

Data Flow in the Mu3e Data Acquisition System



DISSERTATION

zur Erlangung des Grades
„Doktor der Naturwissenschaften“
am Fachbereich Physik, Mathematik und Informatik
der Johannes Gutenberg-Universität in Mainz

vergelegt von

Marius Köppel

geboren in Bühl

Mainz, den 10.01.2024

Abstract

The discovery of the Higgs boson by the ATLAS [1] and Compact Muon Solenoid Experiment (CMS) [2] experiments at CERN's Large Hadron Collider (LHC) [3] in 2012 marked a milestone in particle physics, completing the Standard Model of particle physics (SM) as a self-consistent theory. However, the observation of neutrino oscillations shows violations of lepton-flavour conservation, hinting at physics beyond the SM. Charged lepton-flavour violation (CLFV) decays, such as $\mu^+ \rightarrow e^+e^-e^+$, remain undetected and would be a definitive sign of new physics.

The Mu3e experiment, currently under construction at the Paul Scherrer Institute (PSI), is specifically designed to observe the decay $\mu^+ \rightarrow e^+e^-e^+$, a process exceptionally rare in the extension of the SM with neutrino oscillation, but enhanced in many theories beyond the SM. Achieving high sensitivity is crucial to detect such a rare decay. Mu3e aims for an ultimate sensitivity of one in 10^{16} muon decays. In its initial phase, Mu3e aims to achieve a branching ratio sensitivity of $2 \cdot 10^{-15}$, analysing 1×10^8 muon decays per second over a year of data collection. The experiment uses highly granular detectors consisting of thin High-Voltage Monolithic Active Pixel Sensors (HV-MAPS), the MuPix chip, and scintillating timing detectors, generating approximately 100 Gbit/s of data at these particle rates.

The Mu3e data acquisition (DAQ) system, based on field programmable gate arrays (FPGAs), is a crucial component of the experiment. It employs a trigger-less readout system to deal with the randomly distributed decay particles of muons at rest. The system sorts, time-aligns, and analyses data in real time, using a filter farm of graphics processing units (GPUs) for track reconstruction.

This thesis presents the integration of subdetectors into the Mu3e DAQ system, focussing on the scintillating timing detectors, time-alignment, and data flow within the filter farm. It discusses the requirements for building the Mu3e DAQ system, including data protocols, data flow algorithms, and online data quality checks. Integration runs and testbeams at various facilities, including Deutsches Elektronen-Synchrotron (DESY), Mainz Microtron (MAMI), and PSI, have been instrumental in the refining of the system.

Moreover, this work encompasses irradiation studies of the MuPix10 chip at MAMI to understand its high-rate operation. In addition to research related to the Mu3e experiment, the development of a prototype detector for muon spin rotation (μ SR) techniques using MuPix11 chips is presented. Initial tests in 2021 and a special run at PSI in 2023 demonstrated the first μ SR signal detection using Si-Pixel detectors, marking the start of a novel methodology in this field.

Zusammenfassung

Die Entdeckung des Higgs-Bosons durch die Experimente ATLAS [1] und CMS [2] am LHC des CERNs im Jahr 2012 war ein Meilenstein in der Teilchenphysik und vervollständigte das Standardmodell der Teilchenphysik zu einer selbstkonsistenten Theorie. Die Beobachtung von Neutrino-Oszillationen zeigt jedoch Verstöße gegen die Lepton-Flavour-Erhaltung, was auf eine Physik jenseits des Standardmodells hindeutet. Geladene Lepton-Flavour-verletzende Zerfälle, wie $\mu^+ \rightarrow e^+ e^- e^+$, wurden bisher nicht beobachtet und wären ein eindeutiges Zeichen für neue Physik.

Das Mu3e-Experiment, das derzeit am PSI aufgebaut wird, ist speziell darauf ausgelegt, den Zerfall $\mu^+ \rightarrow e^+ e^- e^+$ zu beobachten, einen Prozess, der in der Erweiterung des Standardmodells mit Neutrino-Oszillation außergewöhnlich selten ist, aber in vielen Theorien jenseits des SM verstärkt auftritt. Um einen solch seltenen Zerfall nachzuweisen, ist eine hohe Sensitivität entscheidend. Mu3e strebt eine ultimative Sensitivität von einem in 10^{16} Myonenzerfällen an. In der Anfangsphase will Mu3e eine Sensitivität für das Verzweigungsverhältnis von $2 \cdot 10^{-15}$ erreichen, indem es 1×10^8 Myonenzerfälle pro Sekunde über ein Jahr der Datenerfassung analysiert. Das Experiment verwendet hochgranulare Detektoren, die auf dünnen HV-MAPS, dem MuPix-Chip und szintillierenden Zeitdetektoren basieren, und erzeugt bei diesen Teilchenraten etwa 100 Gbit/s an Daten.

Das Mu3e DAQ System, welche auf FPGAs aufbaut, ist ein entscheidendes Bauteil des Experiments. Sie verwendet ein triggerloses Auslesesystem, um mit den zufällig verteilten Zerfallsteilchen von ruhenden Myonen umzugehen. Das System sortiert und analysiert die Daten in Echtzeit und nutzt eine Filterfarm aus GPUs für die Spurrekonstruktion.

In dieser Arbeit wird die Integration von Subdetektoren in das Mu3e DAQ-System vorgestellt, wobei der Schwerpunkt auf den szintillierenden Zeitdetektoren, der Zeitausrichtung und dem Datenfluss innerhalb der Filterfarm liegt. Es werden die Anforderungen für den Aufbau des Mu3e DAQ-Systems erörtert, einschließlich Datenprotokollen, Datenflussalgorithmen und Online-Datenqualitätsprüfungen. Integrationsläufe und Teststrahlen an verschiedenen Einrichtungen, darunter DESY, MAMI und PSI, haben maßgeblich zur Evaluierung des Systems beigetragen.

Darüber hinaus umfasst diese Arbeit auch Bestrahlungsstudien des MuPix10-Chips am MAMI Beschleuniger, um dessen Betrieb bei hohen Teilchenraten zu verstehen. Zusätzlich zu den Forschungsarbeiten im Zusammenhang mit dem Mu3e-Experiment wird die Entwicklung eines Prototyps eines Detektors für μ SR-Techniken unter Verwendung von MuPix11-Chips vorgestellt. Erste Tests im Jahr 2021 und ein spezieller Lauf am PSI im Jahr 2023 zeigen die erste μ SR-Signaldetektion mit Si-Pixel-Detektoren und markierten den Beginn einer neuen Methodik in diesem Bereich.

In Erinnerung an Martin W.

Contents

I	Introduction	1
1	Theory	3
1.1	The Standard Model of particle physics	3
1.2	Muon decays	7
1.2.1	Muon decays in the Standard Model	7
1.2.2	Muon decays beyond the Standard Model	8
1.3	Charged lepton-flavour violation experiments	9
1.3.1	SINDRUM	9
1.3.2	MEG and MEG II	10
2	Principles of Particle Detectors	13
2.1	Silicon detectors	13
2.2	Principle of tracking detectors	14
2.3	Multiple Scattering	15
2.4	High Voltage Monolithic Active Pixel Sensors	16
2.5	Scintillating detectors	17
3	Readout Electronics	19
3.1	Field Programmable Gate Arrays (FPGAs)	19
3.2	Data Transmission in FPGAs	21
3.3	Basics of data transmission	21
3.4	Electrical data transmission and Telegrapher's Equations	23
3.5	Serial data links for Inter-device Communication	24
3.6	High speed links in FPGAs	27
3.7	Optical data transmission	31
3.8	Bit Error Rate Tests	32
3.9	Data readout via PCIe	32
4	The Mu3e Experiment	37
4.1	Signal and background processes	38
4.2	Detector design	39
4.2.1	The Compact Muon Beam Line	40
4.2.2	The Mu3e Target	43
4.2.3	The Mu3e Magnet	43
4.3	MuPix Pixel Sensor	44
4.3.1	Pixel cell electronics	46
4.4	The MuTRiG ASIC	48

4.5	The scintillating fibre detector	50
4.6	The scintillating tile detector	52
II The Data Acquisition of the Mu3e Experiment		55
5	Mu3e Data Acquisition System	57
5.1	Overview of the Mu3e DAQ	57
5.2	Clock and Reset System	59
5.3	Readout ASICs	59
5.4	Data rate requirements	60
5.5	Related work of DAQ systems	61
6	Front-end Board	63
6.1	Front-end board	64
6.2	Hit-time sorter	64
6.3	Front-end board communication protocols	65
6.4	Communication front-end boards to switching boards	65
6.4.1	MuPix communication protocol	66
6.4.2	Tile and SciFi communication protocol	67
6.4.3	Slow control communication protocol	67
6.4.4	Run control signals	68
6.4.5	Idle state	69
6.5	Communication of clock and reset system	69
6.6	The MuTRiG datapath	70
7	Switching Board	73
7.1	Overview of the Switching Board	73
7.2	Data flow of the Switching Board	74
7.3	Hit-time-alignment tree	75
7.4	Detector configuration	77
8	Receiving board and Filter Farm	79
8.1	Receiving board	79
8.2	Data flow of the Filter Farm	80
8.3	MIDAS event builder	82
8.4	DMA readout	84
8.5	Mu3e Online Analyzer	85
8.6	Online event selection	86
III Detector Integration		89
9	First Detector Integration	91
9.1	DESY testbeam facility	91
9.2	DESY 2020 experimental setup	91
9.3	DESY 2020 testbeam results	92

CONTENTS

10 Front-end Board Synchronisation	97
10.1 MAMI testbeam facility	98
10.2 MAMI 2020 experimental setup	98
10.3 MAMI 2020 testbeam results	99
11 Mu3e Integration Run 2021	103
11.1 Mu3e Integration Run setup	103
11.2 Mu3e Integration Run DAQ system	104
11.3 Results of the Mu3e Integration Run 2021	106
11.4 Results of the μ SR measurement	108
12 Mu3e Cosmic Run 2022	111
12.1 Mu3e Cosmic Run setup	111
12.2 Mu3e Cosmic Run DAQ system	113
12.3 Data flow during the Mu3e Cosmic Run	113
12.4 Timing studies during the Mu3e Cosmic Run	114
13 Tile Integration	117
13.1 Tile Integration Run setup	118
13.2 Rate monitoring tests	118
13.3 Time correlation studies	119
IV Irradiation and μSR Studies	121
14 MuPix Irradiation Studies	123
14.1 Experimental setup	124
14.2 Analysis flow	126
14.3 Irradiation results	129
15 Advanced muon spin spectroscopy using MuPix chips	135
15.1 Background and introduction	135
15.1.1 Muon beam at PSI	136
15.1.2 Introduction to the μ SR technique	136
15.2 Conceptual design	137
15.2.1 Current status	138
15.3 Monte-Carlo simulation	139
15.4 Prototype and first beam tests	140
V Conclusion and Outlook	143
16 Conclusion and outlook	145
VI Appendices	147
A Acronyms	148

B Additional Material	152
B.1 The Mu3e Experiment	152
B.2 Switching Board	154
B.3 First Detector Integration	155
B.4 Front-end Board Synchronisation	156
B.5 Mu3e Integration Run 2021	157
B.6 MuPix Irradiation Studies	158
C Publications	160
List of Figures	163
List of Tables	167
Bibliography	169
Acknowledgement	179

Part I

Introduction

The first chapter provides a theoretical introduction to lepton-flavour violation (LFV), followed by a brief explanation of the detector principles used in the Mu3e experiment. The part will conclude with a brief discussion of readout electronics and an introduction to field programmable gate arrays (FPGAs), which are the main building blocks of the Mu3e data acquisition (DAQ) system. Finally, an overview of the experimental setup of the Mu3e experiment will be given.

Theory

Particle physics seeks to understand the fundamental laws of nature on a microscopic scale, and the results of this quest are encapsulated in the Standard Model of particle physics (SM). The SM is a relativistic quantum field theory (QFT), specifically a gauge theory with a $U(1) \times SU(2) \times SU(3)$ symmetry, which describes nature in terms of elementary particles and their interactions.

In Section 1.1, a detailed discussion of the SM based on Griffiths [4] is provided. Due to its relevance to this thesis, muon (μ) decays in and beyond the SM are discussed in Section 1.2.1. The last part of the chapter (Section 1.3) gives an overview of different experiments that search for charged lepton-flavour violation (CLFV) decays.

1.1 The Standard Model of particle physics

The world we observe in our daily lives is constructed from a relatively small number of particles. Atoms serve as the fundamental building blocks of all stable matter on Earth. However, atoms themselves are not fundamental; they are composed of electrons, protons, and neutrons. Although electrons are elementary particles, protons and neutrons consist of quarks, specifically the up quark (u) and the down quark (d). Together with the electron neutrino, these particles form the first generation of spin- $1/2$ particles known as fermions. In high-energy interactions, additional generations of fermions emerge. In total, there are 12 fermions, each with its corresponding antiparticle, that possess the same mass but opposite quantum numbers. The difference between the three generations of fermions is caused by their mass and flavour quantum number, all the other properties are the same.

In the SM, every particle is considered a field, either a fermion or a boson field, as illustrated in Figure 1.1. The bosons consist of four kinds of vector bosons with spin 1 and a scalar boson with spin 0, known as the Higgs boson. While vector bosons mediate interactions between particles, the Higgs boson is a pseudo-Goldstone boson resulting from spontaneous symmetry breaking via the Higgs mechanism, which imparts mass to particles. The Higgs field is typically represented by a complex scalar doublet in the fundamental representation of $SU(2)$. This doublet is often written as:

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} \quad (1.1)$$

Here, ϕ^+ and ϕ^0 are the charged and neutral components of the Higgs doublet, respectively. When the Higgs field acquires a vacuum expectation value (VEV), the $SU(2) \times U(1)$ symmetry is spontaneously broken, resulting in the generation of masses for the W and Z bosons as well as fermions through the Yukawa couplings with the Higgs field. $SU(2)$ (or $SU(2)_L$) is the symmetry group associated with weak isospin, and $U(1)$ (or $U(1)_Y$) is the symmetry group associated with hypercharge. The doublet structure allows for the generation of masses in a way that preserves gauge symmetry.

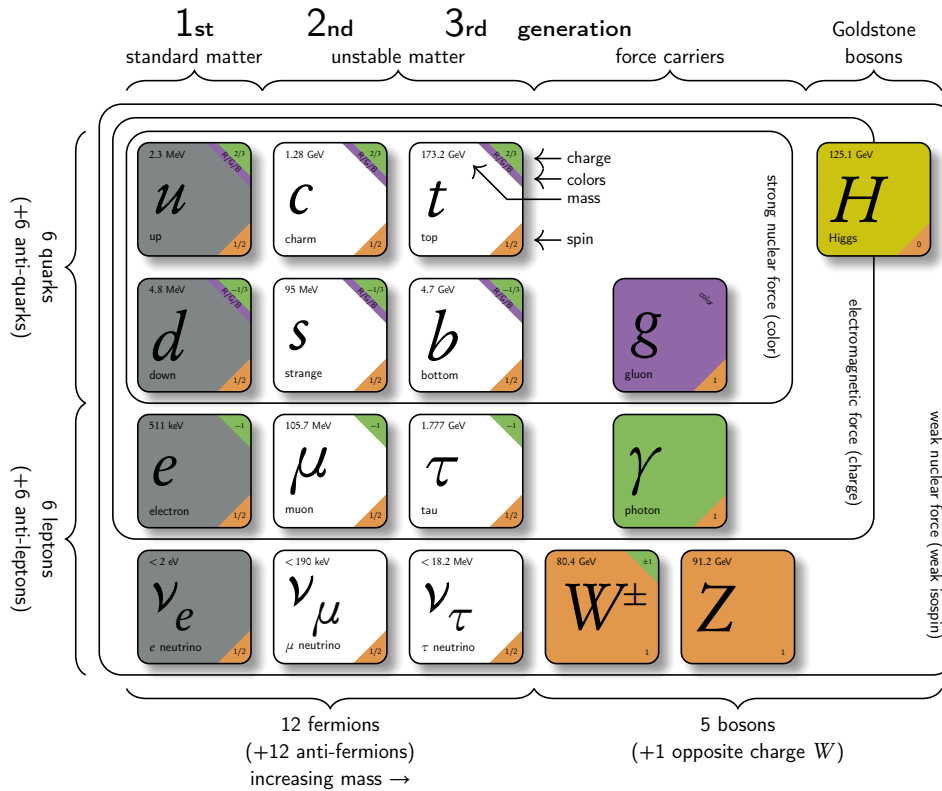


Figure 1.1: Overview of the building blocks of the SM, including the quarks, leptons, and force-carrying bosons, taken from [5].

As mentioned above, the SM describes interactions through four vector bosons, responsible for mediating fundamental forces. The strong interaction is mediated by eight gluons, while the weak interaction is carried out by two charged W^\pm bosons and a neutral Z boson. Electromagnetic interactions are mediated by the photon (γ). Notably, gluons and the γ are massless, while W^\pm and Z bosons have masses of approximately $80.4 \text{ GeV}/c^2$ and $91.2 \text{ GeV}/c^2$, respectively [6]. These boson fields arise from the local gauge invariance of the SM particle fields, as the SM follows a $U(1) \times SU(2) \times SU(3)$ gauge theory. The strong interaction emerges due to the symmetry of local $SU(3)$ transformations.

The SM successfully unifies the weak and electromagnetic forces into the electroweak interaction, including a neutral gauge boson [7]. The Weinberg-Salam theory in the SM describes this unification as a Yang-Mills field with a $U(1) \times SU(2)$ gauge group. The Higgs mechanism spontaneously breaks the symmetry of $U(1) \times SU(2)$ and imparts non-zero masses to the W^\pm and Z bosons [8]. This mechanism also predicts the existence of the Higgs boson, which was discovered at the Large Hadron Collider (LHC) in 2012 [1, 2]. Fermions also acquire mass through Yukawa interactions with the Higgs field.

The three generations of fermions in the SM are classified into quarks and leptons. Each generation of quarks consists of one quark with an electrical charge of $q = 2/3$ and one with $q = -1/3$. Subsequent

generations introduce heavier quarks, with the top quark (t) being particularly massive at approximately $173.2\text{GeV}/c^2$, in contrast to the light up quark (u) with a mass of 2.3MeV . All quarks interact through strong, weak, and electromagnetic forces because of their colour charge, weak isospin, and electric charge. Quarks can also mix with each other via the weak force. This mixing is described by the Cabibbo-Kobayashi-Maskawa (CKM) matrix [9]:

$$\begin{bmatrix} d' \\ s' \\ b' \end{bmatrix} = \begin{bmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{bmatrix} \begin{bmatrix} d \\ s \\ b \end{bmatrix} \quad (1.2)$$

On the left, the weak-interaction doublet partners of down-type quarks are shown. On the right, the CKM matrix is presented, accompanied by a vector depicting the mass eigenstates of the down-type quarks. The CKM matrix serves to characterise the likelihood of a transition from a flavour (i) quark to another flavour (j) quark. These transitions are directly proportional to the square magnitude of the corresponding matrix elements, denoted as $|V_{ij}|^2$ [6].

Leptons are classified into three generations of electrically charged particles: electrons (e 's), muons (μ 's), and taus (τ 's), and three generations of electrically neutral neutrinos: electron neutrino (ν_e), muon neutrino (ν_μ), and tau neutrino (ν_τ). It should be noted that neutrinos are assumed to be massless in the SM.

Historically, the muon was thought to be an excited state of the electron, suggesting that the decay $\mu \rightarrow e\gamma$ should be possible [10]. However, this decay was not observed, leading to the introduction of a new quantum number, called the lepton number, which was assumed to be a conserved quantity [11].

While the SM can accurately describe physical processes up to the scale of electroweak interactions, it relies on the measurement of 19 free parameters. The best estimates of these values at the time of writing are summarised in Table 1.1. The mass of the Higgs boson is not protected from large quantum corrections in the SM, potentially leading to a much higher mass than observed in experiments, a challenge known as the hierarchy problem [4]. Furthermore, the gauge hierarchy problem arises due to the lack of gauge coupling unification at high energy scales in the SM, a feature expected in many theories beyond the SM.

In addition to the theoretical challenges discussed above, the SM faces compelling experimental anomalies. In particular, two observations cannot be explained within the confines of the SM. One such phenomenon is the nature of dark matter, which remains elusive despite constituting a substantial fraction, approximately 25 %, of the total matter content of the universe. This enigmatic substance coexists with the 5 % SM baryonic matter and the enigmatic 70 % attributed to dark energy. The existence of dark matter is inferred from its gravitational effects, prominently observed in phenomena such as the measured velocity distributions of stars orbiting within galaxies [12]. Despite rigorous investigations, the fundamental properties and interactions of dark matter particles remain shrouded in mystery. On the other hand, dark energy is driving the continuous acceleration of the expansion of the universe, while its origin is unknown.

Furthermore, the phenomenon of neutrino oscillations presents a significant challenge to the SM [13]. The SM is assuming massless neutrinos, leading to their flavour eigenstates being equivalent to their mass eigenstates. This equivalence leads to the conclusion that the lepton-flavour is conserved. Consequently, in the SM, the rates for decays such as $\mu \rightarrow e\gamma$ or $\mu^+ \rightarrow e^+e^-e^+$ are zero.

However, experimental evidence has painted a different picture [14, 15, 16]. The neutrino oscillation,

observed by the neutrinos that transition between their different generations as they traverse space, signifies a key departure from the SM predictions. The observed neutrino oscillations are described by the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix [17, 18]. This discovery compels the recognition that neutrinos must possess non-zero masses, a revelation that has far-reaching implications for our understanding of particle physics and the nature of neutrinos themselves. Specifically, the observation of neutrino oscillations implies that the mass differences between the neutrinos must be non-zero, although the mass of the lightest neutrino could still be zero.

#	Symbol	Description	Value
1	m_e	Electron mass	0.511 MeV
2	m_μ	Muon mass	105.7 MeV
3	m_τ	Tau mass	1.78 GeV
4	m_u	Up quark mass	1.9 MeV
5	m_d	Down quark mass	4.4 MeV
6	m_s	Strange quark mass	87 MeV
7	m_c	Charm quark mass	1.32 GeV
8	m_b	Bottom quark mass	4.24 GeV
9	m_t	Top quark mass	173.5 GeV
10	θ_{12}	CKM 12-mixing angle	13.1°
11	θ_{23}	CKM 23-mixing angle	2.4°
12	θ_{13}	CKM 13-mixing angle	0.2°
13	λ	CKM CP violation phase	0.995
14	g_1 or g'	U(1) gauge coupling	0.357
15	g_2 or g	SU(2) gauge coupling	0.652
16	g_3 or g_s	SU(3) gauge coupling	1.221
17	θ_{QCD}	QCD vacuum angle	~ 0
18	v	Higgs VEV	246 GeV
19	m_H	Higgs mass	125.09(24) GeV

Table 1.1: List of the 19 free parameters of the SM, taken and adopted from [19].

Combining the SM with neutrino oscillations, the rate for $\mu \rightarrow e\gamma$ or $\mu^+ \rightarrow e^+e^-e^+$ is no longer zero, but of the order of 10^{-55} [20]. The challenges presented by the neutrino oscillation to the SM emphasise the need for new theoretical frameworks that can address these observations. Since there are only left-handed neutrinos in the SM, no interaction with the Higgs field is possible. Therefore, unlike other fermions, they cannot get their mass from the Higgs mechanism.

To address these issues, ongoing particle physics experiments actively seek deviations from SM predictions. A possibility of exploration is the search for CLFV decays, which violate the conservation of the lepton-flavour. An illustrative example of CLFV is the decay $\mu^+ \rightarrow e^+e^-e^+$, a process strictly forbidden in the SM. As mentioned above, it is even impossible to measure when extending the SM with neutrino oscillations. The detection of such a decay would serve as a clear indication of physics beyond the SM. The subsequent sections delve into a more detailed discussion of muon decays both within and beyond the SM.

1.2 Muon decays

To appreciate the significance of this decay, it is essential to provide context on μ decays within both the SM and beyond it.

1.2.1 Muon decays in the Standard Model

In the SM, μ decay exclusively occurs through the weak interaction. The leading-order process is the Michel decay (see Figure 1.2a):

$$\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu. \quad (1.3)$$




Figure 1.2: In Figure 1.2a the Michel decay of the muon into an electron (e), an electron neutrino (ν_e) and an anti muon neutrino ($\bar{\nu}_\mu$) is shown. In Figure 1.2b the muon decay by the exchange of an ν_e and ν_μ in the expansion of the SM with massive neutrinos is presented.

This process is mediated by the exchange of a virtual W boson. The branching fraction (BF) for this decay is close to 100% [6]. The next-to-leading-order process is the radiative decay:

$$\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \gamma. \quad (1.4)$$

For photon energies above 10 MeV, the BF is approximately $1.4(4) \times 10^{-2}$ [21]. Additionally, in the radiative decay, the emitted γ can convert into an electron-positron pair, giving rise to the internal conversion decay:

$$\mu^+ \rightarrow e^+ e^- e^+ \nu_e \bar{\nu}_\mu. \quad (1.5)$$

This process has a BF of approximately $3.4(4) \times 10^{-5}$ [6]. The decay $\mu^+ \rightarrow e^+ e^- e^+$ is strictly forbidden within the framework of the SM due to its violation of the conservation of the flavour of leptons. However, when the SM is extended to incorporate neutrino oscillations, as depicted in Figure 1.2b, this decay process can become possible. In this extended model with massive neutrinos, the BF for $\mu^+ \rightarrow e^+ e^- e^+$ is astonishingly minuscule, estimated to be on the order of 10^{-55} [22]. This almost infinitesimal probability makes it exceedingly improbable that such a decay could be observed within any reasonable time frame.

1.2.2 Muon decays beyond the Standard Model

In the realm beyond the SM, there exist theoretical frameworks that predict the occurrence of CLFV decays at observable BFs. One class of such theories is supersymmetry (SUSY) [20], which posits the existence of superpartners for every known particle in the SM. These superpartners have different spin values compared to their SM particles. In particular, each fermion in the SM has a supersymmetric boson partner and vice versa. Although these superpartners have yet to be observed, their potential discovery holds promise for resolving some of the lingering questions in particle physics. Within the framework of SUSY, the decay $\mu^+ \rightarrow e^+e^-e^+$ could materialise through a quantum loop involving SUSY particles, as illustrated in Figure 1.3 [23].




Figure 1.3: μ decay via SUSY particles.

The Scotogenic Model, initially proposed by Ernest Ma in 2006 [24], is another theoretical framework designed to address two key puzzles in particle physics: neutrino mass generation and the identification of a dark-matter candidate, all while remaining consistent with the SM. This model extends the SM by introducing three right-handed neutrinos (N_i) and a scalar $SU(2)_L \times U(1)_Y$ doublet. Here, i represents the isospin, L the lepton number, and Y the hypercharge. For these right-handed neutrinos, the quantum numbers for hypercharge and isospin are zero. Note that so far only left-handed neutrinos and right-handed antineutrinos have been observed in nature. Through this extension, it achieves the generation of the SM neutrino masses via a radiative seesaw mechanism that does not necessitate the existence of very massive sterile neutrinos.




Figure 1.4: Possible production channels for the decay $\mu^+ \rightarrow e^+e^-e^+$ considering the Scotogenic Model (cf. [25]) by the loop of the scalar doublet σ^+ .

