: (2018). arXiv: 1803.05466 [hep-ph].
- [135] A. Blondel et al. “Letter of Intent for an Experiment to Search for the Decay $\mu \rightarrow eee$ ”. In: (2012).
- [136] A. Blondel et al. “Technical Design of the Phase I Mu3e Experiment”. to be published. 2018.
- [137] A. Signer et al. “Matrix Element of the Rare Muon Decay”. Private Communication. 2015.
- [138] R. M. Djilkibaev and R. V. Konoplich. “Rare Muon Decay $\mu^+ \rightarrow e^+ e^- e^+ \nu_e \bar{\nu}_\mu$ ”. In: *Phys. Rev. D* 79 (2009), p. 073004. DOI: 10.1103/PhysRevD.79.073004. arXiv: 0812.1355 [hep-ph].

- [139] V. L. Highland. “Some Practical Remarks on Multiple Scattering”. In: *Nucl.Instrum.Meth.* 129 (1975), p. 497. DOI: 10 . 1016/0029-554X(75)90743-0.
- [140] F. A. Berg. “CMBL - A High-intensity Muon Beam Line & Scintillation Target with Monitoring System for Next-generation Charged Lepton Flavour Violation Experiments”. PhD thesis. ETH Zürich, 2017.
- [141] F. Berg et al. “Target Studies for Surface Muon Production”. In: (2015). arXiv: 1511.01288 [physics.ins-det].
- [142] A. M. Baldini et al. “Muon polarization in the MEG experiment: predictions and measurements”. In: *Eur. Phys. J. C* 76.4 (2016), p. 223. DOI: 10 . 1140/epjc/s10052-016-4047-3. arXiv: 1510.04743 [hep-ex].
- [143] S. Haywood. “The ATLAS Inner Detector”. In: *Nucl. Instrum. Meth.* A408 (1998), pp. 242–250. DOI: 10 . 1016 / S0168 - 9002(98)00248-4.
- [144] CMS Collaboration. *CMS, tracker technical design report*. 1998. URL: <http://weplib.cern.ch/abstract?CERN-LHCC-98-6>.
- [145] I. Perić. “A novel monolithic pixelated particle detector implemented in high-voltage CMOS technology”. In: *Nucl. Instrum.Meth.* A582 (2007), p. 876. DOI: 10 . 1016 / j . nima . 2007 . 07 . 115.
- [146] I. Perić et al. “High-voltage pixel detectors in commercial CMOS technologies for ATLAS, CLIC and Mu3e experiments”. In: *Nucl.Instrum.Meth.* A731 (2013), pp. 131–136. DOI: 10 . 1016/j.nima.2013.05.006.
- [147] H. Augustin, N. Berger, S. Bravar, S. Corrodi, A. Damyanova, et al. “The MuPix high voltage monolithic active pixel sensor for the Mu3e experiment”. In: *JINST* 10.03 (2015), p. C03044. DOI: 10 . 1088/1748-0221/10/03/C03044.
- [148] H. Augustin et al. “The MuPix System-on-Chip for the Mu3e Experiment”. In: *Nucl. Instrum. Meth.* A845 (2017), pp. 194–198. DOI: 10 . 1016 / j . nima . 2016 . 06 . 095. arXiv: 1603 . 08751 [physics.ins-det].
- [149] H. Augustin et al. “MuPix7 - A fast monolithic HV-CMOS pixel chip for Mu3e”. In: *submitted to JINST* (2016). arXiv: 1610.02210 [physics.ins-det].