A feature of the Scotogenic Model is the incorporation of sterile right-handed neutrinos (N_i). These neutrinos, being sterile, have the potential to violate lepton-flavour conservation. Possible Feynman diagrams that illustrate how these sterile neutrinos (N_i) can contribute to CLFV in the context of decay $\mu^+ \rightarrow e^+e^-e^+$ are depicted in Figure 1.4.

In addition to these theories, there exist numerous other theories in which the decay $\mu^+ \rightarrow e^+e^-e^+$ may occur. For instance, in left-right symmetric models of the weak interaction, parity is conserved above the scale of symmetry breaking. This conservation leads to the existence of right-handed neutrinos and neutrino masses, among other phenomena. These models can be implemented using a Higgs triplet composed of a neutral scalar (Δ^0), a singly charged scalar (Δ^+), and a doubly charged scalar (Δ^{++}) [26, 27, 28, 29, 30, 31, 32, 33, 34, 35]. Consequently, the decay $\mu^+ \rightarrow e^+e^-e^+$ can occur at the tree level. Furthermore, models that incorporate an extended gauge sector or extra dimensions introduce new heavy-vector bosons, such as Z' [36, 37, 38]. Through flavour-off-diagonal couplings, a Z' boson can mediate the decays $\mu^+ \rightarrow e^+e^-e^+$.

All the theoretical frameworks mentioned above are subject to stringent experimental constraints that restrict their predicted rates of CLFV decays. For instance, the combined Mu to E Gamma (MEG) and MEG II experiment have established the most stringent limit on the CLFV decay $\mu^+ \rightarrow e^+\gamma$, reporting a BF of $\mathcal{B}(\mu^+ \rightarrow e^+\gamma) < 3.1 \times 10^{-13}$ at 90% confidence level (CL) [39]. Similarly, the SINDRUM experiment has placed constraints on the CLFV decay $\mu^+ \rightarrow e^+e^-e^+$, setting a BF limit of $\mathcal{B}(\mu^+ \rightarrow e^+e^-e^+) < 1.0 \times 10^{-12}$ at 90% CL [40]. Furthermore, the SINDRUM II experiment has established limits on the CLFV process $\mu^- \text{Au} \rightarrow e^- \text{Au}$ with a BF of $\mathcal{B}(\mu^- \text{Au} \rightarrow e^- \text{Au}) < 7 \times 10^{-13}$ at 90% CL [41]. In the next section, a more detailed overview of the MEG, MEG II and SINDRUM experiments is provided, based on the paper "Introduction to Charge Lepton-Flavour Violation" [42].

1.3 Charged lepton-flavour violation experiments

The search for CLFV processes is crucial in the quest to discover new physics beyond the SM, given the multitude of different theories that predict varying branching ratios for these processes. Numerous experiments have been conducted in the past to explore various CLFV decays, including rare μ and τ decays, as well as muon-to-electron conversion in nuclei.

Muon-based experiments have played a central role in these efforts due to several advantageous characteristics. Muons can be produced in substantial quantities, possess a relatively long lifetime of approximately 2.2 μs , and exhibit purely leptonic decay modes. Currently, different experiments are being constructed and / or data is taken to search for different LFV decay modes of the μ .

$\mu^+ \rightarrow e^+\gamma$ is, an important process under investigation by the MEG (II) experiment [43, 44]. $\mu^+ N \rightarrow e^+ N$, which is the focus of experiments such as COMET [45] and Mu2e [46]. Meanwhile, the Mu3e experiment, currently under construction, is specifically designed to investigate the decay mode $\mu^+ \rightarrow e^+e^-e^+$.

All these experiments follow a common approach: muons interact or stop on a target, and the resulting decay particles are carefully measured. As mentioned above, the current limit for $\mu^+ \rightarrow e^+e^-e^+$ was established by the SINDRUM experiment [40], and its detector concepts are discussed in this context. The diversity of these CLFV experiments and their associated processes allows for a comprehensive exploration of CLFV and the potential discovery of new phenomena in physics.

1.3.1 SINDRUM

The most recent search for the $\mu^+ \rightarrow e^+e^-e^+$ decay was carried out in 1988 by the SINDRUM experiment at the Paul Scherrer Institute (PSI). In this experiment, no events were observed within the signal region. Consequently, the researchers established an upper limit on the branching ratio

of $\mathcal{B}(\mu^+ \rightarrow e^+e^-e^+) < 1.0 \times 10^{-12}$ at 90 % CL [40]. The experimental setup used in SINDRUM is illustrated in Figure 1.5. The experiment involved stopping a muon beam at a rate of $5 \times 10^6 \mu/s$ on a hollow double cone target. The resulting electrons were detected using five concentric multiwire proportional chambers, with three of them additionally equipped with cathode strips to determine the z coordinate. In addition, an array of scintillation counters was used to measure the timing information of the decay electrons. These detectors operated within a magnetic field of 0.33 T.




Figure 1.5: Sketch of the SINDRUM experiment, taken from [40]. B is the incoming muon beam, S is the beam transport solenoid, T is the muon stopping target, C is the low mass multiwire proportional chamber, H is cylindrical scintillator hodoscope, M is the magnetic coil, L are the light guides, A is the preamplifier and P are the photomultipliers.

Following its initial investigations, SINDRUM was reconfigured to explore neutrinoless $\mu^- \rightarrow e^-$ conversion in muonic atoms. In this context, the experiment established a decay limit with a branching ratio of $\mathcal{B} < 7 \times 10^{-13}$ at 90 % CL using gold as a target [41]. In contrast to the work presented in this thesis, the DAQ system of the SINDRUM and SINDRUM II experiments used a trigger system to read the detectors. To increase the limit set by the SINDRUM experiment, the DAQ system of a potential preprocessor experiment searching for $\mu^+ \rightarrow e^-e^+e^+$ needs to reconstruct all events online to be able to handle the increased accidental background.

1.3.2 MEG and MEG II

The most recent measurement of CLFV was conducted by the MEG and MEG II experiment, which searched for $\mu^+ \rightarrow e^+\gamma$. The MEG experiment operated at PSI from 2009 to 2013, while the MEG II experiment is currently taking data. During the period 2009 to 2013, a total of 7.5×10^{14} muons were stopped in the experiment. However, no excess events were detected, resulting in a new upper limit on this decay of $\mathcal{B}(\mu^+ \rightarrow e^+\gamma) < 4.2 \times 10^{-13}$ at 90 % CL [43].

The MEG experiment used a unique setup: muons were stopped by a thin stopping target, consisting

of a single polyethylene sheet set at an angle of 70° to the beam, which maximised stopping efficiency while minimising multiple scattering and bremsstrahlung. Positrons resulting from muon decays were detected by a drift chamber system comprising 16 independent modules, covering angles from 191.25° to 378.75° [47]. The timing of the positrons was precisely measured using scintillating bars on the timing counters. Photons, integral to the objectives of the experiment, were detected by a liquid xenon calorimeter, specially designed to fully capture a 52.83 MeV photon shower. All components of the experiment, except the liquid xenon calorimeter, were located within a superconducting magnet that generated a field of approximately 1.27 T in the centre. To generate a gradient magnetic field, six coils with different radii are applied. In addition, two additional coils were used to compensate for the residual field around the liquid xenon calorimeter. The experimental setup is depicted in Figure 1.6.




Figure 1.6: Sketch of a simulated event in the MEG experiment, taken from [43].

The MEG experiment has undergone significant improvements that have led to a new experiment called MEG II [44]. These enhancements include the conversion of the drift chambers into a single cylindrical drift volume, improvements in the timing detector using scintillating tiles, replacing the photomultiplier tubes (PMTs) on the front face of the liquid xenon calorimeter with silicon photomultipliers (SiPMs) and the addition of a new radiative decay counter to suppress the background [44]. Resulting in a sensitivity of 6×10^{-14} , MEG II has been actively collecting data since 2021.

The MEG II experiment measured in 2021 a branching ratio limit of $\mathcal{B}(\mu^+ \rightarrow e^+\gamma) < 7.5 \times 10^{-13}$ at 90% CL [48]. Both experiments have produced a combined branching ratio limit of $\mathcal{B}(\mu^+ \rightarrow e^+\gamma) < 3.1 \times 10^{-13}$ at 90% CL. Again, in contrast to the work presented in this thesis, both the MEG and the MEG II experiments use a triggered DAQ system.

Principles of Particle Detectors

In particle physics experiments, tracking detectors play a crucial role, as they enable the determination of both the momentum and direction of charged particles. The ideal tracking detector should minimise particle interaction, avoiding significant deflection, to reconstruct the particle's path accurately based on its interactions with the detector. This requires a minimal material budget.

In the Mu3e experiment, muons decay at rest, producing low-momentum electrons and positrons. Hence, minimising detector material is particularly vital to mitigate multiple-scattering effects, which become more pronounced at lower momenta. Advancements in semiconductor technology enable the construction of extremely thin detectors. In the following sections, the fundamentals of silicon detectors (Section 2.1), tracking systems (Section 2.2), and multiple scattering (Section 2.3), based on Thomson [49], will be explained.

Subsequently, the unique pixel detector employed in the Mu3e experiment will be explored (Section 2.4). Additionally, an introduction to scintillating detectors and an example using scintillating fibres in Section 2.5 will be given, with reference to Knoll's work [50]. Finally, the principles of the readout electronics used in the Mu3e experiment will be elucidated.

2.1 Silicon detectors




Figure 2.1: Sketch of a particle travelling through a silicon pixel detector. When an external voltage is applied, the whole area between the regions doped with p and n becomes active.

In Figure 2.1, a sketch of a silicon detector is provided. This detector comprises different regions doped with implants of type p or n, as well as n^+ doped implants. P-type silicon can be achieved by implanting a phosphorous nucleus that has an extra electron. By doping with a boron nucleus, one can produce n-type silicon because boron has one electron less than silicon. When oppositely doped regions are joined, they form a pn-junction. Inside this junction, charge carriers have the ability to diffuse into the opposing region, where they recombine with their counterparts, resulting in the creation of a charge-free region around the pn-junction, known as the depletion zone. When an external reverse bias voltage is applied, the depletion zone expands and the detector can be depleted, allowing the entire volume to become active.

When a charged particle traverses the detector, it generates free electron-hole pairs. The application of voltage to the pn-junction enhances the motion of these particles towards the electric field, resulting in the generation of an induction signal that can be processed using amplifiers and comparators. Incorporation of small p-doped implants within the detector allows for precise measurement of spatial information. When these sensors are extended in one dimension, they are called strip detectors. If they possess fine segmentation in two dimensions, then they are known as pixel detectors.

2.2 Principle of tracking detectors




Figure 2.2: Sketch of a particle travelling through a tracking detector and important variables of the particle track in the x-y plane.

In Figure 2.2, a schematic representation of the trajectory of a particle is shown within a tracking detector, located in the x-y plane. To determine the momentum of a particle traversing a barrel-shaped tracking detector, one needs to apply a magnetic field perpendicular to the particle's path. Consequently, the Lorentz force equals the centripetal force, yielding the following momentum equation:

$$\frac{m \cdot v^2}{R} = e \cdot v \cdot B \rightarrow p = e \cdot B \cdot R. \quad (2.1)$$

Here, m denotes mass, v represents velocity, e signifies electric charge, and p stands for the particle's momentum. The variable B denotes the magnetic field applied perpendicular to the particle's track, and R represents the radius of the track (cyclotron radius). Note that the left part of this equation is only valid in non-relativistic scenarios. The tracking layer typically comprises multiple silicon

detectors. To assess the resolution of the momentum, the sagitta s is introduced, which can be derived for small ϕ as follows (see right part of Figure 2.2:

$$s = R - R \cdot \cos\left(\frac{\phi}{2}\right) \approx R \cdot \frac{\phi^2}{8} = R \cdot \frac{L^2}{8 \cdot R^2}. \quad (2.2)$$

Furthermore, the momentum resolution is expressed as follows:

$$\frac{\Delta p}{p} = \frac{\Delta R}{R} = \frac{L^2}{8 \cdot R \cdot s} \cdot \frac{\Delta s}{s} \quad (2.3)$$

To achieve a superior momentum resolution, it is essential to have a long path (L), a strong magnetic field (B), and precise sagitta measurements. However, it is important to consider that in the Mu3e experiment, where the decay products possess lower momenta, multiple scattering effects must be taken into account. Furthermore, sagitta is not small in comparison to the radius in the case of the Mu3e experiment.

2.3 Multiple Scattering




Figure 2.3: Particle passing through matter, taken from [6].

In Figure 2.3, the trajectory of a particle traversing matter is shown. The path of the particle is altered because of scattering interactions with the Coulomb fields of nuclei, leading to an offset from its original course and an alteration in its angle. When the detector is extremely thin, the offset can be ignored, focussing solely on the change in angle. The root mean square (RMS) of this angle, denoted as Θ_{MS} , is accurately approximated by the Highland equation [6]:

$$\Theta_{MS} = \frac{13.6 \text{ MeV}}{p \beta c} q \sqrt{\frac{x}{X_0}} \left[1 + 0.038 \log\left(\frac{x}{X_0}\right) \right]. \quad (2.4)$$

Here, p represents momentum, βc signifies the velocity of a particle with a charge q in units of electron charge, x is the path length within the material, and X_0 is the radiation length of the material. Given that the electron has a fixed charge, the only parameter in the equation that can be modified is the radiation length, x . To achieve a precise momentum resolution, it is imperative to minimise the magnitude of the multiple scattering angle, imposing stringent constraints on the material budget.




Figure 2.4: Tracking in the scattering dominated regime, taken from [51].

As Equation 2.4 implies, the deviation is more pronounced for particles with lower momenta and for thicker detector layers. In the context of Mu3e, where particle momenta are in the range of a few tens of MeV, minimising material within the active detector volume is paramount for accurate momentum measurements. Apart from multiple scattering, the spatial uncertainty of individual measurement points along the track also contributes to momentum resolution as discussed above. Figure 2.4 illustrates that these spatial uncertainties can be disregarded when they are significantly smaller than the uncertainties introduced by scattering.

2.4 High Voltage Monolithic Active Pixel Sensors

In Section 2.1, the fundamental principles of silicon detectors were discussed. To extract the signal from these detectors, an amplification circuit can be integrated into the pixels, resulting in Active Pixel Sensors (APS). Alternatively, the entire signal readout can be placed on the chip, leading to Monolithic Active Pixel Sensors (MAPS). These chips cannot be fully depleted, as the required voltage would interfere with the circuits. Thus, they collect the signal via diffusion, which has the disadvantage of a relatively long charge collection time, typically on the order of microseconds, compared to the typical drift time in a bias field of the order of nanoseconds.

In Figure 2.5, a schematic of a High-Voltage Monolithic Active Pixel Sensor (HV-MAPS) is shown [52]. These sensors combine the monolithic characteristics of MAPS with the fast charge collection capability achieved by drift. This is accomplished by using a deep n-well located in the p-substrate, which forms a diode. Pixel electronics can be implemented inside the n-well where they are isolated from the depletion voltage. Due to the use of the deep n-well a high reverse bias voltage can be applied to the substrate, creating a large depletion zone. These deep n-wells are available in commercial high voltage (HV) complementary metal-oxide-semiconductor (CMOS) processes.




Figure 2.5: Sketch of a HV-MAPS sensor [52].

Furthermore, the sensor can be thinned from the back. This technology enables fast charge collection and guarantees a low material budget, ensuring good momentum resolution. The HV-MAPS pixel sensor used in the Mu3e experiment is called MuPix and in Section 4.3, a more specific discussion will give detailed explanations regarding this sensor.

2.5 Scintillating detectors

Scintillators are materials that produce light when they come into contact with ionising radiation. They come in different types and forms. For example, organic scintillators are made up of an activator and a surrounding material. When ionising radiation interacts with these scintillators, the energy of the particles is transformed into light through a fluorescence process that occurs within the activator material. In the case of a plastic scintillator, the main source of fluorescence, known as a fluor, is embedded within a solid polymer matrix called the base.




Figure 2.6: Sketch of light travelling inside a fibre scintillator.

One potential detector design involves shaping these plastic scintillators into fine fibres, capable of efficiently transporting emitted light over considerable distances through total internal reflection. This configuration is shown in Figure 2.6. In this setup, the scintillator is located in the core of the fibre, which is surrounded by a cladding material. The light produced within the core reaches the core-cladding boundary, potentially at an angle that exceeds the critical angle for total reflection. This allows light to propagate along the fibre. At the termination points of the fibres, photomultipliers can be integrated to enhance the signal. Photomultipliers achieve this by initially converting the incoming photons into electrons that are subsequently amplified.

Readout Electronics

To enable the readout of various detectors within a particle physics experiment, a set of essential readout electronics must be in place. In the subsequent sections, the fundamental concepts underlying these components are explained. First, the intricacies of FPGAs, which serve as the primary building blocks in contemporary DAQ systems, will be explored. Following this, the principles of data transmission and the implementation of transceivers utilised within FPGAs will be explored. Lastly, an overview of the principles governing Peripheral Component Interconnect Express (PCIe), a commonly used interface in readout systems, will be provided. These sections are based on previous work by the author [53].

3.1 Field Programmable Gate Arrays (FPGAs)

FPGAs represent integrated circuits endowed with reprogrammable hardware capabilities. These devices enable users to modify their functionality after manufacturing using reconfigurable interconnect lines and Logic Elements (LEs). These LEs are constructed from Look-Up Tables (LUTs) and registers, with interconnect lines located between them.




Figure 3.1: Structure of an FPGA.

In the context of this work, the hardware description language Very High Speed Integrated Circuit Hardware Description Language (VHDL) is harnessed to programme the FPGAs. Figure 3.1 offers a simplified depiction of a FPGA. This architectural marvel comprises a core consisting of LEs

interconnected through a global routing network, surrounded by a perimeter of input/output (I/O) pads. Additionally, FPGAs offer the option of using memory blocks and Hard Intellectual Property Core (IP) cores, which are predefined blocks within the device, thoughtfully designed and provided by the FPGA manufacturer. These IP cores are customised for specific tasks, such as high-speed serialisation or deserialisation.

With the aid of VHDL, the data flow and logic of the FPGA can be programmed. Initially, this design undergoes synthesis, resulting in the generation of a netlist. Subsequently, the fitting tool maps this netlist to the LEs and interconnects within the FPGA. In particular, this process takes into account the desired operating frequency. Ultimately, a configuration file is prepared for upload to the device. In the case of Intel (formerly Altera) FPGAs, this entire procedure is executed within the Intel Quartus Prime software [54].

Logic Elements

In Figure 3.2, an illustration of a LE with four inputs and one output is presented. The output can be connected directly to a LUT or driven by a flip-flop. A LUT is essentially a straightforward array that contains values that can be accessed by simple index operations. On the other hand, a flip-flop serves the purpose of information storage. Specifically, the D-flip-flop delivers data to the output when a clock edge is detected on the clock input. In the absence of a clock signal, changes in the input do not affect the output, leading to the retention of information within the flip-flop until a clock signal is detected. By routing multiple LUTs together, various Boolean functions can be effectively expressed.




Figure 3.2: Illustration of a simple LE containing a four input LUT and a D-flip-flop.

The FPGAs utilised in the Mu3e experiment belong to the category of Static Random Access Memory (SRAM)-based FPGAs. These FPGAs store the configuration of the LUT within SRAM cells. Furthermore, Intel FPGAs have been chosen for their favourable balance between fast and slow I/O, a sufficient quantity of LEs, and the availability of evaluation boards, some of which can be adapted for use in the final detector setup.

Routing

To establish connections between different LEs, input and output pins can be connected to the routing wires. Figure 3.3a provides an illustration of such a connection. Each of these routing wires extends to a switch box positioned at the intersections. Figure 3.3b presents these switch boxes, which are programmable to facilitate the interconnection of multiple LEs.

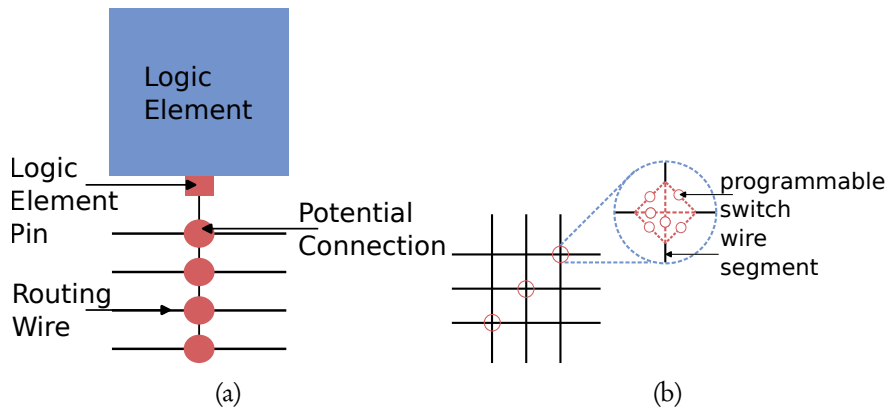


Figure 3.3: Figure 3.3a does show the connection of a LE to a routing wire. Figure 3.3b shows a sketch of a programmable switch sitting on the intersections of different routing wires.

3.2 Data Transmission in FPGAs

In the pursuit of identifying rare events, the study of numerous decays becomes imperative. Consequently, an immense volume of data requires processing. To accommodate such a data rate, high-speed data links and FPGAs capable of real-time data analysis are indispensable.

In this chapter, the fundamentals of data transmission are elucidated in Section 3.3. Subsequently, Section 3.4 delves into electrical data transmission, while Section 3.5 explores the concept of serial data links. Furthermore, the general implementation of high-speed data transmission and the intricacies of transmission within FPGAs will be discussed in Section 3.6. Optical data transmission is covered in Section 3.7. Finally, Section 3.8 describes a test procedure for assessing the performance of such data links followed by the basics of Peripheral Component Interconnect Express (PCIe) and Direct Memory Access (DMA). The key content and images in these sections have been sourced from references such as [55], [56], and [53].

3.3 Basics of data transmission

Electrodynamics

In the field of electrodynamics, the fundamental theory that underpins electrical and optical data transmission is based on classical electrodynamics, as originally formulated by Maxwell. This theory can be succinctly summarised through four key equations:

1. Gauss's Law for Electric Fields:

$$\vec{\nabla} \cdot \vec{D} = \rho \quad (3.1)$$

2. Faraday's Law of Electromagnetic Induction:

$$\vec{\nabla} \times \vec{H} - \frac{\partial \vec{D}}{\partial t} = \vec{j} \quad (3.2)$$

3. Faraday's Law of Magnetic Fields:

$$\vec{\nabla} \times \vec{E} + \frac{\partial \vec{B}}{\partial t} = 0 \quad (3.3)$$

4. Gauss's Law for Magnetic Fields:

$$\vec{\nabla} \cdot \vec{B} = 0 \quad (3.4)$$

Here, \vec{E} represents the electric field, \vec{B} the magnetic field, \vec{D} the electric displacement, \vec{H} the magnetising field, ρ the charge density and \vec{j} the current density. These equations describe the fundamental principles that govern electromagnetic phenomena. Additionally, there are relationships between these fields:

$$\vec{D} = \epsilon_r \epsilon_0 \vec{E} \quad (3.5)$$

$$\vec{B} = \mu_r \mu_0 \vec{H} \quad (3.6)$$

Here, ϵ_0 is the vacuum permittivity, μ_0 is the vacuum permeability, ϵ_r and μ_r are the corresponding values depending on the medium. Furthermore, the conductivity σ is determined by Ohm's law:

$$\vec{j} = \sigma \vec{E}. \quad (3.7)$$

Plane Waves in Materials

Following the calculations of [57], the behaviour of electromagnetic waves in different materials will be examined. Specifically, the wave equations governing their propagation in non-conducting and conducting media will be investigated in the following. In a non-conducting material (where conductivity, σ , and current density, \vec{j} , are both zero), the wave equation takes the form:

$$\vec{\nabla}^2 \vec{E} = \epsilon \mu \frac{\partial^2 \vec{E}}{\partial t^2}. \quad (3.8)$$

In a conducting medium (with σ and $\vec{j} \neq 0$), the wave equation becomes:

$$\vec{\nabla}^2 \vec{E} = \epsilon \mu \frac{\partial^2 \vec{E}}{\partial t^2} + \sigma \mu \frac{\partial \vec{E}}{\partial t}. \quad (3.9)$$

By introducing an ansatz for the electric field \vec{E} being a plane wave, one obtains:

$$\vec{E}(x, t) = \vec{E}_0 e^{i\omega t - ikx}, \quad (3.10)$$

leading to:

$$k^2 = -i\omega\mu\sigma + \omega^2\mu\epsilon. \quad (3.11)$$

One then expresses the wave vector \vec{k} as the sum of real and imaginary parts:

$$\vec{k} = \alpha - i\beta, \quad (3.12)$$

yielding the solution for the plane wave as:

$$\vec{E}(x, t) = \vec{E}_0 e^{i\omega t - \alpha x} e^{\beta x}. \quad (3.13)$$

Here, α determines the wave's length, and β represents the attenuation factor. The dispersion relation connecting the wave vector and frequency is given by:

$$k^2 = \omega^2 \mu \epsilon. \quad (3.14)$$

Solving for α and β , one finds:

$$\alpha = \omega \sqrt{\mu \epsilon} \left(\frac{1}{2} + \frac{1}{2} \sqrt{1 + \frac{\sigma^2}{\omega^2 \epsilon^2}} \right)^{\frac{1}{2}}, \quad (3.15)$$

$$\beta = \frac{\omega \mu \sigma}{2\alpha}. \quad (3.16)$$

This analysis also applies to real dielectric materials where ϵ is a complex number, leading to dispersion and attenuation equations with modified terms.

Dielectrics in High-Speed Data Links

In high-speed electrical data links, dielectric materials play a critical role in reducing power loss. The power loss inside the dielectric is proportional to:

$$P_{\text{loss}} \propto \omega \cdot \tan \delta, \quad (3.17)$$

where $\tan \delta$ represents the loss tangent:

$$\tan \delta = \frac{\sigma + \omega \epsilon''}{\omega \epsilon'}. \quad (3.18)$$

The critical frequency, ω_n , which delineates conducting and insulating modes within a material, is given by:

$$\omega_n = \frac{\sigma}{\epsilon}. \quad (3.19)$$

For many good insulators, conductivity is roughly proportional to frequency. Consequently, dielectric losses significantly impact high-speed data transmission in electrical links.

3.4 Electrical data transmission and Telegrapher's Equations

In the context of high-speed data transmission, electrical links play a crucial role. By modelling an electrical transmission line as an infinite cascade set of two-port systems, it is possible to derive the telegrapher's equations. These systems are constructed from segments with a series impedance (Z), which consists of a resistance (R) and an inductance (L), as well as a parallel admittance (y). The admittance is connected to the ground and is composed of a capacitance (C) and a shunt conductance (G). This configuration is illustrated in Figure 3.4. With the telegrapher's equation, it is possible to determine the wave impedance (Z_0), which is given by:

$$Z_0 = \sqrt{\frac{R + i\omega L}{G + i\omega C}}. \quad (3.20)$$

In the case of a lossless line (i.e., $R = G = 0$), the wave impedance (Z_0) is independent of the frequency (ω).




Figure 3.4: Infinite Cascade model for deriving the telegrapher's equation, taken from [58].

However, in real-world transmission lines, resistance (R) and conductance (G) are non-zero. There are several reasons for this, including dielectric losses and the skin effect. The coefficient β determines the distance a wave travels within a conductor. The inverse of β is known as the depth of the skin (δ), which represents the distance over which the amplitude of the wave is reduced by a factor of $1/e$. This skin effect restricts the area over which a high-speed wave can travel inside a conductor to its outer surface. Consequently, these effects can cause attenuation and dispersion of high-speed data signals, leading to inter-symbol interference.

To ensure successful high-speed data transmission, it is essential to address and correct these effects. The following section explores serial data links in more detail.

3.5 Serial data links for Inter-device Communication

Efficient communication between devices situated on different circuit boards often requires reducing the number of connectors on the boards. Serial data links offer a solution to this challenge by requiring only one line for communication, as opposed to the multiple lines used by parallel buses. However, serial data links must operate at higher frequencies than parallel buses to transmit the same volume of data.

Typically, high-speed data links are constructed using multiple layers, each serving a specific purpose. At a minimum, these links include the physical layer (PHY) and the data link layer. The PHY is responsible for mapping information to physical observables and facilitating the transfer of data to the data link layer. It also ensures electrical compatibility between devices, a crucial aspect of successful communication.

The data link layer plays a vital role in ensuring reliable data transmission and improving signal integrity. It accomplishes this by encoding data, aligning them for proper transmission, and performing clock recovery when necessary. These measures are essential for maintaining the integrity of the data being transmitted.

Beyond these foundational layers, additional layers can be added for error correction or to support protocol-dependent features. These additional layers are essential to address specific communication requirements and ensure the robustness of the data transfer process. In the following sections, the physical and data link layers of high-speed data links will be further explored and discussed in detail.

Physical Layer and non-return-to-zero Coding

In the context of electrical communication, data need to be translated into physical observables. Most often, electrical links use two voltage levels to represent the logical states "1" and "0". This is achieved through non-return-to-zero (NRZ) binary coding, which is a common technique for encoding data for transmission.




Figure 3.5: Illustration of the NRZ code for data transmission.

Figure 3.5 provides an illustration of the NRZ code, showing the two voltage levels applied to the cable to encode 1 and 0 states. High-speed links often implement this encoding using a differential signal with low common mode. One standard for implementing this technology is low-voltage differential signaling (LVDS) [59]. In Intel FPGAs, a proprietary standard called pseudo current mode logic (PCML) [60] is used for very high-speed links, typically in the range of gigabits per second.

To facilitate data transmission, a fast clock is used to serialise the data on the transmitter side. This clock must also be recovered on the receiver side to accurately capture the bits. For clock recovery, a circuit known as a Clock and Data Recovery (CDR) system is used, which typically includes phase interpolation and phase-locked loops (PLLs). Additionally, there are methods where the clock is transferred from the transmitter, and the data is recovered using this clock. In the Mu3e experiment, the second option was chosen.

Phase-Locked Loops

A PLL is a crucial component in clock and data recovery. It generates an output signal whose phase is related to the phase of an input signal. Figure 3.6 illustrates the basic principle of a PLL.

The operation of a PLL involves the following components:

1. **phase frequency detector (PFD)**: It compares the phase of the input signal (reference frequency) with the phase of the feedback signal (output frequency divided by an integer N) and generates a control voltage (U) based on their phase difference.
2. **loop filter (LF)**: The control voltage (U) passes through a LF, which smooths out rapid changes and generates a voltage that adjusts the voltage controlled oscillator (VCO).
3. **voltage controlled oscillator (VCO)**: The VCO generates the output frequency (f_{out}) based on the adjusted control voltage. The VCO frequency is controlled by the control voltage.
4. **frequency divider (N)**: The output frequency is divided by an integer value N , which determines the final output frequency (f_{out}).

In steady state, the PLL generates an output frequency of $f_{\text{out}} = f_{\text{ref}} \cdot N$. The value of N , which is an integer, allows f_{out} to be changed only in the steps of f_{ref} . This limitation can result in very small




Figure 3.6: Illustration of the principle of a PLL.

reference frequencies (f_{ref}) if fine-grained adjustments are needed for f_{out} . A possible approach to this problem involves the incorporation of a divider between f_{ref} and the PFD. Since the frequency divider on the feedback path divides by N and the reference input divider divides by M , the PLL is able to multiply f_{ref} by $\frac{N}{M}$.

Data Link Layer and Encoding Schemes

The data link layer plays a critical role in data transmission by encoding data and managing data protocols, following [61, 58]. Several encoding schemes are commonly used in high-speed data links, including 8b/10b, 64b/66b, and 128b/130b. In this section, the focus is on the 8b/10b encoding scheme and its key features. The 8b/10b encoding scheme [62] converts 8 bits of user data to 10 bits for transmission on the communication line. This conversion allows for the exclusion of some possible 1024-bit patterns, which imposes a run-length limit of five consecutive equal bits. Additionally, the disparity, defined as the sum of one and zero in a word, should not exceed two in this encoding scheme. By implementing these constraints and decoding some 256 possible 8-bit words in two different ways, the bit sequence is designed to contain an equal number of ones and zeros, ensuring direct current (DC) balancing. Furthermore, the running disparity (RD) is continuously measured in the data stream and can only have values of ± 1 .

Previous RD	Next RD	Disparity choices	Disparity chosen
-1	-1	0	0
-1	+1	± 2	+2
+1	+1	0	0
+1	-1	± 2	-2

Table 3.1: Choice of disparity, taken from [63].

This mechanism helps to keep the disparity within the limits specified in Table 3.1. Moreover, the use of the 8b/10b encoding enables the detection of errors caused by a single bit flip.

Control Symbols and DC Balance

In the 8b/10b encoding, there are special control symbols, as detailed in Table 3.2. Among these symbols, K.28.1, K.28.5, and K.28.7 are referred to as "comma symbols". These comma symbols are used to identify word boundaries (finding the alignment of the 8b/10b codes within a bit stream), helping the receiver correctly interpret the data.

Name	8 bit	10 bit	10 bit
	HEX	RD = -1	RD = +1
K.28.0	1C	001111 0100	110000 1011
K.28.1 †	3C	001111 1001	110000 0110
K.28.2	5C	001111 0101	110000 1010
K.28.3	7C	001111 0011	110000 1100
K.28.4	9C	001111 0010	110000 1101
K.28.5 †	BC	001111 1010	110000 0101
K.28.6	DC	001111 0110	110000 1001
K.28.7 †	FC	001111 1000	110000 0111
K.23.7	F7	111010 1000	000101 0111
K.27.7	FB	110110 1000	001001 0111
K.29.7	FD	101110 1000	010001 0111
K.30.7	FE	011110 1000	100001 0111

Table 3.2: List of all control symbols in 8b/10b encoding. All symbols with † are called comma symbols and they are used to align the 8b/10b data stream. The bit order is from the least significant bit to the most significant bit taken from [63].

3.6 High speed links in FPGAs

In this section, the concept of a high-speed link is implemented in an FPGA, focussing on the example of the Intel Arria 10 FPGA, which is a key component in this thesis. Specifically, the examination of the Arria 10 GX Device Transceiver (TX) and the Intel Arria 10 TX PHY. The information provided here is based on the Intel Arria 10 TX PHY User Guide [64].

GX Transceiver

The Intel Arria 10 FPGA is equipped with a substantial number of GX TX channels, with the capability to handle high-speed digital signals and CDR. Each of these channels can transmit data at speeds up to 17.4 Gbps. These channels are organised into banks with up to eight TX in each bank. These banks are strategically placed in the left and right periphery of the FPGA device. In Figure 3.7, an example of one such TX bank containing three channels can be seen.

The overall clock distribution to the three channels is done via the main clock generation block (CGB). Each channel has integrated features, including a PLL and a local CGB. The local CGB divides and distributes the clock to the Physical Coding Sublayer (PCS) and Physical Medium Attachment (PMA) blocks. This is done in non-bonded configuration mode, where the channels are not related to each other. In this mode, the feedback path is local to the PLL.

The PCS plays a crucial role in encoding, decoding, and management of alignment markers within the data stream. Alignment is necessary because the TXs lack knowledge about the word boundaries within the stream. The PCS serves as an interface between the PMA and the FPGA Core Fabric. Additionally, each channel is connected to the Clock Distribution Network and the FPGA Core Fabric. The FPGA features fractional PLLs (fPLLs) that can generate lower clock frequencies, especially for data rates below 12.5 Gbps. There are also Advanced Transmit PLLs (ATXs) that cover the entire

range of supported data rates.

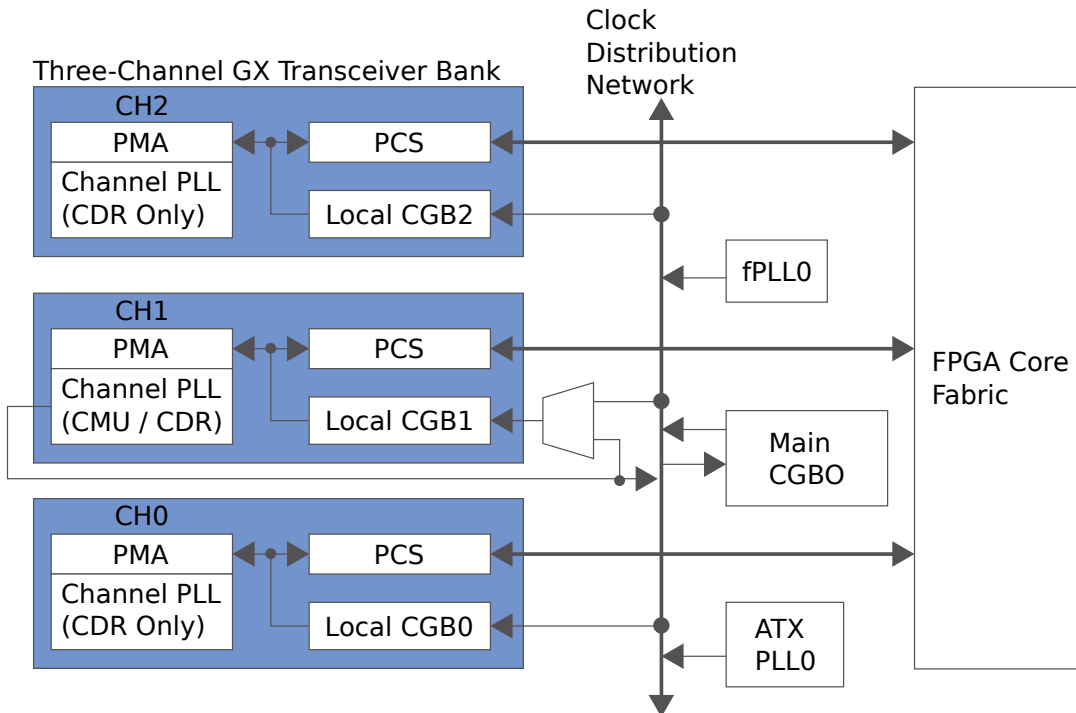


Figure 3.7: High level overview of the TX bank architecture of a bank made out of 3 TX channels, adopted from [64].

One notable advantage of the fPLLs is their ability to generate frequencies that are a rational fraction N/M of the reference frequency. This is achieved by first multiplying the reference frequency by an integer before feeding it into the PFD. In the context of the Mu3e experiment, all TX PLLs use the global 125 MHz clocks, provided by the clock and reset system (as discussed in Section 5.2), as their reference clocks. Since the Front-end board (FEB)¹ TXs operate at a data rate of 6250 Mbps, and the data packages are divided into 32-bit sections with 8b/10b encoding, the transceiver clock must be precisely 156.25 MHz. The use of fPLLs facilitates the achievement of this specific clock frequency requirement.

Intel Arria 10 Transceiver PHY Layer

The Intel Arria 10 TX PHY layer is designed to handle both the PMA and the PCS functions. The PMA serves as the electrical interface to the physical medium and incorporates various essential components, including serializer or deserializer (SERDES), clock and data recovery PLLs, transmit drivers for the analogue front end, and receive buffers for the analogue front end.

Figure 3.8 illustrates a duplex-mode TX channel. In this setup, three types of PCS are available to manage data at different rates. Table 3.3 provides information about the supported PCS types and their corresponding data rates.

¹More information about the FEB will be given in Chapter 6.




Figure 3.8: High level overview of a TX channel with a transmitter and a receiver part, taken from [64].

Additionally, one of the available PCS options is the use of a PCIe IP core, as referenced in [65]. PCIe is a high-speed serial bus interface commonly used to connect peripheral devices with the Central Processing Unit (CPU) in a personal computer (PC). The electrical standard and protocol of PCIe will be discussed later in Section 3.9.

PCS type	data rate
Standard PCS	1.0Gbps to 10.813 44 Gbps
Enhanced PCS	1.0Gbps to 17.4 Gbps
PCIe Gen3 PCS	8.0 Gbps

Table 3.3: Supported PCS types.

Transceiver design IP blocks

Figure 3.9 provides an overview of the fundamental building blocks comprising an Intel Arria 10 TX. Each TX PHY IP Core is equipped with reset ports and reconfiguration registers, as well as controls for configuring the PCS and PMA settings.

The reset ports are connected to a TX PHY Reset Controller IP Core, which serves the purpose of resetting individual channels within the TX PHY IP Core. This IP Core monitors the status of the TX PLL IP Core and triggers a reset if PLL fails to lock onto an input clock signal. The PLL, in turn, generates a clock signal that is supplied to the TX PHY IP Core.

Furthermore, an option for an Avalon Memory Mapped Interface Main (Avalon-MM) interface is available. This interface allows for the reading and writing of reconfiguration registers within the TX PHY IP Core. In the context of the Mu3e experiment, this option is used to monitor the current status of the TXs during operation.

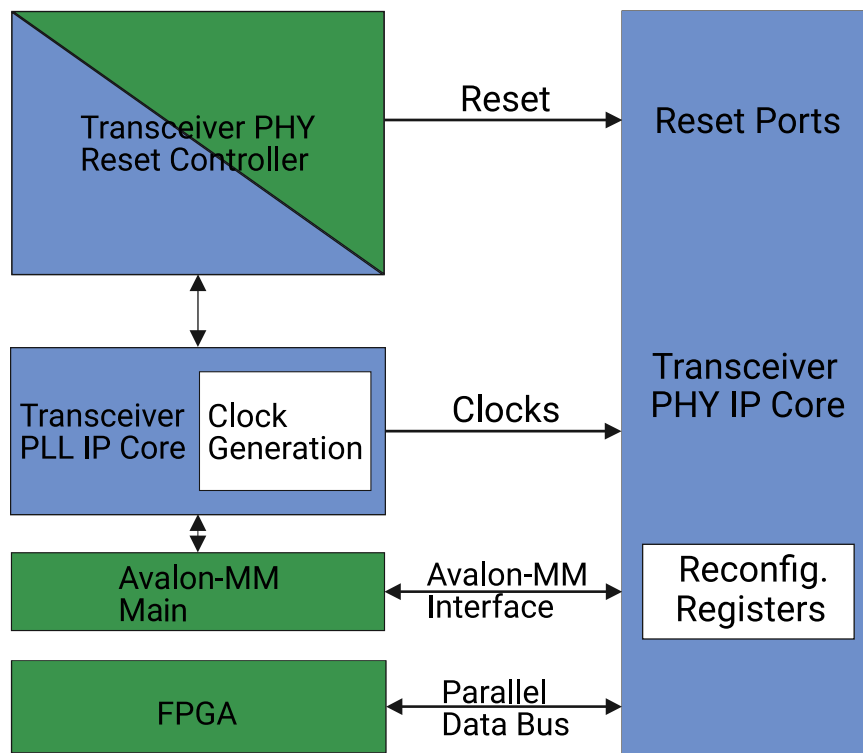


Figure 3.9: Fundamental building blocks of an Intel Arria 10 TX. The blue blocks are generated by Quartus the green blocks can be generated by the user, taken from [64].

Transceiver Channel datapath and clocking

In Figure 3.10, the datapath and clocking configuration of a TX operating at a data rate of 1250 Mbps is illustrated. A similar type of TX is used in the Mu3e experiment to facilitate the distribution of the reset protocol, as discussed in Section 5.2. It is worth noting that all other TXs utilised in this study differ solely in their data rate and parallel data width.

This specific TX operates with an input data width of 8 bits, received as a parallel stream from the FPGA core with a TX core clock (`tx coreclk`) running at 125 MHz. Initially, the data flows into the TX first in first out memory (FIFO), which efficiently organises the data buffer, ensuring that the oldest entry is serviced first upon output. Subsequently, the data at the FIFO output is sampled using a different clock (`tx clkout`) derived from the TX PLL.

Following its journey through the FIFO, the data undergo a 8b/10b encoding process, optionally with the inclusion of a byte serializer. This serializer has the ability to double or quadruple the data width of the PMA serializer, allowing flexibility in running the PCS at a lower parallel clock frequency to accommodate various FPGA interface widths. The resulting 10-bit data are then transmitted to the PMA serializer, implemented as a shift register. This register acquires the parallel data using the low-speed parallel clock and shifts them bit by bit with the high-speed serial clock. Additionally, the TX Bit Slip component can be used to alter the bit position of the 10 bit parallel data before serialisation. Moving to the receiver side of the TX channel, it begins by recovering both the high-

speed serial clock and the low-speed parallel clock from the serial data through CDR mechanisms. Subsequently, the serial data is parallelised within the PMA deserializer. The parallel data then proceed to the word aligner, which is responsible for aligning the bits by identifying the comma words within the 10 bit data. Once synchronised, the rate match process begins, allowing data sets to be inserted or removed to ensure that the rate match FIFO never overflows or underflows.

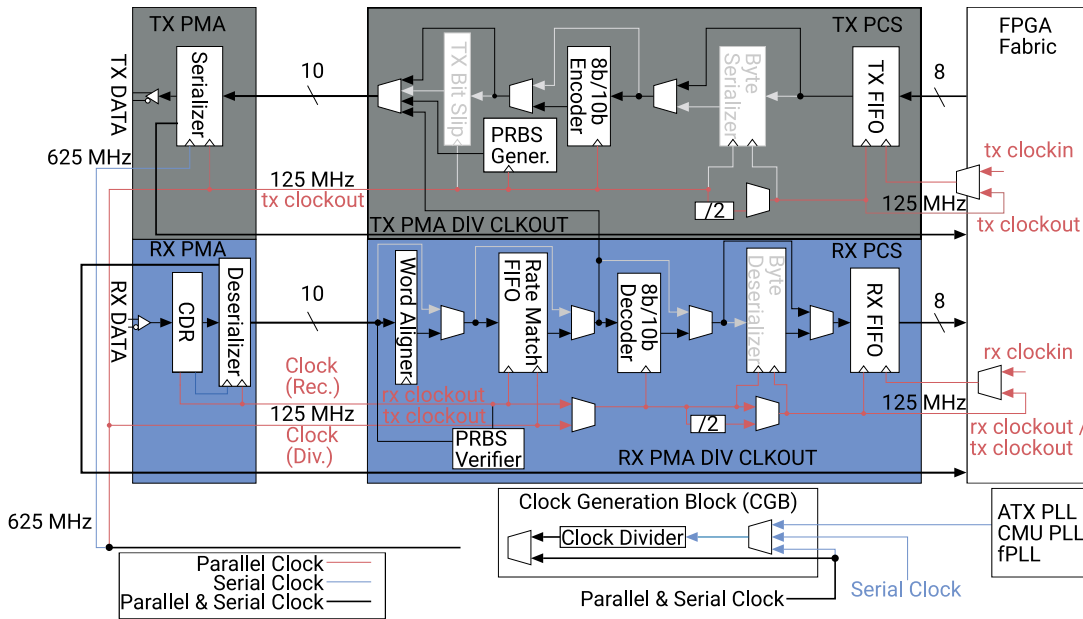


Figure 3.10: Transceiver (TX) channel datapath and clocking at 1250 Mbps, taken from [64].

Furthermore, this can compensate for frequency deviations, typically expressed in part-per-million (ppm), between the transmitter and receiver clocks. This compensation results in a clock difference of approximately $125 \text{ MHz} \pm 100 \text{ ppm}$. In the case of the Mu3e reset system, this rate match process is not used to ensure that the TX is always synchronised with the global 125 MHz system clock. Finally, the remaining portion of the receiver operation is focused on decoding the data to 8 bits and integrating it into the FPGA core.

3.7 Optical data transmission

To mitigate the potential disruptions caused by electromagnetic interference in electrical communication and much longer attenuation lengths, optical fibre technology is used. In this context, pulsed light is transmitted through optical fibre cables. In the transmitter section, light-emitting diodes (LEDs) are commonly used. However, in high-speed data applications, lasers are the preferred choice due to their higher optical output power and faster switching times. On the receiving end, photodiodes are used to convert optical signals back into electrical signals, as detailed in the reference [66].

Given the substantial amount of data that must be read for the Mu3e experiment, optical data transmission is the preferred method. This preference is driven by the fact that the bandwidth of optical communication exceeds that of copper-based systems, as discussed in the reference [58]. Furthermore, a galvanic decoupling of the detectors and the readout electronics can be achieved.

3.8 Bit Error Rate Tests

In the context of testing data links, pseudo-random data patterns can be generated and transmitted through the link. The receiver can predict the incoming data since pseudorandom data form a deterministic series. When comparing the received data with the predicted data, bit errors (BERs) can be determined.

If one considers a data rate of r during a time interval t , the total number of bits transmitted is given by $N_{bits} = r \cdot t$. To calculate BE, one uses the formula:

$$P_{BER} = \frac{N_{err}}{N_{bits}}, \quad (3.21)$$

where N_{err} represents the total number of BERs.

In cases where no BERs are detected in the data stream, an upper limit can be calculated for the BE. This can be accomplished using the Bayesian approach with a flat prior for a Poisson distribution, as discussed in [67]. Therefore, the upper limit for a BE with a CL of 95 % is given by:

$$P_{BER} \leq \frac{\ln(1 - CL)}{N_{bits}} \approx \frac{3}{N_{bits}}. \quad (3.22)$$

3.9 Data readout via PCIe

To facilitate communication between the Switching Boards (SWBs) and the PC interface boards (Receiving boards) within their PCs, as well as to coordinate all components of the DAQ system for the Mu3e experiment, a robust communication system is essential. This system connects the various elements, and since both FPGAs are equipped with PCIe interfaces, the high-speed serial bus standard PCIe is employed.

In the subsequent sections, a general overview of the data readout process based on PCIe is provided. Detailed explanations of PCIe and DMA are presented in Sections 3.9 and 3.9 respectively. These sections are based on references [68] and [69]. The initial implementation of the DMA block is credited to the work of [68].

Chapter 8 delves into a comprehensive explanation of the current version of this readout block. This chapter provides a detailed understanding of the implementation and functionality of this critical component in the Mu3e DAQ system.

Peripheral Component Interconnect Express (PCIe)

PCIe [70, 71, 65, 72] is a high-speed serial computer expansion bus standard specified by the PCIe Special Interest Group (PCI-SIG) [73]. It is commonly used to connect graphics processing units (GPUs) or network cards to the CPU via a computer motherboard.

Unlike older bus standards like Peripheral Component Interconnect (PCI), PCIe is not a bus in the traditional sense. It operates as a point-to-point connection, where each device has its dedicated connection to a network switch. Each of these lanes consists of four wires for transmitting and receiving data (two differential pairs each) and a pair of wires to connect the device to the reference

clock². PCIe provides full-duplex communication between two devices, with up to 32 lanes available for use. As a network-based communication system, data transactions occur in the form of packet transmissions, similar to a local ethernet network, which includes flow control, error detection, and retransmissions.

Version	Introduced	Line code	Transfer rate	Throughput for 8 lines
1.0 [70]	2003	8b/10b	2.5 GT/s	2.0 Gbyte/s
2.0 [71]	2007	8b/10b	5.0 GT/s	4.0 Gbyte/s
3.0 [65]	2010	128b/130b	8.0 GT/s	7.88 Gbyte/s
4.0 [72]	2017	128b/130b	16.0 GT/s	15.75 Gbyte/s

Table 3.4: PCIe link performance, taken from [74].

The Media Access Controller (MAC) address, which serves as the physical address of a device within a network connection, is primarily determined by the device's physical position on the motherboard and later translated to a higher-level address. Table 3.4 illustrates the data rates of different PCIe generations.

Communication via a PCIe link involves three layers: the Transaction Layer, the Data Link Layer, and the Physical Layer. To send write, read, or response (completion) packets, Transaction Layer Packets (TLPs) are utilised. The Data Link Layer ensures that each TLP reaches its intended destination by encapsulating each TLP with its header and a link-level cyclic redundancy check (CRC). Acknowledgements are used to confirm the successful delivery of each TLP. A flow control mechanism ensures that a package is sent only when the link partner is ready to receive data, introducing a slight uncertainty in arrival times.

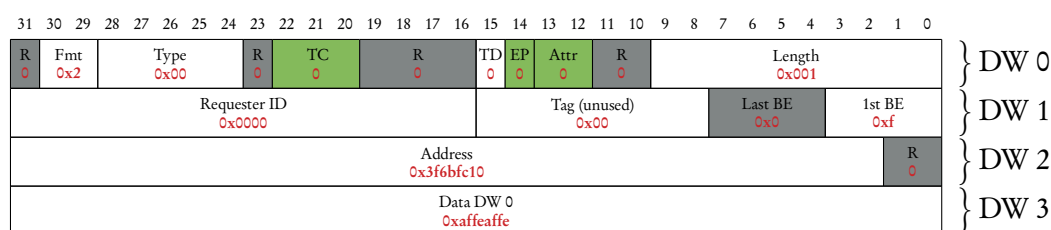


Table 3.5: PCIe write request, taken from [69].

In terms of complexity, a write- TLP is the simplest transaction as it does not require a response. Include the destination device, write address, and data. Table 3.5 provides an example of a write package. In this instance, the CPU writes the value 0xaffeaffe to 0xfdaff040 using 32-bit addressing. Hence, the package consists of four 32-bit words³. The first three DWs form the package header, while the last contains the data to be written. The fields marked in grey (R) are not used in this specific write package, and the green fields may contain non-zero values.

Here, Format (Fmt) and Type indicate that this package is a write command, with the TD bit set to zero to indicate that there are no extra CRC for the TLP data. As the Data Link Layer has its own CRC, no additional one is required, assuming that TLPs remain unchanged during transmission. The length field indicates the number of 32-bit words of data sent, followed by the Requester ID,

²If multiple lanes are utilised, only a single clock is employed.

³One double word (DW) has 32-bit.

which is set to zero, signifying that the packet originates from the PCIe port closest to the CPU (Root Complex). The Tag field can be filled with any value; in this example, it is set to zero. The first 32-bit DW has its 1st Byte Enable (BE) set to 0xf, indicating that all four bytes in the first data DW are valid. Given that the length is one, only one data DW is sent, and thus the last BE must be zero. The write address for the first data DW precedes the data DW, with the last significant bit (LSB) of the second DW set to zero, representing the write address itself. To obtain the correct address, one must multiply 0x3f6bfc10 by four, which results in 0xfdaff040. Regarding data transferred via PCIe, it's essential to note that it is in big-endian format, while Intel and AMD, both x86 processors used in the Mu3e experiment, employ little-endian format. Consequently, the data for the CPU in this example are 0xfeaffeaf. For addresses of more than 4 Gbyte, four DWs are used for the header, comprising two 32-bit words. According to the specification [70], the receiver cannot address 32-bit addresses with four DWs, necessitating the differentiation between 32-bit and 64-bit addresses. The payload for a write command can be much larger than a single 32-bit word, limited only by the peripheral's configuration. For Mu3e, the FPGAs connected via PCIe have a maximum payload size of 2048 bytes [75].