- [150] W. R. Leo. *Techniques for nuclear and particle physics experiments. a how-to approach; with 40 tables and numerous worked examples.* 2., rev. ed. Berlin ; Heidelberg [et al.]: Springer, 1994, XVIII, 378 S.
- [151] A. Weber and I. Perić. “Documentation MuPix8”. preliminary version 1. 2017.
- [152] LHCb Collaboration. *LHCb Tracker Upgrade Technical Design Report.* Tech. rep. CERN-LHCC-2014-001. LHCb-TDR-015. CERN, 2014. URL: <http://cds.cern.ch/record/1647400>.
- [153] H. Chen, K. Briggli, P. Eckert, T. Harion, Y. Munwes, W. Shen, V. Stankova, and H. C. Schultz-Coulon. “MuTRiG: a mixed signal Silicon Photomultiplier readout ASIC with high timing resolution and gigabit data link”. In: *JINST* 12.01 (2017), p. C01043. DOI: 10.1088/1748-0221/12/01/C01043.
- [154] S. Corrodi. PhD thesis in preparation, ETH Zürich.
- [155] P. Durante, N. Neufeld, R. Schwemmer, G. Balbi, and U. Marconi. “100 Gbps PCI-Express readout for the LHCb upgrade”. In: *JINST* 10.04 (2015), p. C04018. DOI: 10.1088/1748-0221/10/04/C04018.
- [156] *PCI Express Base Specification.* Revision 3.1. PCI SIG, 2014.
- [157] D. vom Bruch. “Pixel sensor evaluation and online event selection for the Mu3e experiment”. PhD thesis. Ruperto-Carola-University of Heidelberg, 2017. DOI: 10.11588/heidok.00023689.
- [158] A. Schöning. “Track Fit with Multiple Scattering”. Mu3e Internal Note. 2014.
- [159] A. Schöning. “A Three-Dimensional Helix Fit with Multiple Scattering using Hit Triplets”. Mu3e Internal Note. 2014.
- [160] 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”. In: *JINST* 9.12 (2014), p. C12012. DOI: 10.1088/1748-0221/9/12/C12012.
- [161] 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”. In: *Journal of Instrumentation* 9.12 (2014), p. C12012. URL: <http://stacks.iop.org/1748-0221/9/i=12/a=C12012>.

- [162] N. Berger, M. Kiehn, A. Kozlinskiy, and A. Schöning. “A New Three-Dimensional Track Fit with Multiple Scattering”. In: *Nucl. Instr. Meth A*. 844 C (2017), pp. 135–140. arXiv: 1606.04990 [physics.ins-det].
- [163] S. Schenk. “A Vertex Fit for Low Momentum Particles in a Solenoidal Magnetic Field with Multiple Scattering”. BSc Thesis. Heidelberg University, 2013.
- [164] A. Schöning. “Linearised Vertex 3D Fit in a Solenoidal Magnetic field with Multiple Scattering”. Mu3e Internal Note. 2013.
- [165] M. Kiehn and S. Schenk. “Scattering Vertex Fit Notes”. Mu3e Internal Note.
- [166] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nucl. Instr. Meth. A* 506.3 (2003), pp. 250–303. DOI: DOI : 10.1016/S0168-9002(03)01368-8. URL: <http://www.sciencedirect.com/science/article/pii/S0168900203013688>.
- [167] *Numerical Recipes. The Art of Scientific Computing*. eng. Cambridge: Cambridge Univ. Pr., 1986, XX, 818 S.
- [168] D. E. Knuth. *The Art of Computer Programming*. 3rd. Vol. 2: Seminumerical Algorithms. Addison-Wesley, 2012.
- [169] W.E. Fischer and F. Scheck. “Electron Polarization in Polarized Muon Decay: Radiative Corrections”. In: *Nucl.Phys.* B83 (1974), p. 25. DOI: 10.1016/0550-3213(74)90071-6.
- [170] F. Scheck. “Muon Physics”. In: *Phys.Rept.* 44 (1978), p. 187. DOI: 10.1016/0370-1573(78)90014-5.
- [171] P. Depommier and A. Vacheret. *Radiative Muon Decay*. Tech. rep. TWIST Technote No 55, 2001.
- [172] R. Kleiss, W.J. Stirling, and S.D. Ellis. “A new Monte Carlo treatment of multiparticle phase space at high energies”. In: *Comp. Phys. Commun.* 40 (1986), pp. 359–373. DOI: 10.1016/0010-4655(86)90119-0. URL: <http://www.sciencedirect.com/science/article/pii/0010465586901190>.
- [173] A. Knecht. “The high-intensity muon beam line (HiMB) project at PSI”. In: *PoS NuFact2017* (2018). to be published, p. 100.