In the Mu3e experiment, only write commands are used to transmit data to the PC, and hence, only a brief mention of the read and completion packages is made. The read command includes the read address and length and receives a completion package in response, containing the requested data.

Direct Memory Access (DMA)

In the context of PCIe, any device within the network can perform DMA, allowing access to a system's memory independent of the CPU. Typically, this is achieved through programmed I/O. The advantage of DMA lies in relieving the CPU from the burden of handling read and write operations, which is especially beneficial in systems with multiple concurrent processes, such as the Mu3e experiment's PCs. However, it is possible to send an interrupt from the DMA controller to the CPU if there is a need for CPU intervention in data processing. However, sending such an interrupt interrupts the DMA processes and can lead to a reduction in data transfer rates. Consequently, interrupts are not used in the Mu3e experiment. It is important to note that the firmware utilised for executing DMA transactions in the Mu3e DAQ system was created and implemented by D. vom Bruch [68]. However, the author made certain modifications, which are elaborated in further detail in Section 8.2.

To enable DMA via a PCIe device, the device must know the physical buffer address to which the data should be transferred. This is typically achieved by creating a software driver that informs the peripheral about the buffer's size and address. In a Unix-based system like Linux, this driver is loaded into the Linux kernel, the core component of the operating system that interfaces with various devices. The driver's role is to facilitate communication between the kernel and specific hardware components.

Additionally, in a Linux system, memory is divided into two spaces: user space and kernel space. Normal programmes run in the user space, while the kernel code resides in the kernel space. To access different memory regions, the Linux memory manager can map physical memory from Random-Access Memory (RAM) or hard disk drive (HDD) to virtual memory, making it appear as a block of user space memory for programmes. This mapping is achieved using pages, the smallest units of memory management, which are typically set at 4 kbyte for most Intel and AMD x86 systems. For the DMA firmware used in the Mu3e experiment, the driver allocates page-locked memory at boot time to permanently fix the physical address where the DMA engine can write.

To inform the FPGA about the mapped address regions, the CPU can utilise dedicated memory of

the PCIe device through Memory-mapped input/output (MMIO). The PCIe device provides this functionality through Base Address Registers (BARs). The driver can share page addresses and the size of DMA memory by sending them via a PCIe write command to one of the BARs. The FPGA can also perform this operation in reverse.

In the Mu3e experiment, four BARs are implemented on the FPGA. Two of them for sharing static information such as DMA memory size or status flags (referred to as write- and read-registers from the PC perspective). The other two are used for dynamic information such as slow control communication between Maximum Integrated Data Acquisition System (MIDAS) and the detectors, each with write and read directions when viewed from the PC.

The Mu3e Experiment

Since the observations of neutrino oscillations by experiments such as the Brookhaven Solar Neutrino Experiment [76] Super-Kamioka Neutrino Detection Experiment (SK) [14], Sudbury Neutrino Observatory (SNO) [15], Kamioka Liquid Scintillator Antineutrino Detector (KamLand) [16], and others, it has become evident that lepton-flavour conservation is not an absolute rule in nature. As highlighted in Section 1, the detection of LFV in charged leptons holds the promise of providing crucial insights into new-physics phenomena.




Figure 4.1: History of searches for CLFV in μ and τ decays. Adapted form [77].

Figure 4.1 provides an overview of the historical searches for such rare decays. Specifically, the purple circles on the graph represent experiments dedicated to investigating the decay $\mu \rightarrow 3e$. The upcoming Mu3e experiment, to be conducted at the PSI in Switzerland, aims to find or exclude $\mu^+ \rightarrow e^+e^-e^+$ with a BF in the range of 1×10^{-12} to 1×10^{-16} . Data collection will occur in two phases: Phase I, starting at the end of 2024, and Phase II, beginning at the end of 2027, with the latter phase expected to further increase the sensitivity of the experiment.

The SINDRUM experiment previously established a limit of $\mathcal{B}(\mu^+ \rightarrow e^-e^+e^+) < 1.0 \times 10^{-12}$ at

a 90 % CL [40]. To improve the sensitivity to this BF by three orders of magnitude, a high μ rate of $10^8 \mu/s$ over a year of data collection is required. The anticipated data rate of the detector is so substantial that permanently storing all data would be prohibitively expensive.

In the second phase, the experiment aims to push the sensitivity even further down to $\mathcal{B}(\mu^+ \rightarrow e^-e^+e^+) < 10^{-16}$ with a muon rate of $2 \times 10^9 \mu/s$. However, this thesis focusses primarily on Phase I of the Mu3e experiment.

To gain a comprehensive understanding of how data are acquired in the Mu3e experiment, this chapter proceeds as follows: First, it explains the signal and background processes in Section 4.1, followed by an exploration of the detector concepts in Section 4.2. In addition, in Section 4.2.1 the beam line of the experiment will be explained. The target and magnet in the following are discussed in Section 4.2.2 and Section 4.2.3 respectively. Subsequently, in Section 4.3, the pixel detector is elucidated, which constitutes a significant part of the content of this thesis. At the end the readout application specific integrated circuit (ASIC) for both timing detectors and the timing detectors themselves are explored. This chapter is based on the Technical Design Report of the Mu3e Experiment [51].

4.1 Signal and background processes

The sub-figures in Figure 4.2 illustrate both the signal and the most common background processes in the Mu3e experiment.

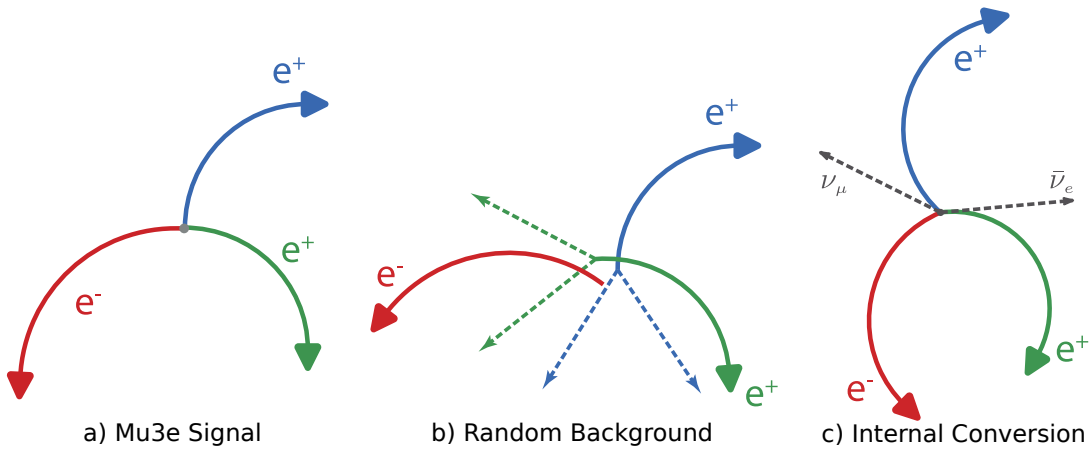


Figure 4.2: Overview of the signal process and the main background processes. In a) the signal process is drawn, b) shows the random combinatorial background and in c) the internal photon conversion is displayed.

Figure 4.2a displays the signal process $\mu^+ \rightarrow e^+e^-e^+$. This is the primary decay that the experiment aims to detect. Figure 4.2b illustrates the random combinatorial background. In this background scenario, various other processes are occurring simultaneously, resulting in particle tracks in the detector that do not converge at the same vertex. Figure 4.2c depicts the process of internal photon conversion, specifically $\mu^+ \rightarrow e^+e^-e^+\nu_e\nu_\mu$. In this process, additional neutrinos are generated. However, because they cannot be detected by the experiment, the measured energies do not sum up to the equivalent of the muon mass. To efficiently distinguish these background processes from the signal, the detector

must have both momentum and time resolutions better than 1.0 MeV/c and 1 ns, respectively. In the following, the detector design is described that was designed to fulfil these requirements.

4.2 Detector design

The schematic layout of the detector is presented in Figure 4.3. The detector is designed in a cylindrical shape and comprises a central station located in the middle, as well as two recurl stations, one upstream (US) and one downstream (DS) of the incoming beam.




Figure 4.3: Sketch of the Mu3e detector concept, taken from [51].

The incoming muon beam is stopped on the Mu3e Target, which is a hollow double cone constructed from Mylar foil to minimise the scattering of the positrons in the target (the muon beam is discussed in more detail in Section 4.2.1 while the Mu3e Target is explained in Section 4.2.2). Stopped muons decay at rest, and their decay products are deflected by the 1 T solenoid magnet field. A silicon pixel tracking detector is used to measure the trajectories of the electrons produced in the decay.

At the recurl stations, the particle trajectories intersect with the detector once again, extending the path length of the track. To maintain a high vertex and momentum resolution while mitigating the influence of multiple scattering within the detector material, the pixel layers are designed to be exceptionally thin. For this purpose, monolithic active-pixel sensors known as MuPix are utilised. Detailed information on the MuPix will be discussed in Section 4.3.

For time measurements at the central detector station, a scintillating fibre detector is employed, while scintillating tiles are used in the recurl stations [78]. Both detectors are read out using SiPMs and an ASIC known as Muon Timing Resolver including Gigabit-link (MuTRiG) [79].

In cases where scattering uncertainties dominate, the relative momentum resolution is directly related to the scattering angle Θ_{MS} and is inversely proportional to the track deflection Ω . This relationship is depicted in Figure 4.4a [51]:

$$\frac{\sigma_p}{p} \propto \frac{\Theta_{MS}}{\Omega}. \quad (4.1)$$




Figure 4.4: Figure 4.4a: Multiple scattering as seen in the plane transverse to the magnetic field direction. The red lines indicate measurement planes. Figure 4.4b: Multiple scattering for a semicircular trajectory, both taken from [51].

To achieve a high-quality momentum resolution, it is advantageous to have a detector geometry with widely spaced layers. However, this must be carefully balanced with considerations regarding the acceptance of the detector, since larger layer spacings reduce the acceptance for low-momentum particles. Additionally, smaller spacings between layers facilitate the process of track finding.

Figure 4.4b illustrates that the trajectory, which is distorted by multiple scattering, overlaps with the undistorted initial trajectory after approximately half a turn, corresponding to $\Omega = \pi$. At this point, the scattering uncertainties cancel out to a first-order approximation, resulting in the best achievable momentum resolution. Consequently, the geometry of the Mu3e detector is optimised to take measurements at $\Omega = \pi$ for most of the tracks, ensuring optimal momentum resolution.

4.2.1 The Compact Muon Beam Line

To achieve its Phase I sensitivity goal of 2×10^{-15} while mitigating the challenges posed by the combinatorial background, the Mu3e experiment is based on specific beam characteristics. It benefits significantly from a continuous beam structure rather than a pulsed one, as this allows for a lower instantaneous muon rate. These conditions are met by the High Intensity Proton Accelerator (HIPA) complex at the PSI, operating at 1.4 MW of beam power.

Mu3e necessitates a muon beam primarily composed of "surface muons", generated from pions that stop and decay on the surface of the primary production target [80]. The intensity of the surface muons, and hence the beam intensity, peaks around 28 MeV/c, which is close to the kinematic edge of the two-body momentum spectrum for the decay of the pion at rest.

The Mu3e Experiment

To meet the intensity target and address the low energy requirements, a beam line capable of efficiently guiding these muons to a small and thin stopping target with minimal losses is essential. At the same time, it must minimise beam-related backgrounds. Achieving this balance involves several considerations.

- **Small Beam Emittance:** To ensure efficient transport, a small beam emittance is required.
- **Moderate Momentum Byte ($\Delta p/p$):** A moderate momentum byte (full width at half maximum of the momentum acceptance) is necessary to maintain a balance between beam intensity and stopping density in the target.
- **Achromatic Final Focus:** An achromatic final focus is employed to optimise the beam intensity and the stopping density within the target.

Additionally, minimising beam-related backgrounds, which can arise from various sources such as Michel positrons from μ^+ -decay or positrons produced from π^0 -decays in the production target, requires strict constraints on the amount of material along the beam path. This includes components such as windows and momentum moderators. Consequently, an extension of the vacuum system is required to extend just in front of the target.

In summary, the Mu3e experiment relies on a continuous beam structure and specific beam line characteristics provided by the HIPA complex at PSI to achieve its sensitivity goals while effectively managing beam-related backgrounds and meeting the unique requirements of the experiment.

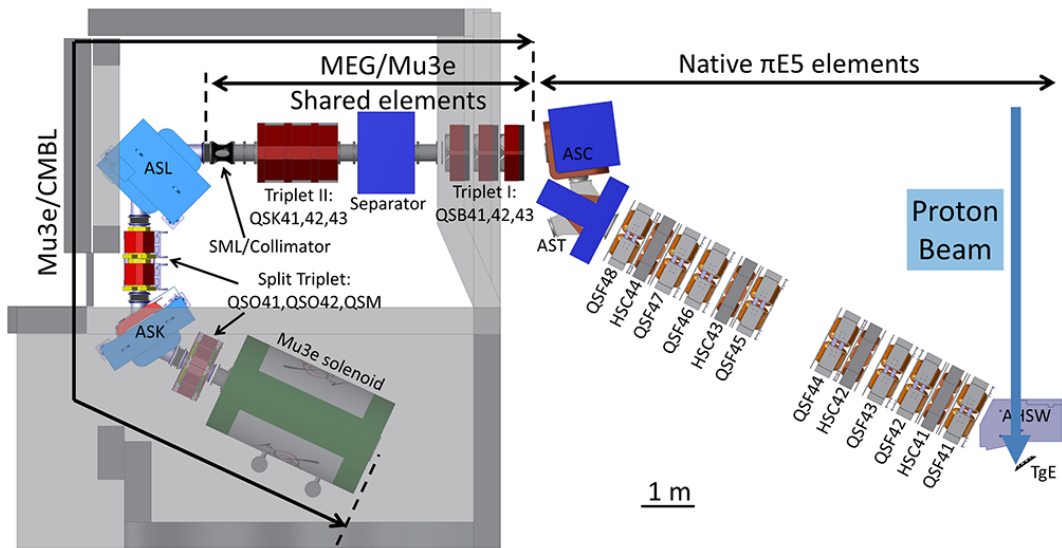


Figure 4.5: Computer-aided design (CAD) model of the entire π E5 channel and Compact Muon Beam Line (CMBL) used as a basis for the G4BL models, taken from [51].

For Mu3e Phase I, achieving muon intensities near 10^8 muons per second is imperative. This requirement narrows down the options for facilities to just one worldwide: the π E5 channel at PSI. However, this channel will be shared with the upgraded version of the MEG experiment, known as MEG II [81]. Notably, MEG II boasts a large detector and established infrastructure permanently situated in the rear portion of the π E5 area.

The innovative CMBL designed for Mu3e offers several advantages. First, it facilitates the placement of the Mu3e solenoid, which measures 3.2 m in length, at the front of the π E5 area (see Figure 4.5). Second, it enables both the MEG II and Mu3e experiments to efficiently share the critical front beam transport components required by both efforts. This dual-use solution streamlines the process of switching between experiments, necessitating only the replacement of MEG II's superconducting beam transport solenoid with a Mu3e-specific dipole magnet, referred to as ASL.

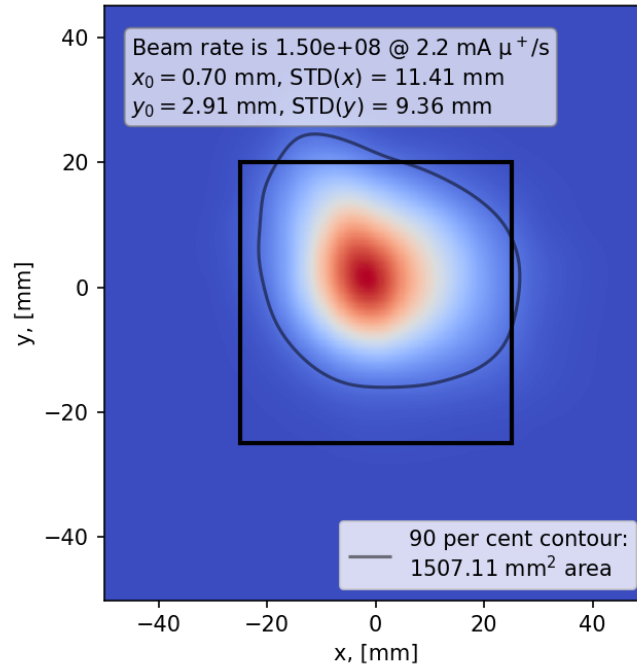


Figure 4.6: Measured beam profile at the collimator, taken from [82].

The estimates for the ultimate muon stopping rate at the target are grounded in the $1\text{-}\sigma$ beam emittances measured at the intermediate collimator system in 2018. These values are $\epsilon_x = 950 \pi \cdot \text{mm} \cdot \text{mrad}$ and $\epsilon_y = 490 \pi \cdot \text{mm} \cdot \text{mrad}$. To complement this, simulations using G4BL [83] were also carried out (as detailed in [84]). Although the intensity of the muon beam upon injection into the solenoid successfully meets the commissioning objective, the critical factors governing the diameter of the stopping target are the inner diameter of the silicon detector and the corresponding beam pipe size. As a result, the radius of the stopping target has been optimised to reach 19 mm. These parameters represent a careful balance between stopping rate, occupancy, and vertex resolution.

The primary losses occur during the transition from the initial beam pipe diameter to the eventual reduction to 40 mm at the target's end. During a second beam test in 2022 the beam spot was measured at the collimator with the full beam pipe insert. The measured beam spot is shown in Figure 4.6. During the same tests, a muon rate of approximately $8.4 \times 10^7 \mu/s$ on the target was measured [82]. A rate dependence with different values of the magnetic field of the spectrometer can be seen in Figure B.1 in the Appendix B.1. Additionally, further improvements in the muon rate are planned, and ongoing research is exploring further enhancements.

4.2.2 The Mu3e Target

Designing the stopping target presents a significant challenge, requiring a delicate balance between optimising stopping power and minimising material usage to reduce backgrounds and their impact on track measurements. The target should contain just enough material in the beam direction to effectively stop most muons. This is achieved in part by introducing a moderator in the final part of the beam line. However, it is crucial to keep the target as thin as possible to minimise the material in the flight path of decay electrons entering the detector's acceptance region. Using a low-Z (low atomic number) material is advantageous because it suppresses photon conversion and large-angle Coulomb scattering. Additionally, spreading out the decay vertices as widely as possible helps reduce accidental coincidences of track vertices and produces a more evenly distributed occupancy in the innermost detector layer.




Figure 4.7: Hollow double-cone muon stopping target made of aluminised Mylar foil, taken from [51].

At PSI, a manufacturing process was developed to create a complete target, as shown in Figure 4.7. Each individual hollow cone within the double-cone structure is manufactured separately, comprising a sandwich structure composed of 2 or 3 thin Mylar foils rolled and glued together with epoxy adhesive. The thickness of the Mylar foils and the combination of several foils are carefully selected to achieve the desired final thickness. Finally, the two individual cones are joined together to form the hollow double-cone structure.

In each sandwich stack, the inner and outer foils are coated with aluminium and their orientation ensures that both the inner and outer surfaces of the cones have an aluminium layer. The presence of conductive surfaces, combined with mounting on a conductive carbon tube, prevents the target from charging up because of the high stopping rate of positive muons.

4.2.3 The Mu3e Magnet

The magnet designed for the Mu3e experiment plays a critical role in ensuring precise momentum determination for muon decay products. It must generate a homogeneous solenoidal magnetic field with a strength of $B = 1\text{ T}$. To perform particle tracking, field uniformity is crucial, with field inhomogeneities along the beam line required to be less than 10^{-3} within a $\pm 60\text{ cm}$ region centred on the magnet's core.

Additionally, the magnet serves as a vital optical element of the beam, guiding the muon beam to the target. To enhance field homogeneity and ensure compatibility with the magnetic field of the last

beam elements in the CMBL, compensating coils are strategically placed on either side of the magnet. In Figure 4.8, one can see the magnet’s delivery to PSI’s experimental hall in July 2020.




Figure 4.8: Picture of the delivery of the Mu3e magnet to PSI’s experimental hall, taken from [51].

4.3 MuPix Pixel Sensor

The Mu3e pixel tracker plays a crucial role in providing precise hit information to reconstruct the tracks of electrons generated in muon decays. The success of the experiment is highly dependent on the achievement of optimal vertex- and momentum-resolution measurements for these electrons. Given the influence of multiple scattering, it is crucial to minimise the material within the active region of the tracking detector. Consequently, the tracker relies on HV-MAPS (see Section 2.4 for a more general introduction), which can be thinned to $50\ \mu\text{m}$ to reduce the radiation length to approximately $0.5 \times 10^{-3} X_0$. To enhance the performance, the detector operates in a dry helium atmosphere and is cooled by a helium gas flow, effectively reducing the impact of multiple scattering.

Ten prototypes have been developed in preparation for the Mu3e experiment, with the latest iteration being MuPix11. Figure 4.9 displays the layouts of the selected prototypes, with MuPix10 being the focus of subsequent discussions. Although MuPix10 and MuPix11 share identical size and performance specifications, MuPix10 had a configuration bug that slowed down the setup time, which was rectified in MuPix11.

The MuPix7 sensor was the first prototype to incorporate all essential functionalities required for

the Mu3e experiment. It featured a fully integrated readout state machine, high speed PLL, and a fast serial output link capable of operating at speeds of up to 1.6 Gbit/s, with the ability to transmit signals over distances greater than 2 m. The active area of the MuPix7 sensor was approximately $3.2 \times 3.2 \text{ mm}^2$.




Figure 4.9: Layouts and size comparison of selected MuPix prototypes. The main area is the active pixel matrix, and the stripe at the bottom contains the digital electronics (periphery). For reference, the size of MuPix10 is $20.66 \times 23.18 \text{ mm}^2$, taken from [51].

Next came the MuPix8 sensor, produced in a large-scale engineering run with varying p-substrate resistivity. MuPix8 was manufactured using the AH18 process by ams AG. It was designed to be more radiation-tolerant and equipped to handle high data rates with three additional serial links. In addition, a second timestamp capability was introduced to measure Time-over-Threshold (ToT), allowing for offline correction of time-walk effects due to pulse height variations. MuPix8 was mainly used to study aspects related to sensor size, including power network design and potential cross-talk issues along long analogue read-out lines.

Addressing specific system considerations for the construction of the pixel tracker, the MuPix9 sensor was developed as a dedicated small-area prototype. It featured fast differential control inputs, a novel chip configuration scheme, and a shunt Low Drop Out (LDO) regulator for studying serial powering [85].

MuPix10, the final prototype, was specifically designed to build pixel modules. It boasts an active area measuring $20.48 \times 20.00 \text{ mm}^2$ and pad layout compatibility with pixel module spatial requirements and High-Density Interconnect (HDI) design guidelines. This sensor was produced in an engineering run at TSI and delivered in March 2020. Figure 4.10 provides a block diagram of the MuPix10 layout, with the primary functional blocks outlined and described in the following sections.




Figure 4.10: MuPix10 block diagram, taken from [51].

4.3.1 Pixel cell electronics

Each pixel in MuPix10 is composed of essential components: the sensor diode, a charge-sensitive amplifier, a source follower to drive the signal to the chip periphery, and a capacitor to allow injection of test charges. Figure 4.11 shows the different components from left to right:

- Charge Injection (C Inj): Here, the test charges can be injected into the pixel.
- Charge Collecting Diode (S): This diode collects the charge generated by particle interactions.
- N-Well Bias Restoration Circuit (B): This circuit helps maintain the bias voltage for the n-well, an important aspect of the pixel's operation.
- Amplifier (A): The charge collected by the diode is amplified here.
- Feedback Line (fb): This line provides feedback to control the amplification process.
- Baseline Circuit with Baseline Restoration (BLR): This circuit helps maintain a stable baseline signal.
- Source Follower (out SF): The output signal is driven to the periphery of the chip through the source follower.

The chosen baseline implementation utilises a p-channel metal-oxide-semiconductor (PMOS)-based amplifier with a source follower. Pulse shaping can be adjusted through bias currents, typically resulting in shaping times of approximately $1 \mu\text{s}$. The dimensions of the charge collecting diode were optimised using TCAD simulation to ensure a homogeneous electrical field, considering a substrate

resistivity of $200 \Omega \text{cm}$ and a depletion voltage of -60V . The guard ring was designed to achieve a breakdown voltage of -120V .

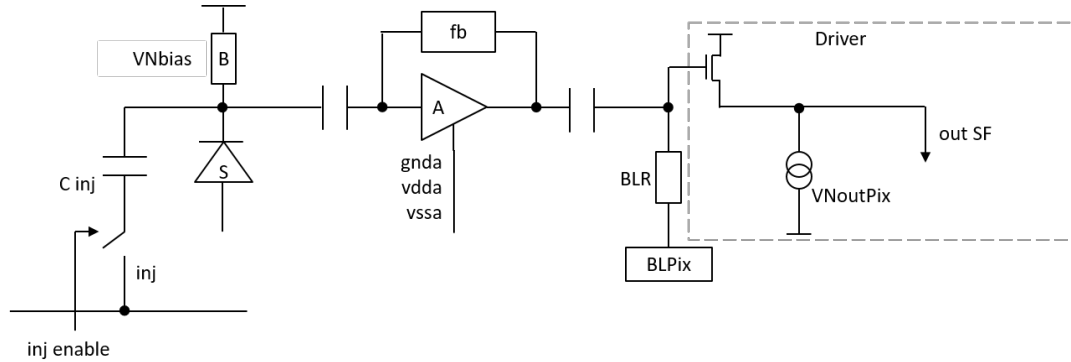


Figure 4.11: The schematic depicts of the analogue electronics within a pixel cell of the MuPix chips, taken from [51].

In terms of the readout buffer cell, the digital electronics were placed at the periphery of the chip in all MuPix designs. This strategic placement aimed to minimise cross-talk between rapidly switching digital signals and the sensitive analogue circuits. The Readout Buffer Cell, which covers about $160 \times 4.2 \mu\text{m}^2$ (approximately 10% of the active cell), spans two pixel columns in width through double column routing. The Readout Buffer Cell plays a crucial role in converting the analogue signal into an arrival time signal using comparators. A common threshold for the comparators is set globally, with individual 3-bit digital-to-analog converters (DACs) allowing fine-tuning of the threshold for each pixel. This feature can be employed to ensure a uniform signal response or noise suppression across the pixel matrix. A fourth enable bit is available to mask out noisy pixels. For testing purposes, the output of the pixel comparator can be monitored through a dedicated output line (hitbus signal).

A hit is defined by the rising edge of the comparator output, and for timestamp generation, the output is sampled with an adjustable frequency derived from the internal clock. In the latest MuPix prototypes, each Readout Buffer Cell has been enhanced with the addition of a second comparator. This 2-comparator threshold scheme enables the implementation of two distinct thresholds: a very low threshold close to noise for measuring the time-of-arrival (ToA) of the rising edge with minimal time-walk and a high threshold well above noise to generate the hit flag. This scheme, as demonstrated with the MuPix8 prototype, significantly improves time resolution. In Figure 4.12 the 2-comparator threshold method implemented in MuPix10 and MuPix11 is shown.

In MuPix10, one has the flexibility to operate with one single comparator or in mixed mode. In mixed mode, the ToA is defined by the lower threshold and the time of fall is defined by the higher threshold. This mixed mode helps mitigate time-walk effects and provides robust ToT information for subsequent time-walk correction. For each ToT bin the mean time is calculated; this can be done by fitting a Gaussian or simply calculating the mean value. The obtained mean is the time-walk per ToT bin and is used to correct the ToA:

In terms of data representation, MuPix10 uses 11 bits for the timestamp and 5 bits for ToT. Both measurements are sampled at adjustable frequencies derived from the internal clock. Grey counters are used for all counters, and ToA and ToT are stored in floating capacitors. MuPix10 can operate in different modes: with only one comparator or with two comparators, where one can choose to take

the ToT measurement from the lower or higher threshold (configurable).




Figure 4.12: The figure illustrates the 2-comparator threshold methods implemented in MuPix10 and MuPix11, taken from [51].

To improve efficiency and avoid long sampling times for signals with large ToT values, MuPix10 includes a hit delay circuit with a programmable timer to generate the hit flag. This circuit serves two purposes: first, it makes hit-flag generation independent of pulse height and helps maintain chronological hit order, and second, it reduces dead-time by excluding extended sampling times. The impact on ToT-based time-walk corrections is minimal, as time-walk effects are typically small for large signals.

4.4 The MuTRiG ASIC

MuTRiG is a 32-channel ASIC designed for reading out SiPM detectors. This ASIC was developed using UMC 180 nm CMOS technology and is primarily intended for use in the Mu3e experiment.

Its purpose is to read out both the fibre and tile detectors in the Mu3e experiment, aiming to achieve the necessary timing resolution for these systems while also being able to handle the high event rate associated with the scintillating fibre detector. MuTRiG represents an evolutionary development based on the STiCv3.1 chip [86], originally developed at the Kirchhoff Institute in Heidelberg for medical applications involving SiPMs, such as EndoTOFPET-US [87].

The analogue processing components of MuTRiG draw from the STiCv3.1 chip, which has shown satisfactory performance in various testing scenarios. However, STiCv3.1 can only transfer data at a rate of approximately 50 kHz per channel through its 160 Mbit/s data link. This transfer rate is insufficient for the timing detectors in the Mu3e experiment, particularly the fibre detector, which needs to handle an event rate of 1 MHz per channel to achieve 100% readout efficiency. MuTRiG improves the excellent timing performance of STiCv3.1 by introducing a fast digital readout system designed for high-rate applications. The expected analogue timing jitter for the common front-end (FE) is around 15 ps. Figure 4.13 illustrates the channel diagram of the MuTRiG chip, which has the following components:

- Input Stage: The SiPM signal current is buffered with low input impedance to minimise the input time constant. It also allows one to adjust the SiPM bias voltage per channel.
- T- and E-Trigger branches: The signal is split into two branches: T-Trigger (for time discrimination) and E-Trigger (for energy discrimination).
- Hit Logic Module: Signals from the T-Trigger and E-Trigger branches are processed to generate a combined hit signal, see also Figure B.2 of the Appendix B.1.
- TDC Module: The time information from the T-Trigger branch is measured using a time-to-digital converter (TDC) module.




Figure 4.13: The diagram illustrates a MuTRiG channel’s components and signal flow, taken from [51].

Besides the features shown in Figure 4.13 the MuTRiG has on-chip memory to buffer the digitised time stamps before the data are transferred. For the configuration of the analogue FE, TDC, and digital modules a Serial Peripheral Interface (SPI) interface is used. Figure B.2 of Appendix B.1 provides a schematic representation of the chip’s functionality ¹.

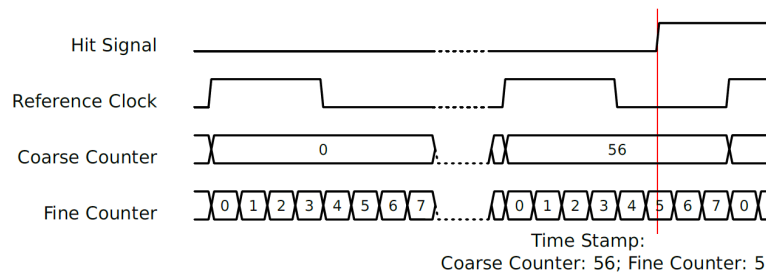


Figure 4.14: Working principle of the TDC, showing the fine and coarse counters, reference clock and example arrive of a hit signal, taken from [51].

The MuTRiG chip achieves its excellent timing resolution through its differential analog FE and the 50 ps binning TDC, which were inherited from the STiCv3.1 chip. The fundamental working principle of a TDC is illustrated in Figure 4.14. When a hit signal exceeds a certain threshold, the TDC module samples the state of a coarse counter, which increments at a rate of 625 MHz, driven by a reference clock.

¹More comprehensive information can be found in the associated reference [88].




Figure 4.15: Schematic of the MuTRiG TDC, taken from [51].

A fine counter with 50 ps bins is utilised to make a more precise measurement of the hit-time within the 1.6 ns coarse counter bin. The fine time data are obtained from sampling the state of 16 phase-shifted versions of the 625 MHz clock as produced by the VCO, see lower left of Figure 4.15. The values of both the coarse and fine counters are recorded as the time stamp of the hit signal. The TDC also records the time at which the signal drops back below the threshold, which is used for energy measurements via the ToT. The Global TimeBase Unit provides common coarse and fine counter values for all channels for time stamping, as shown in Figure 4.15. The TDC requires approximately 30 ns to reset after a hit, resulting in a maximum occupancy of around 30 million events per second per channel.

To accommodate high-rate data readout, MuTRiG features a double data rate serializer and a custom LVDS transmitter, developed to establish a gigabit data link with the DAQ system for data transmission. Event data from all channels are buffered and sent out in frames via a 1.25 Gbit/s LVDS serial data link. To enhance the chip's event rate capability, the output event structure can be switched from the standard 48 bits (containing both the time stamps when a hit signal exceeds and returns below the threshold) to a shorter event structure of 27 bits, containing only the first of these times and a 1-bit energy flag for the hit.

4.5 The scintillating fibre detector

To eliminate any potential sources of combinatorial background arising from tracks with different timing, Mu3e requires a detector dedicated to excellent time measurement of the tracks in its central region. Furthermore, this detector must be thin, offer a high level of efficiency, and be able to handle high event rates. To fulfil these requirements, Mu3e has developed a scintillating fibre detector (referred to as the SciFi detector). The time resolution of the SciFi detector was measured to be 250 ps, while achieving an efficiency of 95%. The spatial resolution of the detector is around 100 μm and the thickness is minimal, at 0.2% of radiation length [89].




Figure 4.16: In Figure 4.16a a full size SciFi ribbon prototype with preliminary holding structure is shown. The SciFi ribbon is made up of four staggering layers of round scintillating fibres. In Figure 4.16b the front view of the prototype is given, taken from [51].

In general, the SciFi detector has been designed to meet the stringent requirements of Mu3e, effectively reducing the combinatorial background and allowing precise tracking and timing measurements in the central region of the experiment. The mechanical design of the SciFi detector is explored below. The scintillating fibres in the Mu3e experiment are positioned between the second and third layers of the pixel detector. These fibres are organised into ribbons, each of which comprises four layers, with a total of 128 fibres per ribbon. Each ribbon has specific dimensions, measuring 300 mm in length and 32.5 mm in width. To visualise this setup, see Figure 4.16, which provides an image of a prototype ribbon used in the experiment.




Figure 4.17: In Figure 4.17a the connection between the fibre ribbons with the SiPMs via the MuTRiG PCB is shown. In Figure 4.17b a CAD rendering of the full SciFi detector used in the Mu3e experiment is shown. A SciFi super-module is made of two fibre ribbons.

To facilitate data readout from the SciFi detector, the ribbons are connected to arrays of SiPMs at both ends. The MuTRiG ASIC is then connected to a custom-made printed circuit board (PCB) (SciFi module board (SMB)) that connects to the SiPMs via a flexible printed circuit. Figure 4.17 provides a CAD rendering that illustrates the position of the SciFi detector within the Mu3e experiment setup.

4.6 The scintillating tile detector

The tile detector plays a crucial role in providing the most precise timing information for particle tracks, especially since it is located at the very end of curving particle trajectories. Unlike scintillating fibres, there are no restrictions on the amount of detector material for the tile detector. However, the placement of the tile detector inside the recurl pixel detectors imposes tight spatial constraints. A sketch of the modular detector can be seen in Figure 4.18.




Figure 4.18: a) Tile matrix mounted on a tile module with readout electronics. b) Side view of the US tile station. c) Tile stations inside the Mu3e experiment.

The tile detector comprises plastic scintillators segmented into small tiles, each tile being read out by a SiPM directly attached to the scintillator. The primary objective of the tile detector is to achieve a remarkable time resolution of better than 100 ps and a high detection efficiency close to 100 % to further suppress accidental background [51, 78].




Figure 4.19: Tile scintillator geometry: (left) edge tile, (right) central tile, taken from [51].

The tile detector is organised into two identical stations, with each station placed in one of the recurl stations. Each segment of the tile detector is shaped like a hollow cylinder that surrounds the beam tube. The length of a segment extends to 34.2 cm along the beam direction (z-direction), including

the endrings. The outer radius of the segment is limited to 6.4 cm, determined by the surrounding layers of the pixel sensor.

Within each recurl station, the tile detector is divided into 52 tiles in the z direction and 56 tiles along the azimuthal angle (φ -direction). This segmentation scheme represents the highest feasible channel density while considering the space requirements for the readout electronics. High granularity is critical to achieving low occupancy and high time resolution in the detector.

The technical design of the tile detector follows a modular concept, where the detector is composed of small, independent detector units. The smallest part is a tile matrix made up of 4×4 tiles. 26 of these matrices are built together to one tile module and one station has in total seven modules.

The individual tiles used in the detector are constructed from Eljen Technology EJ-228 plastic scintillator and have dimensions of approximately $6.3 \times 6.2 \times 5.0 \text{ mm}^3$, as shown in Figure 4.19. The edges of the two outer rows of tiles in an array are bevelled at an angle of approximately 25.7° . This bevelling allows for the arrangement of seven base units approximately in a circular configuration.




Figure 4.20: Individual ESR foils for two types of scintillator tiles: (left) edge tile, (right) central tile, taken from [51].

Each individual tile in the detector is wrapped with an Enhanced Specular Reflector (ESR) foil for optical isolation and to increase the light yield. The foil is designed to cover the entire tile except for an opening window that matches the size of the SiPM surface. This configuration can be observed in Figure 4.20. The SiPMs are soldered onto a flexible PCB and connected to the ASICs on the readout board, referred to as the Tile module board (TMB).

Part II

The Data Acquisition of the Mu3e Experiment

The hunt for new physics is done in multiple ways. One of them is to investigate for small deviations from the Standard Model of particle physics (SM). The search for rare decays is a possible way to achieve this. These investigations challenge not only the SM but also state-of-the-art technologies to process the large amount of physical events that must be studied.

The Mu3e experiment is one of the precision experiments exploring decay $\mu^+ \rightarrow e^+e^-e^+$. In the following part, the concept of the data acquisition (DAQ) system of the experiment is discussed, which is the key part in handling the large amount of data produced.

Mu3e Data Acquisition System

In the next chapter, the general setup of the DAQ system of the Mu3e experiment is described. It will provide an overview of the various building blocks and dive into the rates and data flow within the system. The chapter ends with a short related work section about triggerless DAQ systems. It is important to note that the chapter builds on previously published work by the author, namely references [90, 91].

While the development of a complex system involves collaborative efforts, certain aspects were specifically undertaken by the author during the course of this thesis work. The author notably contributed to the development of the SciFi detector datapath on the Front-end board (FEB), the implementation of the time-alignment algorithm, and the design of the data flow on the Switching Board (SWB), as well as the entire datapath on the PC interface board (Receiving board). Furthermore, the author contributed to the Maximum Integrated Data Acquisition System (MIDAS) software backend and developed the Mu3e Online Analyzer framework, which is used to perform online data quality checks.

5.1 Overview of the Mu3e DAQ

Figure 5.1 illustrates the full Mu3e DAQ system, which consists of three layers of field programmable gate array (FPGA) boards. Due to the high granularity and particle rates of the detector, the DAQ system must handle a substantial amount of data. It is designed to accommodate an expected data rate of more than 100 Gbit/s at 10^8 muon stops per second.

The custom-developed FEBs are used to read out the different detectors and are the first layer of the DAQ system. These boards are placed within the magnetic field and connected to the various subdetectors via 1.25 Gbit/s low-voltage differential signaling (LVDS) links. The main functions of the FEBs are to read out and configure the detector application specific integrated circuits (ASICs) and time sort the received hits into one stream from multiple detectors. Chapter 6 provides a detailed description of the firmware used to read out the SciFi detector and an introduction to the sorting algorithm within the FEBs.

The second layer of the DAQ system consists of the PCIe40 Board (referred to as SWB), originally developed for the A Large Ion Collider Experiment (ALICE) and Large Hadron Collider beauty (LHCb) upgrades. These boards are placed outside the magnet and connected to the FEBs via 6.25 Gbit/s optical links. They are responsible for the time-alignment of the different data streams and the transmission of detector configurations to the FEBs. Chapter 7 provides a detailed explanation of the different firmware blocks of the SWB.

The final layer of the system comprises commercial Terasic DE5a-Net-DDR4 boards (referred to

as Receiving board) with an Arria10 FPGA and onboard Double Data Rate Synchronous Dynamic Random-Access Memory (DDR SDRAM). These boards use DDR SDRAM to buffer the data from the full detector, while the hits from the four layers of the central pixel detector are sent to the graphics processing unit (GPU) via Direct Memory Access (DMA) at a maximum speed of 38 Gbit/s [53]. In order to be able to process the whole 100 Gbit/s multiple Receiving boards are daisy chained, each processing a fraction of the data. The GPU processes these data for particle track and vertex reconstruction.




Figure 5.1: A sketch of the final Mu3e DAQ system is shown. For the Integration Run 2021 and Cosmic Run 2022, the dashed parts of the system were used.

In the Mu3e experiment, the low momentum of the decay particles and the strong magnetic field lead to highly curved tracks. These particles can produce hits in physically distant parts of the detector. To accurately reconstruct the tracks, it is essential to have access to the data from the entire detector, even on individual farm nodes. The readout network rearranges the data, allowing the farm nodes to process the complete detector data from different time slices. Since the Mu3e experiment searches for a three-body decay that is randomly distributed throughout the detector, a traditional trigger system cannot be applied to select interesting events. Specifically, the rate is too high for a 3-fold timing coincidence. Therefore, a complete online track reconstruction must be performed using pixel hits from the central detector to select events of interest. After the reconstruction of the track, the selected events are read and sent to the Maximum Integrated Data Acquisition System (MIDAS), which controls the entire DAQ system and sends the data to mass storage. The total filter farm of the final system consists of twelve units of farm servers daisy-chained together to handle the expected data rate of 100 Gbit/s from the central pixel detector. After the selection of the event, the rate of the selected events will be of the order of 100 Mbyte/s. The farm servers are connected by 10 Gbit/s Ethernet to transmit the reduced data rate to MIDAS. Chapter 8 provides a detailed explanation of the data flow within the farm server and a short overview of the online event selection.

For data quality checks and more precise track reconstruction, during this work a framework called the Mu3e Online Analyzer was developed, based on the analyzer [92]. This framework processes the selected GPU events using the complete Mu3e reconstruction [93]. It can be executed on the personal computers (PCs) hosting the SWBs to perform quality checks at each stage of the DAQ system. Section 8.5 offers a more detailed discussion of this framework.

5.2 Clock and Reset System

To synchronise the different detectors, a dedicated clock and reset system was constructed. The Genesys-2 FPGA board [94] serves as the central component, generating a 125 MHz clock and synchronised reset signals for start and stop runs. Additional electronics create 144 copies of the clock and reset signals, which are distributed to each part of the DAQ system. For a comprehensive description of the final Mu3e DAQ system, please refer to the publication [95]. In Figure 5.2 a 3D representation of the full clock and reset distribution system and the final system inside a 19-inch rack-mountable box is shown.




Figure 5.2: Left: A 3D representation of the full clock and reset distribution system. On the left side is the Genesys 2 board with the clock and reset FPGA Mezzanine Card (FMC) distribution board. On the right side is the active splitting motherboard with an optical receiver (centre of the motherboard) that accepts the clock or reset lines from the FMC board via an optical fibre. The motherboard electrically routes the eight signals to the fan-out daughter boards, where each board generates 36 optical copies of the routed signal. Right: The full clock and reset system in a 19-inch rack-mountable box, taken from [51].

5.3 Readout ASICs

The Mu3e Phase I detector consists of a total of 2844 MuPix sensors and 278 Muon Timing Resolver including Gigabit-link (MuTRiG) ASICs, which read scintillating fibres and scintillating tiles. The scintillating fibres are read out by 3072 silicon photomultiplier (SiPM) readout channels, while the tile detector has 5824 scintillating tiles. The MuPix chips have three links in total, one link per submatrix for the inner vertex detector, and one link for all three matrices for the outer layers. Each

link is electrically connected to the corresponding FEB via a chain of connectors. On the other hand, the MuTRiG chip has only one link through which it sends data to the FEB.




Figure 5.3: Cross-section of the detector showing the active detector elements and the readout cabling. The scale is kept except for the thickness of active layers, flexes, cables and printed circuit boards (PCBs), taken from [95].

Due to the design of the detector, with a particle tracking volume outside the detector tube and all signal lines routed inside, there are tight space constraints for signal cabling. In Figure 5.3 an overview of the connections is given. The signals need to be transported out of the active tracking region, while the material is minimised to reduce the multiple Coulomb scattering of decay particles. To achieve this, the data signals from the pixel detector are first transported out of the active region using thin aluminium-polyimide High-Density Interconnects (HDIs) [96]. Flexible PCBs connected by interposers are then used to transfer the signals to micro twisted-pair cables, which lead to the Service Support Wheels (SSW) located close to the ends of the bore of the 3 m long 1 T solenoid magnet.

The MuPix chips from the inner pixel layer transmit data at a rate of up to 30 Mbits/s each over three links, while the outer layers are connected through one link. For the SciFi detector, each 128-channel SiPM array is connected to one SciFi module board (SMB) with four MuTRiG ASICs. These module boards are then connected to the SSW using microtwisted-pair cables. In the case of the tile detector, a MuTRiG chip reads 32 individual SiPMs, and thirteen MuTRiGs are collected on a tile module board. The tile module board is then connected to SSW using a ribbon cable. In general, signal cabling routing and readout path design in the Mu3e detector system are designed with the aim of minimising material and ensuring efficient data transfer from the detectors to the respective FEBs and SSW.

5.4 Data rate requirements

The bandwidth requirements for the different parts of the Mu3e DAQ system are summarised in Table 5.1. The calculations are based on a detailed simulation using Geant4 [97], taking into account the 8 bit/10 bit encoding [62] and a 75 % protocol efficiency (for every three data bits, four protocol bits are sent).

For the pixel detector, each FEB has a maximum hit rate of approximately 58 MHz. Each pixel hit contains 8 bit for column information, 8 bit for row information, 16 bit for chip address, 11 bit for

time information, and 5 bit bits for Time-over-Threshold (ToT) information. Therefore, the FEB requires a maximum bandwidth of around 4.6 Gbit/s. Of the total 88 FEBs for the full pixel detector, the 10 inner vertex FEBs plus 26 central FEBs are connected to the central SWB. The central SWB is the most critical in terms of bandwidth and has 8 links connected to the first farm server. Taking into account the encoding and efficiency, the maximum data output for the central pixel detector is 48 Gbit/s. The central pixel data must be sent to GPUs for track reconstruction via DMA, and the maximum data throughput using DMA to one GPU is measured to be 38 Gbit/s per farm server [53]. Note that the overall rate of all detectors is 100 Gbit/s but for the online track reconstruction only the data from the central pixel detector (48 Gbit/s) are needed to send to the GPUs. Since farm servers are daisy chained, each of them can process a part of 48 Gbit/s of the data.

Subdetector	Max. hit rate per FEB	Max. bandwidth per FEB	# of FEB	Max. output per SWB	Max. throughput to one GPU
Pixel	58 MHz	4.6 Gbit/s	88	48 Gbit/s	38 Gbit/s
Fibres	28 MHz	2.3 Gbit/s	12	24 Gbit/s	-
Tiles	15 MHz	1.2 Gbit/s	14	6 Gbit/s	-

Table 5.1: Data rate estimation from the detector simulation [97]. For the SciFi detector, clustering in the SWB FPGA is performed. For the bandwidth, 75 % protocol efficiency and 8 bit/10 bit [62] encoding are assumed.

For the SciFi detector, the maximum hit rate is 28 MHz. On the FEB each hit contains 2 bit for the ASIC ID, 5 bit for the channel ID, 15 bit for the coarse counter, 5 bit bits for the fine counter and 1 bit if the hit was an energy hit. Each SciFi detector FEB has two 6.25 Gbit/s optical link connected to the SWB, with a maximum bandwidth per FEB of around 2.3 Gbit/s. The complete SciFi detector consists of 12 FEBs connected to one SWB, which has four links to the first farm server. The maximum output of this is SWB is 24 Gbit/s. During GPU track reconstruction, the SciFi detector data is stored in the DDR SDRAM, and only the selected hits are read out via a second DMA engine with a total rate of around 100 Mbit/s.

For the scintillating tile detector, the maximum hit rate is 15 MHz, and each hit has the same format as the SciFi detector hits. On the FEB each hit contains 4 bit for the ASIC ID, 5 bit for the channel ID, 15 bit for the coarse counter, 5 bit bits for fine counter, and 15 bit for energy information. The maximum bandwidth per FEB for the tile detector is around 1.2 Gbit/s. To fit the tile hit into a 32 bit word only a part of the energy information will be sent to the SWB. Upstream (US) and downstream (DS) tile detectors are connected to two separate SWBs, each with a maximum output rate of 6 Gbit/s.

These bandwidth requirements are crucial for ensuring efficient data transfer and processing in the Mu3e DAQ system, taking into account high hit rates and the need to handle large amounts of data from different detector components.

5.5 Related work of DAQ systems

The main topic of this thesis is the design and integration of the DAQ system used in the Mu3e experiment. Therefore, some related work for DAQ systems in particle physics will be discussed below. Since the search for the $\mu^+ \rightarrow e^+e^-e^+$ process is a three-body decay at rest in the Mu3e experiment, it requires the full readout of the detector to select physics events. This is why the Mu3e DAQ system is designed to be trigger-free and synchronous [98].

In a trigger-free system, all detector data are continuously streamed without the need for a predefined

trigger condition. This is in contrast to trigger-based DAQ systems, where data are only read if a certain detector response exceeds a specific threshold determined by a trigger condition. Examples of trigger-based DAQ systems can be found in ATLAS [99] and Compact Muon Solenoid Experiment (CMS) [100] at the Large Hadron Collider (LHC). In the context of charged lepton-flavour violation (CLFV) experiments searching for muon decays, both the SINDRUM and the Mu to E Gamma (MEG) DAQ system had triggered systems, as described in Section 1.3.1 and 1.3.2.

The Mu3e DAQ system utilises a network of FPGAs and fast optical links to transport all data outside of the detector volume. This allows for online reconstruction of the entire detector data in real time. It is worth mentioning that the concept of full online reconstruction of the entire detector data is also being implemented in other experiments. For example, the LHCb [101] experiment is undergoing an update that includes the implementation of a complete online reconstruction of all detector data. Additionally, parts of the ALICE experiment [102] at the LHC have also adopted a similar approach for their detectors.

Overall, the triggerless and streaming design of the Mu3e DAQ system allows for a comprehensive analysis of the data, allowing the search for rare decays such as $\mu^+ \rightarrow e^+ e^- e^+$. In the following, all the different parts of the DAQ system developed during this work will be explained.

6

Front-end Board

In the next chapter, the main firmware components for the FEB in the Mu3e system are explained. This chapter will provide a comprehensive explanation of the hit-time sorter and the different communication protocols utilised.

The hit-time sorter is a crucial component of the firmware responsible for sorting the hits from the detectors in time. Understanding its functionality and implementation is essential for the proper operation of the FEB.




Figure 6.1: Picture of the FEB.

Additionally, the chapter will cover various communication protocols employed in the system. These protocols facilitate the exchange of data between different components, such as the FEBs and the SWBs. A first version of these protocols has already been published [53]. Gaining further insight into these protocols will allow us to comprehend the communication flow within the system.

Lastly, the chapter will discuss the development of the MuTRiG data path. The MuTRiG data path, which was developed during the course of the thesis work and published by the author [90], will be presented, providing valuable information and insight into its implementation.

6.1 Front-end board

Figure 6.1 provides a picture of the FEB, featuring its core element, the Intel ArriaV [103] FPGA. The ArriaV FPGA plays a critical role, acting as a bridge between the electrical readout of the detector through LVDS links and the optical connection through SAMTEC FireFlies [104] to the SWB. Accompanying the ArriaV FPGA is the Intel Max10 [105] FPGA, responsible for the ArriaV configuration and diligent monitoring of board temperature and voltage. The configuration of the ArriaV FPGA is stored in flash memory. The board employs two clock chips that are driven by the global 125 MHz system clock. These chips ensure a precise clock signal is delivered to both the FPGAs and the detectors.

Efficient power delivery and data management are ensured through a connection to a backplane, which supplies power to the board and facilitates the readout of voltage and temperature information. To power the board with electrical power direct current (DC)-DC buck converters are used. Commercial buck converters use inductors with a ferrite core. Since the board is located inside the 1 T Mu3e magnet, the inductor would be saturated and thus inhibit the function of the converters. Therefore, air coils are used instead, which do not suffer from this effect. Additionally, for debugging purposes, both FPGAs can be programmed via two Joint Test Action Group (JTAG) connectors.

6.2 Hit-time sorter

The hit sorter plays a crucial role in the data processing of the Mu3e system. The first version of the hit sorter was designed for the MuPix7 and MuPix8 chips [106], while the final version was developed for the MuPix10 and MuPix11 chips [107]. In Figure 6.2 a sketch of the final pixel hit sorter is shown.




Figure 6.2: Sketch of the hit sorter.

The final pixel hit sorter was designed to accommodate the readout of twelve MuPix chips with three links each for the inner vertex detector, as well as 36 MuPix chips with one link each for the outer layers. The MuPix protocol requires four cycles of a 125 MHz clock to transmit one hit. To handle this, a 3-to-1 multiplexer is used for the inner layers, and a 4-to-1 multiplexer is used for the outer layers, reducing the number of inputs to the sorter to twelve or nine, respectively.

For each input, the sorter uses a memory with 16 slots for each of the 2048 timestamps of the eleven-bit time counter of the MuPix chip. The timestamp information is stored in the memory addresses, with only 21 bits required, consisting of 8 bits for the column, 8 bits for row, and 5 bits for ToT.

To track the number of hits written per timestamp, a counter memory is used, counting all the hits written to the storage memory. During operation, a dynamic time window is used to define the

acceptable timestamp range. Hits within this window are stored in the Random-Access Memory (RAM), while hits outside the range are rejected and counted using an out-of-range counter. The position in time of the window can be adjusted during data taking to prevent loss of hits. If there are more than 16 hits per timestamp, another RAM is used to count the timestamps that have overflowed, providing this information DS in the readout chain. In addition, the number of hits per timestamp is counted and stored in counter memory.

Once the time window is closed, the hits can be read from memory. A sequence of hits to be read out is generated using the counters within the counter memories. To avoid falling behind in the reading process, a credit system is employed. If the reading process becomes too slow, timestamps may be skipped and marked as overflowed. This overflow information is transmitted after every 16 timestamps are sent.