- [174] G. Pellegrini et al. “Technology developments and first measurements of Low Gain Avalanche Detectors (LGAD) for high energy physics applications”. In: *Nucl. Instrum. Meth.* A765 (2014), pp. 12–16. DOI: 10.1016/j.nima.2014.06.008.
- [175] G. Pellegrini et al. “Recent Technological Developments on LGAD and iLGAD Detectors for Tracking and Timing Applications”. In: *Nucl. Instrum. Meth.* A831 (2016), pp. 24–28. DOI: 10.1016/j.nima.2016.05.066. arXiv: 1511.07175 [physics.ins-det].
- [176] H. F. W. Sadrozinski et al. “Ultra-fast silicon detectors (UFSD)”. In: *Nucl. Instrum. Meth.* A831 (2016), pp. 18–23. DOI: 10.1016/j.nima.2016.03.093.
- [177] A. Schöning. “Mu3e-gamma”. Talk at the Mu3e external collaboration meeting, 2017.
- [178] C. Cheng, B. Echenard, and D. G. Hitlin. “The next generation of $\mu \rightarrow e\gamma$ and $\mu \rightarrow 3e$ CLFV search experiments”. In: *Proceedings, 2013 Community Summer Study on the Future of U.S. Particle Physics: Snowmass on the Mississippi (CSS2013)*. 2013. arXiv: 1309.7679 [physics.ins-det]. URL: <http://inspirehep.net/record/1256023/files/arXiv:1309.7679.pdf>.
- [179] F. Gray. *Pulse code communication*. US Patent 2,632,058. 1953. URL: <http://www.google.com/patents/US2632058>.
- [180] K. Boyette. *OpenCores 8bit/10bit encoder and decoder*. URL: http://opencores.org/project,8b10b_encdec.
- [181] P. A. Franaszek and A. X. Widmer. *Byte oriented DC balanced (0,4) 8B/10B partitioned block transmission code*. US Patent 4,486,739. 1984. URL: <http://www.google.com/patents/US4486739>.
- [182] A. X. Widmer and P. A. Franaszek. “A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code”. In: *IBM Journal of Research and Development* 27 (1983), p. 440.
- [183] Altera Corporation. *Arria V Device Datasheet*. 2017. URL: https://www.altera.com/en_US/pdfs/literature/hb/arria-v/av_51002.pdf.
- [184] H. Augustin et al. “The MuPix Telescope: A Thin, high Rate Tracking Telescope”. In: *JINST* 12.01 (2017), p. C01087. DOI: 10.1088/1748-0221/12/01/C01087. arXiv: 1611.03102 [physics.ins-det].