For the tile and SciFi detectors, the sorter principle remains the same, with the number of timestamps increased to 8192. The data stored in the sorter increases to a maximum of 25 bits, and the number of inputs is reduced to three for the tile detector and two for the SciFi detector.

6.3 Front-end board communication protocols

Communication protocols are essential for successful data transmission between devices. In the context of the Mu3e experiment, several protocols were developed to facilitate communication with the FEBs, clock and reset systems, and to configure the detectors.

To ensure effective data transmission and configuration between the FEBs and the SWBs, a specific protocol was developed, which is discussed in Section 6.4. This protocol enables data reception and detector configuration between the FEBs and the SWBs.

Additionally, a protocol was devised for run-control communication with the clock and reset system, as described in Section 6.5. This protocol allows for coordinated control between the FEBs and the clock and reset system during the experiment.

6.4 Communication front-end boards to switching boards

A common communication protocol was designed to merge detector data and configurations during the readout process. This protocol is important for dynamic firmware design and efficient data transfer, considering the large amount of data involved in the Mu3e experiment. The specification of this common protocol, which facilitates communication between the FEBs and the SWBs, is detailed below.

Each data word consists of 32 bits, with the bits transmitted by the FPGA firmware sent through the link in little-endian order (the least significant bit first). All packets begin with a preamble that includes the packet type and the K-symbol K28.5, and end with the K-symbol K28.4. The MuPix data protocol is described in Section 6.4.1, while the formats for scintillating tile and SciFi detectors are presented in Sections 6.4.2.

For detector configuration, a slow control protocol (Section 6.4.3) is used, and the run control signals are covered in Section 6.4.4. Finally, the idle state is explained in Section 6.4.5.

In the following discussions, the bit counting is performed setting the least significant bit to 0 and the most significant bit to 31.

6.4.1 MuPix communication protocol

The communication protocol for the transmission of MuPix data follows a specific structure and is shown in Table 6.1. In the preamble, the least significant byte contains the comma word. The preamble (also referred as start of package (SOP)) begins with an identification pattern (111010) indicating the type of package being sent, followed by the FEB FPGA ID (bits 23 to 8).

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0					
Type	-	FPGA ID		K28.5	} preamble
ts (47:16)					
ts (15:0)		package counter			} debug word
-	SUB counter		hit counter		
-	send ts counter (30:0)				
ts (11:4)		overflow		K23.7	} SUB
ts (3:0)	chipID	row	col	tot	} hit
-				K28.4	} trailer

Table 6.1: Structure of the MuPix data protocol.

Next, the MuPix Data Header is sent, which includes a 48-bit FPGA timestamp, providing sufficient run-time coverage of approximately 3 hours. The second word of the data header contains a package counter (bits 16 to 31), keeping track of the number of packages sent during the run. Two debug words are sent to cross-check the number of sub-headers (SUBs), the number of hits sent and the time when the package was transmitted from the sorter. This is useful information to debug the protocol while in development.

Afterwards the SUB is sent, which contains the indicator (K-symbol K23.7), the bits 11 to 4 of the FPGA timestamp and the 16 overflow indication bits, indicating which timestamps experienced overflow during sorting. By adding SUBs the hit size can be kept to be 32 bit leading to a more efficient data format. Following SUB, the hits are transmitted. For each SUB, all hits within the lower 16 timestamp are sent. Each hit consists of chipID (5 bits), column and row information (8 bits), and ToT information (5 bits).

The SUB and hits iteration continues until the eleven bits of the SUB timestamp overflow. The next SUBs and hits are then sent using the same format.

Finally, the package is closed by sending the K-symbol K28.4, indicating the end of the transmission (also referred as end of package (EOP)). By adhering to this communication protocol, the MuPix data can be reliably transmitted between the FEBs and the SWBs, allowing accurate analysis and processing of the acquired data.

6.4.2 Tile and SciFi communication protocol

Similarly to the MuPix communication protocol, the Tile and SciFi data transmission protocol has been defined. The difference is only in the hit information, which is shown in Table 6.2. The preamble for these data types is identified by the pattern (111000).

Bits	Value
31:28	ts (3:0)
27:22	ASIC ID
12:17	channel ID
16:14	ts-remainder
14:10	fine counter
8:0	energy

Table 6.2: Structure of the MuTRiG hit data send from the FEB to the SWB.

For a SciFi hit the ASIC ID only needs five bits, which is sufficient to encode the 96 MuTRiG chips on the four links foreseen in the SWB to Receiving board connection. On the other hand, the tile hit needs the full six bits for the ASIC ID to be able to encode the 182 MuTRiG chips on the four links. Furthermore, both hits contain a channel ID (5 bits), a coarse counter (4 bits) running at 8 ns, a 3 bit coarse counter running at 1.6 ns and a fine counter (5 bits). Since the SciFi detector uses the short hit mode of the MuTRiG, the energy-flag bit is only sent occasionally. To encode if the hit has an energy-flag, the energy bits are set to 0x1FF otherwise they are set to 0x000. The energy-flag is used to mark hits where the time information comes from the second timestamp. The tile detector is using the long hit mode, and therefore the second timestamp is sent immediately afterward. This allows one to calculate a ToT information (energy 9 bits). In Section 6.6 the firmware used for this subtraction is explained in more detail.

6.4.3 Slow control communication protocol

The preamble of the slow control data communication protocol contains the pattern (000111), which serves as an indication of the type of slow control operation. After the preamble, the "start address" is sent. This address refers to the actual register inside the RAM of the FPGA where the user wants to read or write. The slow control settings for write-command and read-command are stored in this RAM or registers, and the addresses assigned for write and read consecutive operations are implemented as first in first out memorys (FIFOs) or consecutive addresses in the RAM. The next command in the protocol is "length", which indicates the number of consecutive registers after the start address to which the user wants to read or write. In the case of a write-command, the data for these registers will be sent following the length field. Finally, the package is closed with a trailer word, marking the end of slow control communication. In Table 6.3 the whole package is defined. For more detailed information on how slow control packets are processed within the FPGA, please refer to the reference provided in the thesis document [53].

When pattern (00) is used, it means an incrementing read command. In this case, the receiver reads the number of addresses specified in the "length" field, starting from the "start address". On the other hand, if a non-incrementing read command is intended, pattern (10) is used. Here, the receiver reads from the "start address" multiple times, and the number of reads is determined by the value in the "length" field. When pattern (01) is used, it represents an incrementing write command. In

this case, the receiver writes the payload (data) to the "start address" and the following registers. The size of the payload must match the value specified in the "length" field. Lastly, pattern (11) is used for a non-incrementing write command. The receiver consecutively writes the payload to the "start address". Similarly to the previous cases, the "length" field must match the number of data words in the packet.

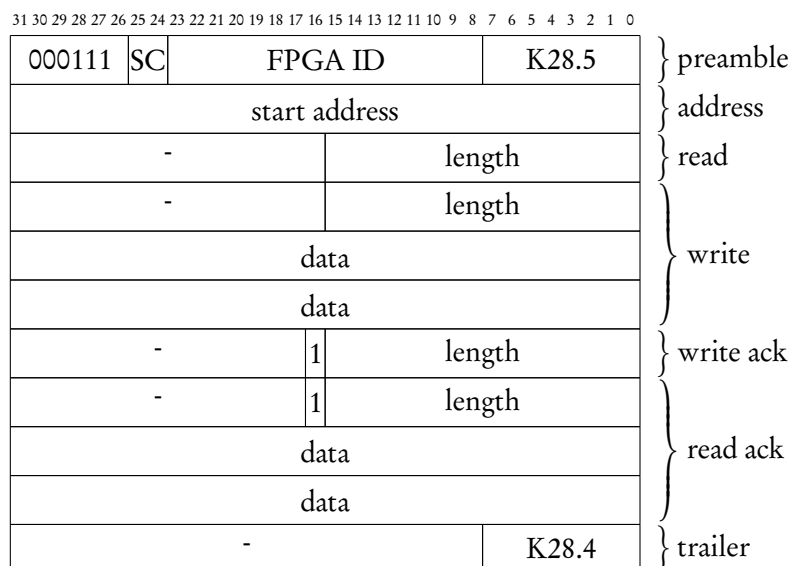


Table 6.3: Structure of the Slow Control packet. The preamble, address, and trailer parts are always fixed. Depending on the command (write, read, write acknowledgement or read acknowledgement) the payload part is different.

Upon receipt of a slow control request, the FEB that is responsible for handling the request acknowledges it by sending a response. The response is structured according to the same protocol scheme as the preamble and start address. In the case of an incrementing or non-incrementing read, the response consists of the preamble, start address, 16 zeros, a 16-bit length field, the payload (which contains the data for each address), and finally the trailer. On the other hand, an acknowledgement for an incrementing or non-incrementing write consists of the preamble, start address, 15 zeroes, a 1-bit acknowledge field, a 16-bit length field, and the trailer.

The acknowledgement mechanism ensures that the sender receives confirmation of completion of the requested slow control operation, allowing proper synchronisation and error handling in the communication between the FEBs and the SWBs.

6.4.4 Run control signals

To acknowledge the reset command sent by the clock and reset system to the FEBs, the run control signals are utilised. These signals are transmitted from the FEBs to the SWBs and then read out by MIDAS.

The run control signals are denoted by the comma word K30.7, which is reserved specifically for these signals. In Table 6.4, the structure of the run control signals is presented.

31 30 29 28 27 26 25 24 23 22 21 20 19 18 17 16 15 14 13 12 11 10 9 8 7 6 5 4 3 2 1 0	run control signal	K30.7
---	--------------------	-------

Table 6.4: Structure of run control signals.

Run control signals are sent as a single word containing the appropriate run command. The specific run commands are further detailed in the reference document [108]. The purpose of these commands is to control the execution of the run, including starting, stopping, pausing, or configuring the DAQ system.

6.4.5 Idle state

When there are no actual data to be transmitted, idle packages are sent instead. These idle packages serve several purposes within the communication system.

Firstly, idle packages consist solely of the K-symbol K28.5. By sending these words, byte alignment is achieved. Byte alignment ensures that the bits within the transmitted data are correctly aligned within bytes, following the defined communication protocol. This alignment is crucial for proper interpretation and processing of the data at the receiving end.

Secondly, the idle packages cause a change in the states of the state machines running on the different FPGAs. By receiving the idle packages, the state machines are informed to enter an idle state and not perform any calculations or data processing operations. This helps to minimise unnecessary computational overhead when there is no actual data to be handled.

To achieve byte alignment, an entity is implemented after the transceiver. This entity is responsible for shifting the bytes of the 32-bit word, adjusting the position of the K-symbol K28.5 according to the predefined communication protocol. By ensuring proper byte alignment, the transmitted data can be correctly interpreted and processed by the receiving devices.

6.5 Communication of clock and reset system

As described in Section 5.2, the Mu3e DAQ system utilises a specialised clock and reset system. This system enables communication with all FEB through a dedicated optical reset link operating at a speed of 1.25 Gbit/s. The link employs 8 bit/10 bit decoding for 8-bit word transmission. During the idle state, the K-symbol K28.5 is transmitted. The commands corresponding to various run states can be found in Table 6.5. To initiate a normal run sequence starting from the idle state, the following steps are required:

1. By sending the command **Run Prepare**, the run number is distributed to all FEBs, and it also initiates the transition to the **Run Prepare** state. In this state, various components such as FIFOs and buffers are cleared. In addition, an active signal is transmitted to the MIDAS system through the SWB. Once all active signals have been received, the next transition can be initiated.
2. The subsequent signal is **Sync**, which triggers a reset for all detector ASICs, ensuring that they remain in this state for at least one complete timestamp cycle.

3. Upon issuing the **Start Run** command, the reset is triggered, allowing the detector ASICs to resume normal operation. The timestamp counter begins counting from 0, and the data is subsequently transmitted to the SWB.
4. To conclude a run, the **End Run** signal is sent. Subsequently, the system transitions to the terminating state and no new data is transmitted from the FEBs. Following the last data package, the run trailer is sent.

Command	Code	Payload
Run Prepare	0x10	32 bit run number
Sync	0x11	-
Start Run	0x12	-
End Run	0x13	-
Abort Run	0x14	-
Start Link Test	0x20	-
Stop Link Test	0x21	-
Start Sync Test	0x24	-
Stop Sync Test	0x25	-
Test Sync	0x26	-
Reset	0x30	16 bit mask
Stop Reset	0x31	16 bit mask
Enable	0x32	-
Disable	0x33	-
Address	0x40	16 bit address

Table 6.5: Reset link protocol.

To abort a run, the **Run abort** command can be sent, which triggers a transition to idle state. Apart from these states, there are additional options available. The **Start Link Test** command initiates a Bit Error Rate Test (BERT) on all links except the reset link, while the **Stop Link Test** command stops the ongoing BERT. When the **Start Sync Test** signal is sent, the logic on the FEBs begins measuring the phase of the test signals received over the reset link. This measurement can be stopped by sending the **Stop Sync Test** command. The **Reset** signal, combined with mask bits in the payload, allows specific parts of the system to be reset and can be released by using the **Stop Reset** command. The combination of the command **Enable** or **Disable** command and the **Address** command enables or disables specific parts of the DAQ system. On the firmware side, all these functionalities are implemented within the state controller entity located on the common part of the FEB. A detailed description of how this is implemented in the firmware can be found in the Master Thesis of M. Müller [108]. The specification of the entire system is provided in [109, 110, 111]. The software component for controlling these transitions is implemented in MIDAS [112].

6.6 The MuTRiG datapath

In this thesis, the datapath of the SciFi detector was developed and tested. A schematic representation of the datapath is shown in Figure 6.3. The presented datapath was previously published by the author [90]. The main tasks of the datapath are to maintain the 32-bit size and to sort the hits within

a 8 ns resolution. This is necessary because the MuPix has an 8 ns resolution, while the MuTRiG chip generates a timestamp with a much finer binning of 50 ps.

The coarse counter of the MuTRiG ASIC is implemented as a 15-stage linear shift register, optimised for operation at 625 MHz [79] and it counts up to $2^{15} - 1$. The counting is achieved by changing the stages of the shift register, represented by pseudorandom values (e.g., 0x1234, 0x5678, 0x4242). These values can be assigned to binary values (e.g., 0x1234 = 0, 0x5678 = 1, 0x4242 = 2) using a lookup RAM on the FEB (PRBS T/E block).

In the case of the SciFi detector, two SMBs each containing four MuTRiG ASICs are connected to a FEB via a specific Detector Adaptor Board (DAB). Once the data are received in the Arria V FPGA, the hit data is unpacked and organised into a detector-specific record type called Rec1. The data lines from the different ASIC are merged into groups of two for further processing. A dual-port RAM for the timestamp-mapping (PRBS T/E block) is used to process the two data streams simultaneously. To ensure that each hit is sorted with 8 ns resolution, the coarse counter from the MuTRiG needs to be divided into 8 ns bins before sorting the hits in time.




Figure 6.3: Sketch of the FEB datapath of the SciFi detector. The E-T part is only needed for the Tile detector but was tested in this thesis with the SciFi datapath.

In the first step the coarse counter needs to be corrected to count up to 2^{15} . Therefore, a lapse correction block (Lapse CC) is used. This block compares the coarse counter values with the arrival time on the FEB to determine the number of lapses. A waveform diagram of the lapse correction process is shown in Figure 6.4 (for visualisation a counter with two bit width is shown). After this counter value correction, the 625 MHz counter needs to be divided by five to obtain a part that follows the 125 MHz clock and a remainder part.

In the tile case, the MuTRiG is configured to run in long mode. Therefore, each hit will have an additional 15 bit of the falling edge of the signal (E). Since the tile datapath only uses one sorter the dual-port RAM for the timestamp-mapping (PRBS T/E block) can convert the E and T parts in

parallel. To further calculate the energy information or, more precisely, a ToT, the first timestamp is subtracted from the second (E-T block). Performing the E-T operation before the sorter allows to reduce the hit size by limiting the maximum ToT. This is mandatory to be able to fit the sorter RAMs on the Arria V.

In the SciFi case, the MuTRiG is configured to run in short mode. The second timestamp is sometimes sent with a delay to the first timestamp. In this mode, the second timestamp is sent out with an energy-flag to calculate the energy information offline. The E-T part in Figure 6.1 is transparent in this mode and only the energy-flag is packaged with the hit information as described in Section 6.4.2. Following all the necessary corrections, the various hits can be sorted and transmitted to the SWB via up to two 6.25 Gbit/s optical links.

The datapath for the tile detector employs the same blocks, with the only difference being the number of ASICs, which changes from 8 to 13. For further information on the MuPix datapath and the common components of the FEB, refer to the work of [113].




Figure 6.4: Waveform of the lapse correction with a two bit width to visualise the principle.

Switching Board

The upcoming chapter will dive into the firmware components of the SWB in the Mu3e system. This chapter aims to provide a comprehensive explanation of the datapath, hit-time-alignment, and detector configuration that have been developed as part of this thesis.

7.1 Overview of the Switching Board




Figure 7.1: Picture of the PCIe40 Board which was developed for the ALICE and LHCb upgrades [114, 115]. Inside the Mu3e DAQ system it is called SWB.

Figure 7.1 shows the SWB, with its central component being the Arria10 [116] FPGA board covered by a passive heat sink for efficient cooling. It facilitates optical connections between the inputs from the FEB to the optical outputs to the Receiving board. The board holds a remarkable amount of 50 optical channels, provided by eight MiniPods [117] and two Enhanced Small Form-factor Pluggable (SFP+) connectors [118]. Complementing the Arria10 FPGA, a Max10 [105] FPGA is utilised to monitor voltage and temperature and a Max V [119] FPGA is used for flash programming. To streamline operations, the entire board is housed within a PC and seamlessly connected via two Peripheral Component Interconnect Express (PCIe) interfaces, facilitating detector configuration and efficient data readout via DMA.

7.2 Data flow of the Switching Board

Table 7.1: Optical fibre cabling from the FEBs to the SWBs and from the SWBs to the Receiving boards.

Sub-Region	Central Pixel Detector	SciFi Detector	Recurl Station US	Recurl Station DS
Input SWB	$36 \times 6.25 \text{ Gbit/s}$	$12 \times 6.25 \text{ Gbit/s}$	$33 \times 6.25 \text{ Gbit/s}$	$33 \times 6.25 \text{ Gbit/s}$
Output SWB	$8 \times 10 \text{ Gbit/s}$	$4 \times 10 \text{ Gbit/s}$	$2 \times 10 \text{ Gbit/s}$	$2 \times 10 \text{ Gbit/s}$

The role of the SWB is to collect data from the FEBs and merge the various data streams for the subsequent layer of the system, known as the Receiving board. Additionally, the SWB is responsible for transmitting the detector configuration to the individual FEBs (for more details, see Section 7.4) and for reading parts of the data for quality checks. Table 7.1 provides an overview of the different input and output links associated with the four SWBs. In particular, the output links of the recurl SWBs are divided into tile and pixel data.

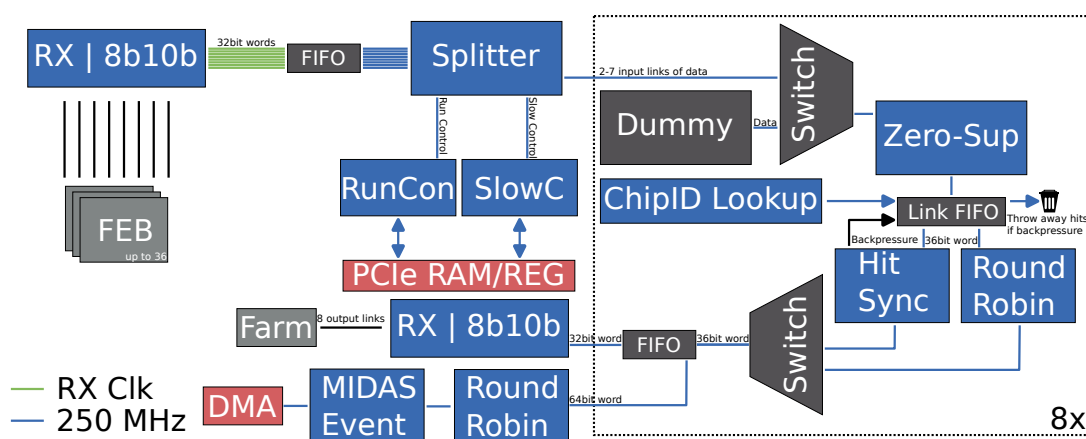


Figure 7.2: Sketch of the SWB datapath of the central pixel detector.

Figure 7.2 illustrates the various components of the SWB firmware. The figure depicts the data path designed for the central pixel detector, featuring $36 \times 6.25 \text{ Gbit/s}$ input links and $8 \times 10 \text{ Gbit/s}$ output links. The firmware is built modular and can be adapted for the other three SWBs by configuring the input and output link mapping accordingly.

In addition to the SWB-specific firmware, the transceiver (RX | 8b10b) and the PCIe blocks (depicted in red) are shared components between the SWBs and the Receiving boards. The fundamental functionalities of these blocks are elaborated in Section 3.1.

Once the hits from the 36 input links are received, the clock domain is adjusted from 156.25 MHz to 250 MHz. This transition is needed since the output links to the Receiving boards has a data rate of 10 Gbit/s with 8 bit/10 bit decoding and the firmware is designed to have 32 bit words. As discussed in Chapter 6, there are three distinct protocols employed for run control, slow control, and data transmission. In Figure 7.2, the splitter block segregates the three protocols and routes them to their respective firmware blocks. Run and slow control blocks are managed through PCIe memory and

Switching Board

registers. For a detailed explanation of the run control flow within the entire DAQ system, refer to [113].

Table 7.2: Input mapping of the different detector regions of the central pixel SWB. The layers are split in upstream (US) and downstream (DS).

Pixel Layers	US Layer 0	US Layer 1	US Layer 2	US Layer 3	DS Layer 0	DS Layer 1	DS Layer 2	DS Layer 3
# Input Links SWB	2	3	6	7	2	3	6	7
# Chips per Half-Ladder	3	3	8	9	3	3	9	9
# Half Ladder per FEB	4	3/4	4	4	4	3/4	4	4
# Input Links per FEB	36	27/30	32	36	36	27/30	36	36
# Links per MuPix	3	3	1	1	3	3	1	1
# MuPix Chips Total	24	30	192	252	24	30	216	252

After the splitter, the data is assigned to different regions of the central pixel detector, as outlined in Table 7.2. These regions are processed by eight data path blocks, each block corresponding to an output link. The number of input links utilised by each block varies depending on the type of SWB.

The initial block on the data path is a dummy generator, which can generate synthetic data for system tests. The following component is a zero suppression block, which removes all empty SUB and sorter packages when no hits are present. With this the overhead of the protocol can be reduced, which is necessary during detector commissioning or tuning when the overall hit rate is reduced. Furthermore, when only parts of the filter farm selection are operational, the zero suppression block can be used to reduce the data rate.

In the subsequent part of the system, all data are stored in a FIFO to enable buffering during readout from subsequent stages. In addition, a local chip mapping is performed from 6 bit to 8 bit. This conversion is necessary since the 6 bit chipID of the MuPix communication protocol is insufficient to store information for up to 252 chips in the data path responsible for handling hits from layer 3. In this stage of the data path, backpressure from subsequent components is managed by discarding hits if not all of the data can be processed. Consequently, hits under the current SUB are discarded and marked as overflow.

The next steps in the data path involve a hit-time-alignment. Two modes can be employed: time-alignment of all input links or a simple round-robin readout. The time-alignment mode will be discussed in Section 7.3. The round-robin readout is utilised for debugging or system testing purposes, while time-alignment is employed during data taking.

At the output of the data path, the 36 bit hit is converted into a 32 bit word for transmission to the Receiving board. For debugging and histogramming, the hit is converted to the 64 bit format used on the Receiving board, and another round-robin entity is used to read out the different data paths. These hits are then packed into MIDAS events and transmitted via DMA for analysis using the Mu3e Online Analyzer (see Section 8.5). The MIDAS event format is discussed in Section 8.3. Next, the hit-time-alignment entity will be explained.

7.3 Hit-time-alignment tree

The author of this work published the hit-time-alignment firmware algorithm in [90]. The main task of the SWB is to perform time-alignment on the different data streams. Time-alignment of the hit

data is essential to have the hits from the central pixel detector at the same time present to perform full track reconstruction on the GPU.

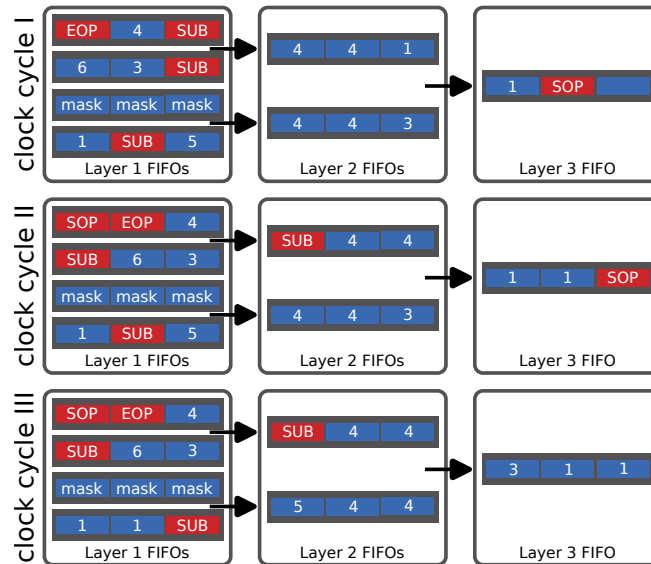


Figure 7.3: Example of a four-to-one time-alignment firmware on the SWB using pairwise comparison in a tree, taken from [90]. The different markers are start of package (SOP), end of package (EOP) and sub-header (SUB) as defined in the Table 6.1 of Section 6.4.1.

In Figure 7.3, a schematic of the hit-time-alignment firmware for the SWB is presented. The entire state machine can be found in Figure B.3 of the Appendix B.2. The numbers on the left indicate three example steps (clock cycles) of the algorithm. The SOP and EOP markers signify the start and end of a data package sent from the FEBS. The SUB marker represents the subheader, which contains additional timestamp information. The values shown between the markers indicate the hit-times within a SUB.

The alignment firmware combines the various packages into a unified stream while keeping the hit-times sorted. Each data package from the different FEBS contains hit data sorted into 8 ns bins and every 16 μ s a new package is transmitted. The data package from the FEBS is designed to hold the complete hit information using 32 bit. To maintain a size of 32 bit, only four bits are allocated for the time information (in Figure 7.3, the white numbers within the blue rectangles represent the time information of the 32 bit hit word). Consequently, 128 SUBs are sent, containing the upper seven bits of the global timestamp, providing a total of 2^{11} possible timestamps, when the four bits of the hit-time overflow.

The time-alignment firmware utilises pairwise comparison of hit-timestamps in a tree-based architecture. Each layer of the tree consists of two input streams and one output stream. The streams are buffered in FIFO queues. The example in Figure 7.3 demonstrates a queue size of three, while in the final system the queue size is of the order of 2^{10} . At each layer of the tree, both streams are buffered until each contains an SOP, and then the hit with the lowest timestamp is forwarded to the next layer. If a stream contains a SUB, hits from the other stream are forwarded until both streams contain the same SUB. Since the FEB always sends the 128 SUBs for each package, even in cases where there are no hits for those timestamps, the firmware only needs to compare the four bits of individual hits. This

process is replicated throughout the tree until the EOP marker is reached. Figure 7.3 shows three steps (clock cycles) of a four-to-one tree as an example, while the actual implementation employs an eight-to-one tree following the same principle.

As described above, the firmware combines eight input streams into one output stream. Since there is only one clock frequency change from 125 MHz (frequency of ASICs) to 250 MHz (frequency of the SWB firmware), the time-alignment firmware becomes a bottleneck with an effective ratio of $4 - 1$. As discussed in Section 5.4, the estimated average data rate for the central pixel detector is around 48 Gbit/s. The central pixel detector SWB has 36 input links from the FEBs and eight output links to the Receiving board. Thus, the potential throughput of the time-alignment firmware for the central pixel detector SWB is approximately 80 Gbit/s. This reduces the time-alignment firmware bottleneck to approximately $2 - 1$ for the central pixel detector. This bottleneck is really rare, since it assumes full saturation of all 36 input links. However, for the vertex detector where the chips see the highest rate, this can occasionally happen. Therefore, all input links are buffered to encounter such rate bursts.

For the other parts of the detector, the overall bandwidth requirements are even more relaxed. The fibre detector produces an overall data rate of 24 Gbit/s, and each SWB processes this data with four 10 Gbit/s links to the Receiving board. The two stations of the tile detector generate an overall rate of 6 Gbit/s per SWB, while each recursive SWB has a 10 Gbit/s link to the PCIe board for the tile detector.

7.4 Detector configuration

To be able to configure the different detectors, the slow control protocol, defined in Section 6.4.3, is used. The configuration principle was developed during the Master's Thesis of the author [53]. In this section a short summary of the configuration scheme for the timing detectors is given. The configuration of the pixel detector is explained in more detail in [113].

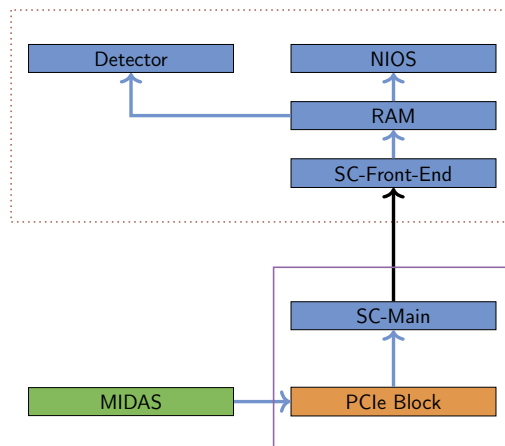


Figure 7.4: Design blocks for the slow control communication from the SWB to the FEB. The dotted line marks the firmware inside the FEB, the normal line inside the SWB.

The communication from the SWB to the FEB is solely for slow control purposes. In this setup, the MIDAS system transmits a slow control command through PCIe to the SWB. On the SWB,

the SC-Main entity receives this command and transmits it to the corresponding FEB via an optical link. The SC-Front-End entity on the FEB then unpacks this command and sends it to a RAM (see Figure 7.4). The RAM addresses are mapped either to the NIOS embedded processor or to firmware-specific components such as FIFOs or registers.




Figure 7.5: Design blocks for writing a slow control command from MIDAS to the FEB.

The design for sending slow control commands from MIDAS to the FEB is illustrated in Figure 7.5. First, the slow control protocol is written into the write memory of the PCIe interface. A start marker is then written before the preamble address. When the SC-Main entity detects the start marker, it begins reading out the slow control package and sends it to the transceiver, which transfers it to the FEB. On the FEB, the SC-Front-End decodes the package.

In the case of a write command, the entity starts at the address specified in the protocol and writes the data to a RAM. This RAM is divided into three sections. The first section contains 64 k of 32 bit words. The second part consists of 256 32 bit words directly connected to the firmware components. The last region also contains 256 32 bit words. The first address in this region, known as **cmdlen**, triggers an interrupt to the NIOS soft-core processor when written, and the second address, **offset**, is then read by the NIOS. This process can be seen as a Remote Procedure Call (RPC). In the first half of the **cmdlen** word, one can specify the function to be executed on the NIOS. For example, setting it to 0x0101 would prompt the NIOS to read the specified number of words from the first part of the RAM, starting from the address set in **offset**. This approach can be used to store the detector configuration for the Tile module board (TMB) or SMB in a specific RAM region and configure the detectors. Using this protocol, different control sequences can be set up and controlled from MIDAS.

After the slow control package is processed by the SC-Front-End entity, an acknowledgement is sent back to the switching board. The SC-Secondary, on the SWB, decodes this acknowledgement, and the package is written in the read memory of the PCIe interface. The MIDAS front-end (FE) continuously monitors this memory and reads new packages. As this entire process involves multiple devices, the software running on the PC needs to wait for acknowledgement before proceeding. The next chapter will provide a description of data processing on the PCIe board.

Receiving board and Filter Farm

The chapter presents a detailed exploration of the firmware components within the Receiving board. The objective is to fully explain the data path and the seamless integration of the GPU and the Mu3e Online Analyzer. It concludes with a brief overview of the online event selection.

8.1 Receiving board

Figure 8.1 displays the Receiving board, featuring the central component, the Arria10 [116] FPGA. To facilitate communication, the board hosts four quad small-form-factor pluggable ports (QSFP) connectors, providing a total of 16 optical input and output links. Positioned on either side of the Arria10 FPGA, two DDR SDRAM slots are available, offering the potential to equip the board with up to 4 Gbyte of RAM. For seamless integration, the entire board is enclosed within a PC and connected through a PCIe interface, enabling board control and data readout via DMA. The board is a commercial development board produced by Terasic [120].




Figure 8.1: Picture of the Receiving board [120]. In the picture the two DDR SDRAM slots are not equipped.

8.2 Data flow of the Filter Farm

After processing the data, the four SWBs, it is sent via 16 10 Gbit/s links to the first Receiving board within the filter farm. As a result, the initial Receiving board in the filter farm receives the complete 100 Gbit/s detector data from the four SWBs. Since all data are synchronised and sorted to 8 ns, the board receives a continuous stream of time-sorted data.

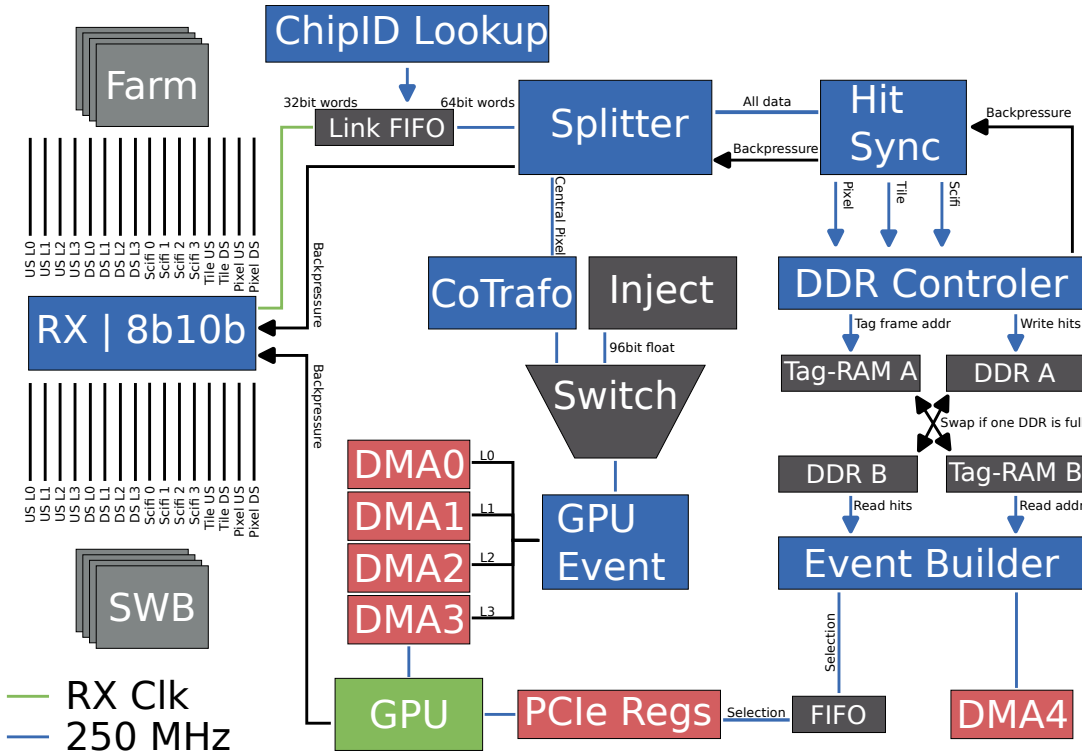


Figure 8.2: Sketch of the Receiving board datapath.

Figure 8.2 illustrates a schematic of the various firmware blocks developed during this thesis for the Receiving board. Starting with the common transceiver (RX | 8b10b), which was also used on the SWBs, this block sends the decoded 32 bit word to the first buffer (Link FIFO). In this process, the 32 bit stream is converted into 64 bit words by mapping the local chipID information to global chipIDs.

After passing through the first buffer, the hits from the central pixel detector undergo conversion into global 32 bit floating point values of x , y , and z , utilising a total of 96 bits. This conversion is crucial to enable efficient tracking on the GPU. In Figure 8.2, the component responsible for this conversion is denoted as CoTrafo (the original firmware was developed by S. Corrodi and adapted to work in the final firmware). Additionally, for comprehensive system testing, an injection entity (Inject) was developed to receive simulated data.

To perform the coordinate transformation, the chipID is used as an address to access the global corner value (row=0, column=0) of the corresponding pixel chip in a RAM-based lookup-table. With knowledge of the corner value of each chip and the column (col) and row information, the global

position x, y, z of the i th chip can be calculated using the equation:

$$\vec{h}_i = \vec{s}_i + \text{col}_i \cdot \vec{c}_i + \text{row}_i \cdot \vec{r}_i, \quad (8.1)$$

where \vec{h}_i represents the hit position vector (x, y, z) , \vec{s}_i is the global corner position, and \vec{c}_i and \vec{r}_i are the column and row directions, respectively.




Figure 8.3: Sketch of the transformed hits located in the PC RAM.

Following this transformation, the hit data is organised into frames of 8 ns and packed in 256 bits to be processed efficiently by the DMA engine operating at 256 bits \times 250 MHz. Hits within each GPU frame are sorted in time for each of the four layers and stored in a 0.5 Mbyte sub-package per layer, within the overall 2 Mbyte package (see Figure 8.3). References to the 8 ns borders of the hits are transmitted at the end of each sub-package. Each layer is readout via one DMA engine, which is explained in Section 8.4.

By employing this packaging approach, the GPU can dynamically create time frames in multiples of 8 ns. Having 8 ns level references allows the inclusion of hits from neighbouring frames ($\pm n \times 8$ ns) without redundant copying. These overlaps are essential for tracking using all four layers [68, 121].

During the GPU selection (the online selection will be explained in more detail in Section 8.6 of this chapter), the data is stored on one of the two onboard memory interfaces. When the buffer is full, the data is forwarded to the next Receiving board in the daisy chain configuration. As the entire farm consists of twelve farm servers, each only needs to process $1/12$ of the total 100 Gbit/s rate.

After the selection a reference of the selected events is sent back to the Receiving board using PCIe registers and stored in a request FIFO to trigger the readout of the data buffer. By utilising two DDR SDRAMs, one can be used to buffer the data while the other can be read out.

8.3 MIDAS event builder

In the following the event structure for the selected events is explained. The overall MIDAS event structure can be seen in [122].




Table 8.1: Structure of the farm MIDAS event header, adapted from [122].

Each MIDAS event is accompanied by a 16-byte header, which serves to differentiate between different types of events. The Event Header, presented in Table 8.1, has the following structure:

- **Event ID:** The event ID indicates the type of event. Each FPGA board on the farm will have the same ID.
- **Trigger Mask:** The trigger mask allows one to specify the subtype of an event. Each farm server has a specific trigger mask.
- **Serial Number:** The serial number starts at 0 for each run and is incremented by the event builder on the Receiving board.
- **Timestamp:** The event builder writes the time stamp before reading an event, providing the time in 8 ns intervals since the run began.
- **Event Data Size (Bytes):** The event data size represents the size of the event in bytes, excluding the header.




Table 8.2: Structure of the farm MIDAS Bank Header, adapted from [122].

Following the Event Header, a global Bank Header, as depicted in Table 8.2, is written. The global Bank Header has the following structure:

- **All Bank Size (Bytes):** Represents the size, in bytes, of the subsequent data along with the size of the bank header itself.
- **Flags:** The four last significant bit (LSB) (0:3) indicate the version of the bank structures, currently set as 1, used for endian detection (byte ordering). The fifth bit indicates that the banks are 32-bit banks (Bank Size being 4 bytes long), while the sixth bit indicates that the banks are 64-bit aligned using the BANK32A bank header.

For all events built on the FPGA boards, the BANK32A type is used. On the Receiving board four banks are written, while on the SWB one bank is written. Each bank has an individual bank header (BANK32A) that contains a four-character name that uniquely identifies the detector type. The bank header, illustrated in Table 8.3, comprises the following components:

- Bank Name: This four-character identifier is used to distinguish the type of detector associated with the bank.
- Type: The bank's type specifies the data type of the subsequent data. For all events, 64 bit unsigned integers are used.
- Bank Size (Bytes): Indicates the size, in bytes, of the following data, along with the size of the bank header itself.
- Reserved: This field is not utilised.




Table 8.3: Structure of the MIDAS BANK32A, adapted from [122].

Furthermore, the data area of each bank must have a size that is a multiple of 8 and 32 bytes. The first requirement ensures that the next bank header is aligned on an 8-byte limit in memory, allowing faster access by the Central Processing Unit (CPU). The 32-byte alignment is needed since the DMA engines uses a 256 bit width.




Figure 8.4: State machine for the MIDAS event builder firmware on the Receiving board.

Figure 8.4 depicts the state machine for the MIDAS event builder. In each step, the state machine writes either the hit data or the protocol information to a RAM on the Receiving board to enable event and bank size counting.

IDLE state: The state machine remains in the IDLE state until a selection is triggered from the GPU. Upon selecting the trigger, the state machine waits until the first hit of the selected event is read back from the DDR SDRAM.

Header State: In this state, the event header and the all bank header information are written to memory.

BaH0, BaH1 and Data States: The bank header and bank data are written in these states. The state machine transitions from Data over SBSIZE back to BaH0, iterating until a bank for each subdetector (Pixel, Tile, SciFi) is written out.

SBSIZE State: The size of the bank is written to the address stored in the BaH1 state.

DEBG State: In the DEBG state, the state machine writes the bank with the debug information.

ALIGN State: The state machine enters the ALIGN state and fills the event size until it becomes a multiple of 256 bit.

SABSIZE and SESIZE: Using the stored addresses, the state machine writes out the entire bank size (SABSIZE) and the event size (SESIZE).

END State: Finally, the state machine returns to the IDLE state and awaits the next selection.

8.4 DMA readout




Figure 8.5: State machine for the DMA readout firmware on the Receiving board.

The entire event is constructed around the selected hits inside the Receiving board and sent to the PC RAM via DMA. This transfer requires a triggered DMA engine. Each of the five DMA streams is controlled by the same state machine, as shown in Figure 8.5. Here's how the state machine operates:

WAIT State: The state machine waits until the software requests data.

DATA State: If the RAM contains data and not all requested data are read out, the state machine enters the DATA state. This state is necessary because the RAM requires two cycles to provide the stored data. For selected hit readout, the serial number of the MIDAS event is incremented and set in the SER state. However, for the four engines reading out the GPU events, this state is skipped. It is important to check if the RAM has enough data to fill the entire frame, and if not, the state machine

goes back to WAIT or enters the LAST state.

RUN State: In the RUN state, all the data from the RAM is read out. If there are no more data, the state machine transitions to LAST if all requested words are read and no more data remains in the RAM. Otherwise, it goes back to WAITING.

LAST State: In the LAST state, the last address of the last valid data word is sent out via PCIe, and the state machine enters the PAD state.

PAD State: In the PAD state, 4 kbyte of data is written to trigger a DMA readout cycle, and the "done" signal is set to one.

SKIP State: In situations where no request is started, and the state machine has data, the events are skipped.

Since the Receiving board features five DMA readout engines, the driver must provide them with different PC RAM regions. For the four engines responsible for the GPU event, the RAM is sliced into four 0.5 Mbyte blocks, one for each engine. These blocks are repeated with an offset of 2 Mbyte until the full requested RAM size is filled. On the other hand, a single continuous block of RAM is utilised for the selected data readout. The use of fair round-robin readout enables each engine to write 4 kbyte of the data, and the next engine starts to write 4 kbyte afterwards. Once the data is transferred to MIDAS, it can be analysed using the Mu3e Online Analyzer, which was developed during this thesis and a first version was explained in [90].

8.5 Mu3e Online Analyzer

The Mu3e Online Analyzer software was developed to analyse data quality and perform tracking. It is built on the MIDAS-compatible analyzer framework [92], allowing seamless integration into the MIDAS system to receive data from various detectors.

To provide interactive ROOT-like graphics in web browsers, the software utilises the THttpServer class of ROOT combined with JavaScript ROOT [123]. This enables convenient visualisation and exploration of the data. For debugging and performance testing, the system can be fed with simulation data from the Mu3e software [93].




Figure 8.6: Sketch of the Mu3e Online Analyzer framework, taken from [90].

Figure 8.6 illustrates the different steps involved in the Analyzer's operation. The first part of the process involves connecting to MIDAS and checking for the presence of detector events. The data banks within the events are first decoded and sent to a queue, where different cleaning and clustering steps are executed for each detector type. This is done using a parallel readout pipeline. Online tracking is performed using Mu3e Reconstruction software [93], generating tracks that are written to a ROOT file [123]. The tracks can be used to compare the GPU selection. The reconstructed tracks are then displayed on an event display to provide a live view of specific events [124]. At the same time, the raw data is buffered using a second queue to produce data quality plots (DataQuality) of the system.

For commissioning, test stands, or readout of the data from the SWB the Analyzer is able to remove noise and/or sort the data between different detector modules. Following the cleaning process, the software can also be used to reconstructed cosmic tracks in the detector. The first tests of this framework were performed during the Mu3e Cosmic Run 2022 (see Chapter 12).

8.6 Online event selection

As explained in Section 8.2 detector data need to be significantly reduced for permanent storage. Only event candidates with physical relevance are retained. In the context of Mu3e particle rates, requiring three concurrent tracks in time alone does not sufficiently decrease the data rate. As a result, no hardware triggers are used. Instead, the online filter farm reconstructs all tracks and applies a software-based selection algorithm. This selection requires the presence of three tracks that occur simultaneously, originating from a common vertex, and exhibiting the expected kinematic properties of signal events. To handle the computational demands of this task, GPUs are employed, benefiting from the rapid technological advancements in the gaming and deep learning markets.

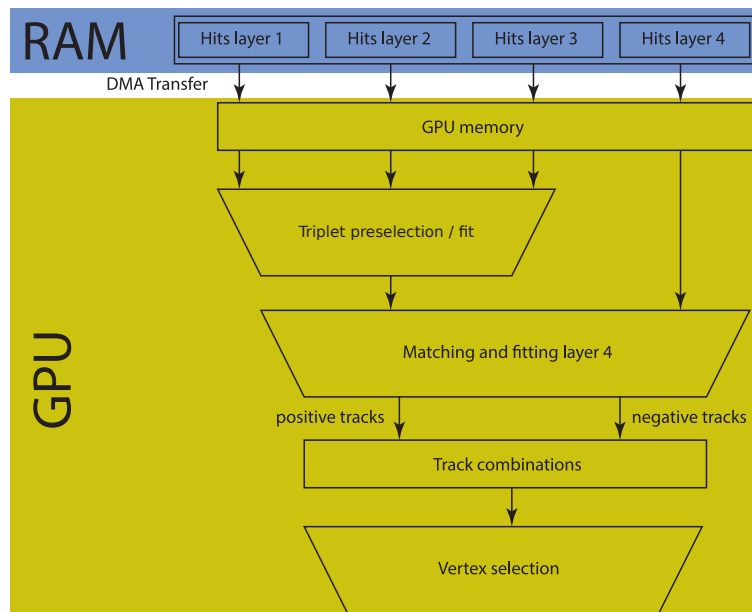


Figure 8.7: Flow diagram of the online reconstruction software after the GPU events are transferred to the PC RAM, taken and adapted from [51].

For a fast track fitting process, a simplified version of the fast linear fit based on multiple scattering (details in Chapter 19 of the Technical Design Report of the Mu3e experiment [51]) is implemented on the GPUs. Additionally, events containing at least two positive and one negative electron tracks are inspected for common vertex and signal-like kinematics. This selection is executed on individual farm servers frame by frame and the chosen frames are merged into the global data flow. Figure 8.7 shows an overview of the GPU part of the data flow after the events are transferred to the PC RAM as explained in Section 8.2.

During online reconstruction, only hits from the central station of the pixel detector are considered. This choice is made because matching the recoiling tracks and the time information of the tiles and

fibres is computationally expensive and unnecessary for the initial selection process. Triplets are formed by combining hits from the first three detector layers. Before the actual fitting procedure, a set of basic geometric selection criteria (triplet preselection) is applied to reduce the number of combinations by a factor of approximately 50.

The fitting of triplets is non-recursive and linearised, making it suitable for parallelisation on GPUs. With numerous computing kernels but limited memory capacity, GPUs excel at tasks involving repetitive computations on the same memory content. Assuming a muon rate of 2.466×10^8 Hz and 64 ns time frames, approximately 10 hits per layer are expected. This leads to combinations of the order of 10^3 . With a code optimised for these conditions [68, 125, 121], a rate of 2.3 billion fits per second on an NVIDIA Geforce GTX 3080Ti was measured [126]. This performance is sufficient to handle the expected combinations.

For each triplet that meets the criteria of χ^2 and radius, the track is extrapolated to the fourth detector layer. If at least one hit exists within a specific transverse radius and z-distance, the closest hit to the extrapolated position is used to form a second triplet from hits in layers two, three, and four, which is then fitted. An improved curvature value for the track is derived from averaging the results of the two triplet fits. Subsequently, the particle's charge is determined from the track curvature, and all combinations of two positive tracks and one negative track are examined for a common vertex.

The vertex position is calculated by averaging the intersections of the tracks in the transverse plane (perpendicular to the magnetic field), weighted by the uncertainty arising from multiple scattering in the first layer and hit resolution. A χ^2 -like variable is defined based on the distances of the closest approach of each track to the mean intersection position and their uncertainties, both in the transverse and r-z plane. Vertices are selected on the basis of their proximity to the target and the χ^2 value. In addition, criteria are applied for the total kinetic energy and combined momentum of the three tracks at the closest approach points. After all these cuts, the frame rate is reduced by approximately a factor of 200 [121, 68].

In addition to identifying signal candidate events, cosmic ray muon candidates and random frames are saved for calibration, alignment, and studies related to selection efficiency. The parameters of all reconstructed tracks are histogrammed for monitoring purposes and for searches, such as for two-body muon decays [127].

The triplet fit, propagation to the fourth layer, vertex fit, and monitoring have been implemented and optimised for performance. Extensive testing has been carried out on NVIDIA Geforce GTX 1080Ti and 3080Ti cards, demonstrating that 12 of these cards are sufficient for the Phase I detector [68, 121, 125].

Part III

Detector Integration

This part discusses the various testbeam campaigns conducted to validate the integration of the detectors into the data acquisition (DAQ) system. Accomplishing these tasks required the collaboration of multiple individuals. The author actively participated in all of the tests described in the subsequent chapters and performed tasks such as configuring the readout system, preparing the experimental setup, and analysing the collected data.

First Detector Integration

The first complete detector integration test of the Mu3e experiment occurred in February 2020 at the Deutsches Elektronen-Synchrotron (DESY) testbeam facility [128]. The primary objective of this test was to operate all three detectors simultaneously with a shared readout system and to observe correlations between them. This section provides an explanation of the testbeam facility, followed by a description of the detectors and the DAQ system used in the experiment. Furthermore, the data collected during this testbeam campaign are analysed and the results are presented.

The author participated in the test beam and was responsible for the DAQ system and data analysis. One of the primary objectives of the DAQ system integration was to develop firmware on the field programmable gate array (FPGA) for generating Maximum Integrated Data Acquisition System (MIDAS) events, as described in Chapter 8.

9.1 DESY testbeam facility

The DESY testbeam facility, located on the Hamburg-Bahrenfeld campus, operates within building 27, also known as 'Halle 2,' which is one of the experimental halls at DESY. This facility provides three independent beam lines capable of producing electron or positron particles with selectable momenta ranging from 1 GeV/c to 6 GeV/c. The beam lines are connected to the DESY II synchrotron, which typically generates electron beams with an oscillating energy between 0.45 GeV and 6.3 GeV. The electron/positron beam of the DESY II synchrotron hits a carbon-fibre target to generate a bremsstrahlung beam. This beam consists of photons that are then converted into electron/positron pairs by using a metal plate target. The beam is then horizontally expanded into a fan shape using a dipole magnet. Finally, a collimator is used to select a specific portion of this fan-shaped beam, creating the final beam for the different testbeam areas. It is worth noting that this testbeam facility is among the few worldwide that grants users access to multi-GeV beams. Consequently, it has the necessary infrastructure for the development, testing, and research and development (R&D) of nuclear and particle physics detectors.

9.2 DESY 2020 experimental setup

In Figure 9.1, a sketch of the setup is presented. For this particular test, a 4×4 tile matrix, a SciFi ribbon, and four MuPix8 chips were used. The 3 GeV electron beam first passed through the MuPix chips (upstream (US)), then through the SciFi detector, and finally reached the Tile matrix (downstream (DS)).

The timing detectors were read using the initial version of the Muon Timing Resolver including Gigabit-link (MuTRiG) chip, while the pixel detector used the MuPix8 chip. The MuPix8 chips had

the values of the rows in the x direction and the column in the y direction. Furthermore, the MuTRiG chip provided 32 readout channels for the SciFi ribbon in the y direction. Like the MuPix8 chips, the tile matrix is aligned in the x/y direction. All detectors were facing their active area in the US beam direction.

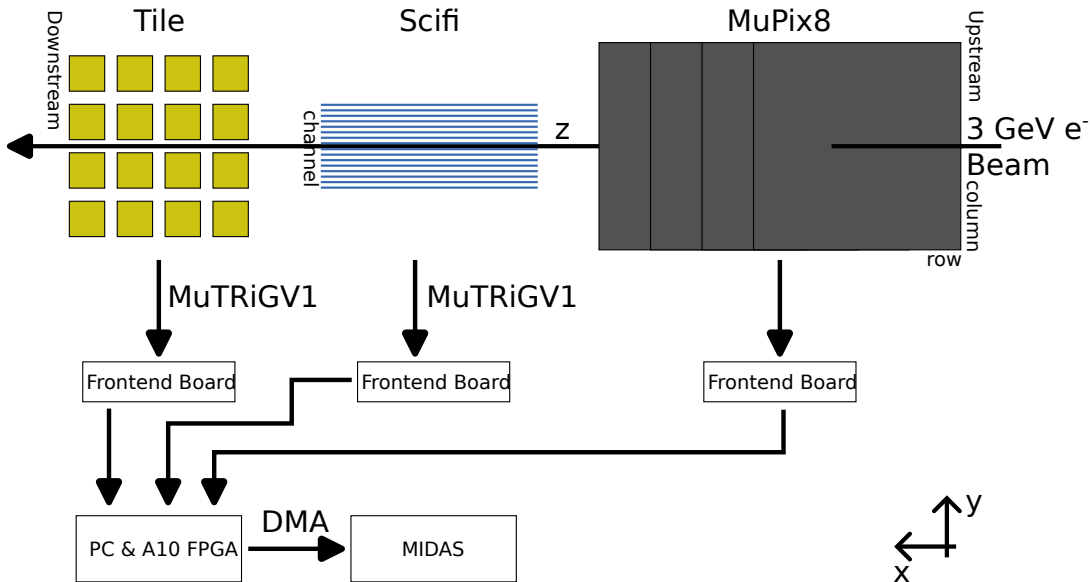


Figure 9.1: Sketch of the setup used for the DESY integration testbeam in February 2020.