- [185] Altera Corporation. *Stratix IV GX FPGA Development Board – Reference Manual*. 2012. URL: https://www.altera.com/en_US/pdfs/literature/manual/rm_sivgx_fpga_dev_board.pdf.
- [186] Altera Corporation. *Stratix IV Device Handbook – Volume 1*. 2016. URL: https://www.altera.com/en_US/pdfs/literature/hb/stratix-iv/stratix4_handbook.pdf.
- [187] L. Huth. “Development of a Tracking Telescope for Low Momentum Particles and High Rates consisting of HV-MAPS”. MA thesis. Ruprecht-Karls Universität Heidelberg, 2014.
- [188] Mentor Graphics. *ModelSim Altera Starter Edition 10.4b*. URL: <https://www.mentor.com/products/fv/modelsim/>.
- [189] Wikipedia. *Pseudorandom binary sequence*. 2015. URL: https://en.wikipedia.org/wiki/Pseudorandom_binary_sequence.
- [190] J. Kröger. “Readout Hardware for the MuPix8 and a Firmware-based MuPix8 Emulator”. MA thesis. Ruprecht-Karls Universität Heidelberg, 2017.
- [191] L. Landau. “On the energy loss of fast particles by ionization”. In: *J. Phys.(USSR)* 8 (1944), pp. 201–205.
- [192] M. Fael, L. Mercolli, and M. Passera. “Radiative μ and τ leptonic decays at NLO”. In: *JHEP* 07 (2015), p. 153. DOI: 10.1007/JHEP07(2015)153. arXiv: 1506.03416 [hep-ph].
- [193] G. M. Pruna, A. Signer, and Y. Ulrich. “Fully differential NLO predictions for the radiative decay of muons and taus”. In: *Phys. Lett. B* 772 (2017), pp. 452–458. DOI: 10.1016/j.physletb.2017.07.008. arXiv: 1705.03782 [hep-ph].
- [194] G. M. Pruna, A. Signer, and Y. Ulrich. “Fully differential NLO predictions for the rare muon decay”. In: (2016). arXiv: 1611.03617 [hep-ph].
- [195] M. Fael and C. Greub. “Next-to-leading order prediction for the decay $\mu \rightarrow e(e^+e^-)\nu\bar{\nu}$ ”. In: (2016). arXiv: 1611.03726 [hep-ph].
- [196] A. Signer, G. M. Pruna, and Y. Ulrich. Private Communication. 2017.
- [197] G. Rutar, C. Bemporad, P. W. Cattaneo, F. Cei, L. Galli, P. R. Kettle, and A. Papa. “A dedicated calibration tool for the MEG and MEG II positron spectrometer”. In: *Nucl. Instrum. Meth.* A824 (2016), pp. 575–577. DOI: 10.1016/j.nima.2015.11.121.

- [198] W. Verkerke and D. P. Kirkby. “The RooFit toolkit for data modeling”. In: *eConf C0303241* (2003). [,186(2003)], MOLT007. arXiv: physics/0306116 [physics].
- [199] R. Brun et al. *ROOT User’s Guide – An Object-Oriented Data Analysis Framework*. 5.34/36. Geneva, 2013.
- [200] J. H. Friedman. “Data Analysis Techniques for High-Energy Particle Physics”. In: *1974 CERN School of Computing, Godoyssund, Norway, 11-24 Aug 1974: Proceedings*. 1974, p. 271.
- [201] G. Cowan. *Statistical data analysis. [with applications from particle physics]*. Oxford science publications. Oxford: Clarendon Press, 1998.
- [202] F. Wauters. “Muon Beam Polarisation at PSI”. Talk at the Mu3e external collaboration meeting. 2017.
- [203] L. Gerritzen. “Reconstruction of Two-Prong Signatures with a Linearised Multiple Scattering Vertex Fit for Mu3e”. MA thesis. Ruprecht-Karls Universität Heidelberg, 2017.
- [204] C. Kresse. “Track Reconstruction of Photon Conversion Electrons from Displaced Vertices in the Mu3e Detector”. MA thesis. Ruprecht-Karls Universität Heidelberg, 2017.
- [205] H. Leuschner. Masters thesis in preparation, Ruprecht-Karls Universität Heidelberg. 2018.
- [206] M. Bauer. “Familon”. Talk at the Mu3e external collaboration meeting. 2017.
- [207] M. Bauer, P. Foldenauer, J. Jaeckel, and A. Schöning. Private communication. 2017.
- [208] A. Schöning. “Dark Photon Searches”. Talk at the Mu3e external collaboration meeting. 2018.
- [209] U. Hartenstein. PhD thesis in preparation, Johannes Gutenberg University Mainz.
- [210] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr, and B. Fuks. “FeynRules 2.0 - A complete toolbox for tree-level phenomenology”. In: *Comput. Phys. Commun.* 185 (2014), pp. 2250–2300. DOI: 10.1016/j.cpc.2014.04.012. arXiv: 1310.1921 [hep-ph].
- [211] C. Degrande, C. Duhr, B. Fuks, D. Grellscheid, O. Mattelaer, and T. Reiter. “UFO - The Universal FeynRules Output”. In: *Comput. Phys. Commun.* 183 (2012), pp. 1201–1214. DOI: 10.1016/j.cpc.2012.01.022. arXiv: 1108.2040 [hep-ph].

- [212] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. “The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations”. In: *JHEP* 07 (2014), p. 079. DOI: 10 . 1007/JHEP07(2014)079. arXiv: 1405 . 0301 [hep-ph].
- [213] J. Alwall et al. “A Standard format for Les Houches event files”. In: *Comput. Phys. Commun.* 176 (2007), pp. 300–304. DOI: 10 . 1016/j . cpc . 2006 . 11 . 010. arXiv: hep - ph / 0609017 [hep-ph].
- [214] J. Alwall, E. Boos, L. Dudko, M. Gigg, M. Herquet, A. Pukhov, P. Richardson, A. Sherstnev, and P. Z. Skands. “A Les Houches Interface for BSM Generators”. In: 2007. DOI: 10 . 2172 / 921331. arXiv: 0712 . 3311 [hep-ph]. URL: http://lss.fnal.gov/cgi-bin/find_paper.pl?conf-07-621.
- [215] M. Baszczyk et al. “SuperB Technical Design Report”. In: (2013). arXiv: 1306.5655 [physics.ins-det].
- [216] M. Ablikim et al. “Dark Photon Search in the Mass Range Between 1.5 and 3.4 GeV/c²”. In: *Phys. Lett. B* 774 (2017), pp. 252–257. DOI: 10 . 1016/j . physletb . 2017 . 09 . 067. arXiv: 1705 . 04265 [hep-ex].
- [217] D. Becker et al. “The P2 Experiment - A future high-precision measurement of the electroweak mixing angle at low momentum transfer”. In: (2018). arXiv: 1802 . 04759 [nucl-ex].
- [218] K. T. Brinkmann, P. Gianotti, and I. Lehmann. “Exploring the mysteries of strong interactions: The PANDA experiment”. In: *Nucl. Phys. News* 16 (2006), pp. 15–18. DOI: 10 . 1080 / 10506890600579868. arXiv: physics/0701090 [physics].
- [219] I. Perić and N. Berger. “MUPIX7 Data Format”. Mu3e Internal Note.
- [220] N. Berger. “MUPIX8 Data Format”. Mu3e Internal Note.
- [221] R. H. Dalitz. “Decay of tau mesons of known charge”. In: *Phys. Rev.* 94 (1954), pp. 1046–1051. DOI: 10 . 1103/PhysRev . 94 . 1046.

ACKNOWLEDGMENTS

The final page shall be dedicated to all those who supported me during the doctoral studies.

First of all, I would like to thank my supervisor Niklaus Berger for offering this exciting topic, and for his guidance and support throughout the whole time of my PhD studies. I am impressed by his expertise on so many fields, and I am glad he shared it with me.

I am grateful that Stephanie Hansmann-Menzemer agreed to be the second referee.

Thanks to André Schöning, whose sense for interesting physics I will always admire, for initially drawing my attention to the *Mu3e* experiment and for constant advice during the PhD.

I would like to express my thanks to the *Mu3e* collaboration. It is a pleasure to work with such passionate people.

Special thanks go to the working group in Heidelberg for the pleasant atmosphere, and countless coffee breaks, dinners, barbecues, test beam campaigns, ...

My dear fellow doctoral students, I enjoyed working together with you from day one on. Thanks for being such a great team: Heiko Augustin, Dorothea vom Bruch, Sebastian Dittmeier, Lennart Huth, Adrian Herkert, and Moritz Kiehn.

I had the great opportunity to supervise outstanding master students. Lukas Gerritzen, Caren Kresse, and Hendrik Leuschner, thank you for teaching me just as much as I taught you.

A special thanks goes to all the diligent proof-readers: Heiko Augustin, Sebastian Dittmeier, Lukas Gerritzen, Adrian Herkert, Lennart Huth, Ivan Perić, and André Schöning.

For their patience with my questions and their passion for the topic, I like to thank my colleagues from theoretical physics: Partick Foldenauer, Giovanni Marco Pruna, Adrian Signer, and Yannick Ulrich.

And finally, I would like to thank my mother. Without her constant support this work would not have been possible.