To connect the detectors, four MuTRiG chips were used for the SciFi detector and one MuTRiG chip was used for the Tile detector. All three detectors were connected to three first-version Frontend boards (FEBs)¹ via low-voltage differential signaling (LVDS) links. The acquired data were then transmitted through optical links to a personal computer (PC) equipped with a PC interface board (Receiving board).

The MIDAS events built on the Receiving board encompassed three banks, one for each detector. Subsequently, the data was transferred from the Receiving board to the Random-Access Memory (RAM) of the PC using Direct Memory Access (DMA). On the PC, the data was inspected before being handed over to MIDAS, which stored the data on disk and controlled the entire setup. Throughout the testbeam, an overall data rate of 150 MB/s to disk was achieved. The absence of a data selection process and the capability of each FEB to transmit data at a rate of 6.25 Gbit/s resulted in the main limitation being the writing of data to disk. To address this, a backpressure was implemented to limit the rate of data written from the FPGA via DMA by discarding events in the firmware.

9.3 DESY 2020 testbeam results

Figure 9.2 displays the hitmaps of all four MuPix8 chips employed in the setup. The first MuPix8 chip in Figure 9.2 was placed US, while the fourth was placed DS in relation to the direction of the beam. The fourth chip was placed in front of the SciFi detector. The performance of the chips appears to be

¹As a predecessor to the FEBs described in Section 6.1 a prototype board was designed for firmware and lab tests.

satisfactory, with only a few hot pixels observed at the edges of the column for chip one (Figure 9.2a) and chip three (Figure 9.2c). However, the beam mostly hits the higher row values of the chips, since only a coarse space alignment was performed during the testbeam.

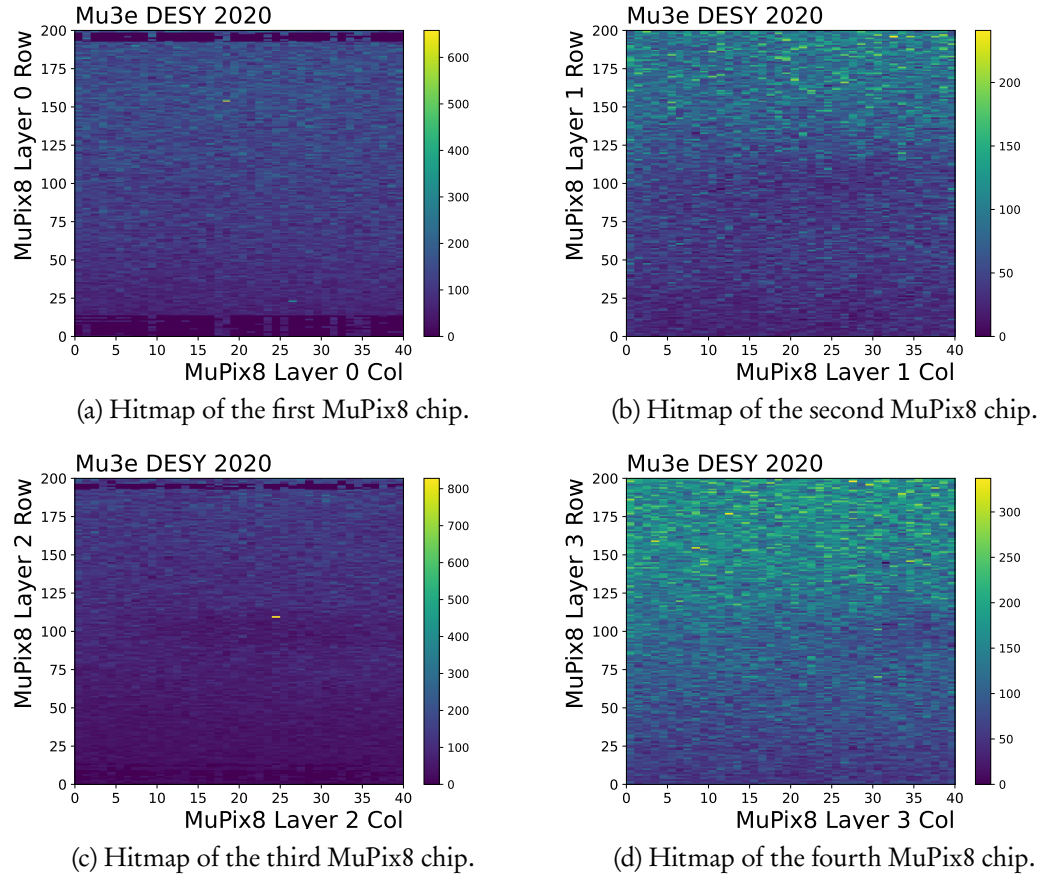


Figure 9.2: Hitmaps of the all four MuPix8 chips used in the setup.

Since the primary objective of the testbeam was to operate all three detectors simultaneously, no offline alignment was conducted. Instead, a manual adjustment of the chip positions was performed during the testbeam to achieve a basic alignment. Figure 9.3 illustrates the spatial correlations between the first and second chips². Both correlations, shown in Figure 9.3a for column-to-column and Figure 9.3b for row-to-row, indicate a favourable spatial alignment of the chips.

After the testbeam, two issues in the setup were identified only during the subsequent analysis of the data. Firstly, the MuTRiG chip used to read out the tile matrix was misconfigured, resulting in problems with the time information of the hits. Unfortunately, this issue could not be resolved by offline analysis. Consequently, the data collected during the beamtime could not be used to establish correlations with the other two detectors. Secondly, the SciFi ribbon was read only from one side, since the other side had some electrical issues, leading to a high noise rate because left-right coincidences could not be performed. However, utilising the information from the pixel chips during

²Further correlations between layer 1, 2 and 3 are show in Figure B.4 of the Appendix B.3.

offline analysis may assist in reducing the noise rate of the SciFi detector.

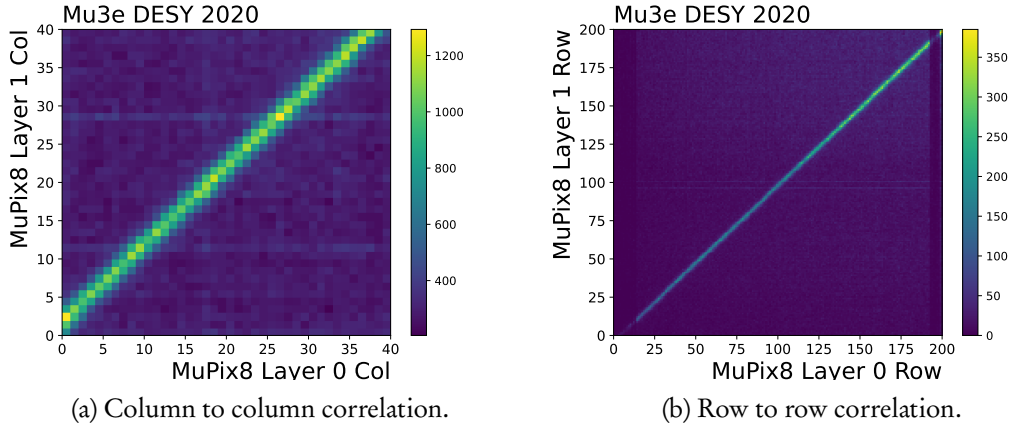


Figure 9.3: Space correlations of layer 0 and layer 1 of the MuPix8 chips.

The initial step of the analysis involved aligning the global time produced by the MuTRiG chip with the MuPix8 timestamp. Since the coarse counter correction for the MuTRiG chip, as explained in Chapter 6, was not implemented during the testbeam, an offline correction was applied. The MuTRiG coarse counter operates at 650 MHz and lapses after reaching $2^{15} - 1$ counts. The MuPix8 timestamp counter operates at 125 MHz and lapses after reaching 2^{15} counts. To achieve alignment between these different counters, the FPGA timestamp, running at a frequency of 125 MHz, was plotted against the full MuTRiG timestamp.




Figure 9.4: Left: Global time fit of the MuTRiG timestamp to correct for different counter lapsing in contrast to the MuPix8 chips. Right: Difference in the global timestamp of the third MuPix8 chip and the corrected global SciFi timestamp.

On the left side of Figure 9.4, a fit of one lap is displayed. The early lapsing of the coarse counter is evident here, as the MuTRiG timestamp restarts slightly shifted after reaching its maximum counter value. The multiple lines inside a lapsing batch represent the different MuTRiG chips used in the setup. The linear function fitted in the left part of Figure 9.4 exhibits a slope of 156.5 ± 4.4 , which is

consistent with the expected value of $8 \text{ ns}/50 \text{ ps} = 160$. By multiplying the timestamp of the SciFi detector by the fitted value, the lapsing can be corrected, effectively shifting it to 125 MHz.

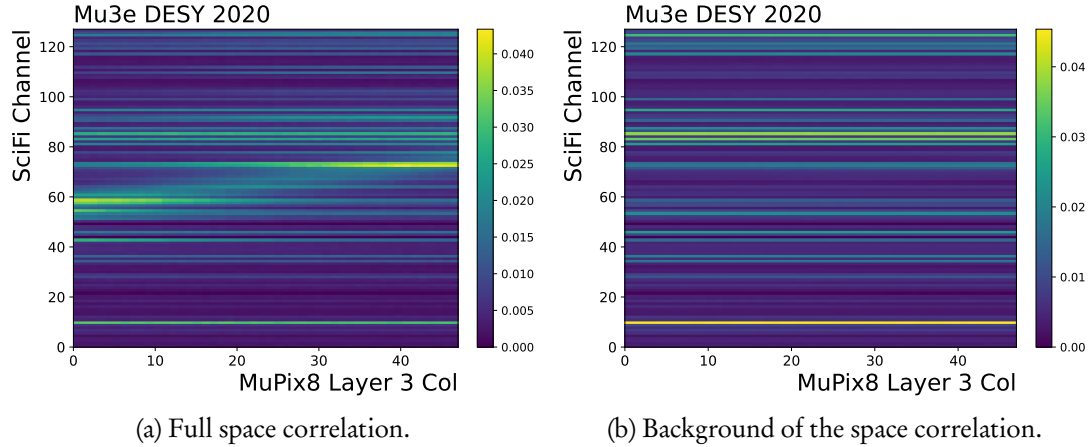


Figure 9.5: Figure 9.5a shows the complete space correlation between SciFi and MuPix. Figure 9.5b shows the background of the space correlation between SciFi and MuPix.

To mitigate noise on the SciFi ribbon, the timestamp of the MuPix8 layer, located in front of the SciFi ribbon, was subtracted from the corrected SciFi timestamp. The right part of Figure 9.4 illustrates the difference between the two timestamps. A clear peak is observed around $35 \mu\text{s}$. Furthermore, a smaller peak is visible around $70 \mu\text{s}$, which is an overflow effect of the $35 \mu\text{s}$ peak. The overflow occurs because the time difference between the two detectors was calculated by correlating all hits within a time window of $\pm 80 \mu\text{s}$, which led to the appearance of a second correlation with the particles in the next beam-bunch. The need for such a large search window was due to the detectors having different latency. The $35 \mu\text{s}$ peak can be used to establish a cut-off point to subtract the SciFi noise. Subsequent testbeam runs addressed this issue by employing the hit sorter, as described in Section 6.2, which can be adjusted to compensate for these latency offsets.




Figure 9.6: Substraction of full space correlation and background.

Figures 9.5a, 9.5b and 9.6 depict the normalised spatial correlations between the SciFi channel number and the MuPix8 row number of the pixel layer in front of the SciFi ribbon. In Figure 9.5a, the

normalised spatial correlation is presented for all data. Furthermore, Figure 9.5b shows the normalised spatial correlation for regions outside the red dotted line on the right of Figure 9.4. The difference between the two histograms is shown in Figure 9.6, which reveals a clear spatial correlation that becomes even more prominent after subtraction. The ≈ 15 channels of the active SciFi ribbons in the correlation match the dimensions of the MuPix8 chip. The MuPix8 chip has 48 columns, each with a size of $80\ \mu\text{m}$, while the SciFi ribbon is $32.5\ \text{mm}$ wide and covered by 128 channels. This leads to an area of $48 \times 80\ \mu\text{m} = 3.84\ \text{mm}$ for the MuPix8 chip and an area of $\frac{32.5\ \text{mm}}{128} \times 15 = 3.809\ \text{mm}$ for the SciFi ribbon.

In summary, the testbeam campaign demonstrated that the DAQ system is capable of reading the different detectors and correlations between two different types of detector can be observed. In addition, first time synchronisation studies were performed between the different detectors. However, since no specific detector tuning was performed and not all parts of the readout chain were fully functional, further testing was required to validate the system. In the following chapter, a more comprehensive time synchronisation study between different FEBs is conducted using MuPix8 chips in at the Mainz Microtron (MAMI) accelerator. The goal of the follow-up MAMI testbeam was to study the time synchronisation over multiple runs and to investigate the overall stability of the system.

10

Front-end Board Synchronisation

After conducting initial detector integration tests, it became necessary to assess the time synchronisation between different detectors. In [108], a jitter of less than 5 ps was measured between the various output clocks of the clock and reset system after optical transmission and re-conversion to differential electrical signals. This measurement easily satisfies the required 30 ps specification for the Mu3e experiment.




Figure 10.1: Floorplan of the MAMI facility.

In July 2020, a follow-up study was conducted, using the first version of the FEB, to evaluate time synchronisation under realistic conditions, closely similar to the Mu3e experiment. The following section provides an explanation of the MAMI testbeam facility, followed by a description of the experimental setup. The chapter concludes with an analysis of the testbeam data.

The author participated in the testbeam and was responsible for the setup of the DAQ system. The entire analysis of the testbeam data was carried out by the author.

10.1 MAMI testbeam facility

MAMI is a continuous wave microtron that generates a high-intensity, polarised electron beam with energies of up to 1.6 GeV. A floorplan of MAMI can be seen in Figure 10.1. The electron beam at MAMI can reach currents of up to 100 μA . The capability to achieve high rates was later utilised in Chapter 14 for conducting irradiation studies of the MuPix sensors.

MAMI serves as the core of an experimental facility for particle, nuclear, and X-ray radiation physics at Johannes Gutenberg University in Mainz. It is among the largest accelerator facilities on a university campus in Europe, supporting fundamental research. Experiments carried out at MAMI involve approximately 200 physicists from various countries who collaborate on international projects. In addition to physics experiments, the beam at MAMI can be used to investigate detector prototypes.

For the synchronisation tests, a 855 MeV electron beam was utilised after the microtron stage in Halle B and before the X1 experiment (see Figure 10.1).

10.2 MAMI 2020 experimental setup




Figure 10.2: Sketch of the setup used for the MAMI integration testbeam in July 2020.

In Figure 10.2, a schematic diagram of the setup is presented. The test involved the use of four MuPix8 chips [129]. An electron beam, with an energy of approximately 855 MeV, passed through these four

MuPix chips. The first two chips were connected to one FEB, while the last two chips were connected to a second FEB. Both FEBs were synchronised using the clock and reset system [109] described in Chapter 5.

Similar to the DESY testbeam in 2020, the detectors were linked through LVDS connections, and the data was transmitted via optical links to a PC hosting a Receiving board. Two experiments were carried out to test the synchronisation between the different detectors.

The first part of the experiment involved reconfiguring the MuPix8 chips at the start of each run¹ and subsequently resetting the entire system using the clock and reset system. In the second part, the entire system was reset at the beginning of each run, while the configuration of the chips remained unchanged. The first experiment should study whether reconfiguring the chip at run start has an influence on time synchronisation, while the second study studies the influence of reset behaviour of the DAQ system. Both tests should not change the time synchronisation between the detectors.

10.3 MAMI 2020 testbeam results

In Figure 10.3 the hitmaps of the first MuPix8 chips, each connected to a different FEB, used in the setup are displayed. In Figure B.5 of the Appendix B.4 the two remaining MuPix8 chips are shown. The first and second chips (chipID 0, 1) were placed US, while the last two chips (chipID 2, 3) were positioned DS with respect to the direction of the beam. Despite a few hot pixels, the relatively small beamspot of the MAMI beam is visible on all four hitmaps. It is important to note that no offline detector tuning or alignment was performed and that the setup was manually adjusted during operation.

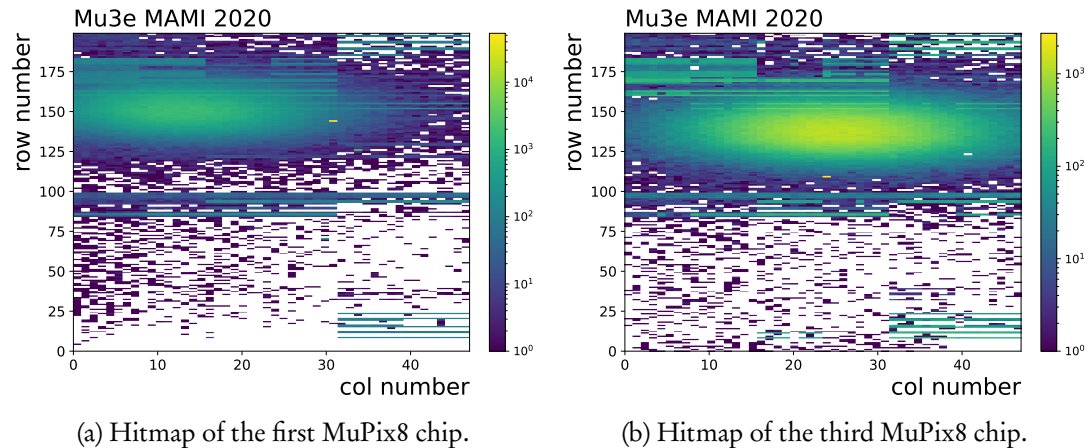


Figure 10.3: Hitmaps of the first MuPix8 chips of the two FEBs in the MAMI synchronisation setup.

In Figure 10.4a, the time correlation between chip 0 and chip 2 is displayed as an example. In addition to the timestamps around 200×8 ns and 800×8 ns, which are caused by noisy pixels generating the same timestamp, a clear correlation is observed.

¹The full run sequence of the Mu3e experiment is explained in Section 6.5.




Figure 10.4: Time correlations of chip 0 and chip 1.

To investigate the time correlation between chips connected to different FEBs, the difference in hit-time was calculated. Figure 10.4b illustrates the difference between the hit-time of chip 0 and chip 2. To search for the correction between the chips, two search windows per chip with $\pm 200 \times 8$ ns are created. Due to the double correlations within the two windows, an unintended triangular-shaped background was created. Consequently, the Gaussian fit function was extended by $f_{bg}(x) = c(d - |x|)$ to account for this background. The fit performed, excluding the timestamps around 200×8 ns and 800×8 ns, yielded a total time resolution of approximately 70 ns. That is, concerning that no Time-over-Threshold (ToT) corrections or other calibrations were performed, the expected time resolution of MuPix8 [129].




Figure 10.5: Time synchronisation with reconfiguration. The error bars represent the sigma values obtained from the Gaussian fit.

Figure 10.5 shows the time synchronisation between different chips over multiple runs. For each run, all chips were reconfigured. The error bars represent the sigma values obtained from the fit.

It can be observed that the time difference between the two chips falls within the range of 8 ns. Since the MuPix chips have a clock frequency of 125 MHz to sample the hit-timestamp, the granularity of the time of the chips has a lower limit of 8 ns. All runs fall within this 8 ns range, which satisfies the specification required for the Mu3e experiment.

However, one caveat of the tests is the behaviour of chip 3 after run 131. The mean time difference was lower than for the other runs. This is marked with a grey line only fitted on the runs after 131. This effect is not fully understood, but it is possible that a different configuration setting caused a change in its timestamp.




Figure 10.6: Time synchronisation without reconfiguration. The error bars represent the sigma values obtained from the Gaussian fit.

Figure 10.6 presents the time synchronisation between different chips over multiple runs without reconfiguration at run start. It can be seen that the time difference between the two chips remains within the range of 8 ns. Additionally, the strange behaviour of chip 3 after run 131, which was observed in the previous test, is no longer present in this test. This supports the hypothesis that the problem was caused by a faulty configuration for this specific run.

In addition, a linear function ($f(x) = a \cdot x + b$) was fitted to all tests, excluding the faulty runs of chip 3. For all tests, this yields an overall zero slope within the error bars. Only the runs of chip 1-2 from Figure 10.5 do show a negative slope. Again, this could be caused by a faulty configuration during these tests.

Overall, all tests demonstrate synchronisation between different detectors and FEBs using the clock and reset functions of the system as expected. Furthermore, the reset sequence of the DAQ system did not influence this synchronisation. This lays the foundation for the evaluation of a prototype of the Mu3e inner vertex detector in a larger integration run at Paul Scherrer Institute (PSI). The DAQ system was scaled to accommodate the readout of more than 100 MuPix10 chips and is described in Chapter 11.

Mu3e Integration Run 2021

During the Mu3e Integration Run 2021, detector prototypes were tested at the $\pi E5$ muon beam line at PSI. The purpose of this run was to operate the detectors under conditions as close as possible to the final Mu3e experiment. This involved cooling the inner vertex detector prototype with gaseous helium [130], operating the entire setup inside a 1 T solenoid magnet and using a muon beam. Furthermore, parts of the final DAQ system were used during this integration run.

In addition to the primary Mu3e integration tests, a secondary test was conducted towards the end of the beam time to investigate the use of the inner vertex detector for the muon spin rotation (μ SR) measurements. In Chapter 15 the design and test of a dedicated Si-Pixel-based spectrometer for μ SR measurements based on these initial tests will be discussed. This chapter builds on previously published articles by the author [131], as well as one by Rudzki [132], but it provides more detailed information and additional insights regarding the DAQ system compared to the published work. The author actively participated in the Mu3e Integration Run 2021 and was involved throughout the three-month period, including setting up the experiment at PSI, providing the DAQ system, and being part of the analysis process.

11.1 Mu3e Integration Run setup




Figure 11.1: Picture of the Mu3e Integration Run 2021 detector prototypes.

Figure 11.1 illustrates the setup used during the Mu3e Integration Run. In this configuration, the muon beam enters the detector setup from the left side. The central part of the setup consists of

the vertex detector and the SciFi detector, which are housed within a cylindrical support frame that surrounds the beam pipe and the target. The support frame incorporates various supply chains for cooling, electrical power, and optical readout cables. The two detectors are read out by the final version of the FEBs.

The detector configuration included the inner vertex detector prototype and two ribbons from the SciFi detector. The vertex detector consisted of two layers. Each layer was constructed from ladders with a total of six chips each. In addition, each ladder was divided into two half ladders that were read individually from left and right. Two different target configurations were used: the hollow double cone Mu3e Target (introduced in Section 4.2.2), and an additional target was inserted to investigate the use of the detector for μ SR measurements. This second target consisted of two discs. The first disc housed a scintillator that was used as an entrance counter, while the second disc contained the actual target. The target was composed of a silver probe placed inside a permanent magnet. It is important to note that the μ SR measurements were conducted without the application of an external magnetic field. However, it should be mentioned that since the test was performed after the Mu3e integration run, the presence of a residual magnetic field cannot be excluded.

11.2 Mu3e Integration Run DAQ system

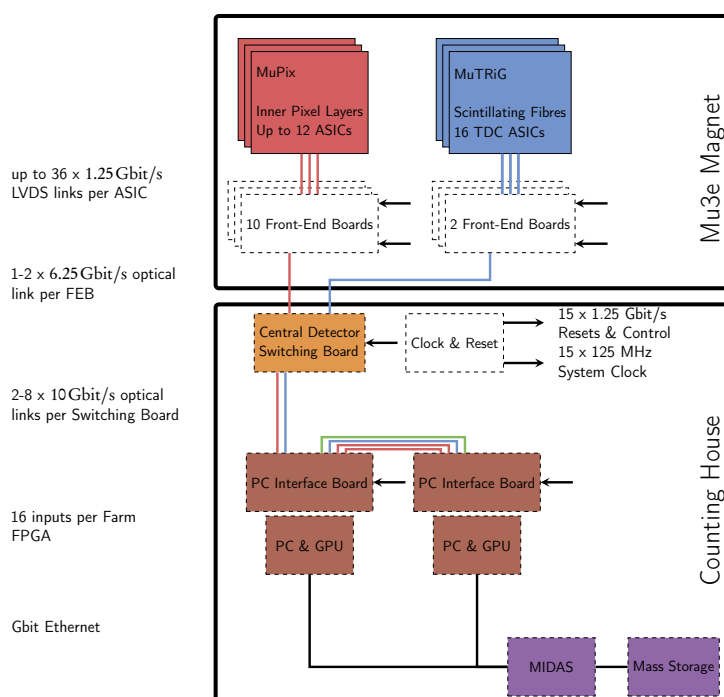


Figure 11.2: A sketch of the Mu3e DAQ system used in the Integration and Cosmic Run is shown.

Figure 11.2 provides a schematic representation of the DAQ system employed in the Mu3e Integration Run. As explained in Chapter 5 the DAQ system consists of three layers of FPGA boards. In the Mu3e Integration Run the first layer, which utilises the custom-developed FEB, was tested.

The FEB is positioned within the magnetic field and connected to the various subdetectors via electrical 1.25 Gbit/s LVDS. The primary function of the FEB is to read out and configure the different detectors and sort the received hits in time.

During the Mu3e Integration Run 2021 one of the main achievements was the sorting of the pixel detector hits in firmware on the FEBs (the sorter is explained in Section 6.2). Identical to the final DAQ system, the Mu3e Integration Run 2021 contained ten FEBs to read the inner vertex detector. On the other side, only two FEBs for the SciFi detector compared to the final twelve. Furthermore, there were no tiles and none of the outer pixel detectors were present. The first layer of the DAQ system was located inside the Mu3e magnet while the rest of the system was placed inside the Mu3e Counting House outside the experimental area.

The second layer comprises PCIe40 boards, which were originally designed for A Large Ion Collider Experiment (ALICE) and Large Hadron Collider beauty (LHCb) upgrades [133]. These boards are located outside the magnet and are connected to the FEBs via 6 Gbit/s optical links. All twelve FEBs are read out with only one Switching Board (SWB), which contains a third of the final links for the central pixel detector and a small fraction for the SciFi detector. The PCIe40 boards serve to combine and time-align the different data streams and facilitate the transmission of detector configurations to the FEBs through optical connections. The time-alignment of the different FEBs, developed in this work and explained in Section 7.3, was also tested for the first time during the Mu3e Integration Run 2021.




Figure 11.3: Chip to chip correlation between the different layers of the inner vertex detector using the Mu3e Target.

The third layer encompasses commercial Terasic DE5a-Net-DDR4 boards, installed in a PC equipped with a graphics processing unit (GPU). The DDR4-RAM onboard of these boards is used for data buffering, while hits from the two layers of the central detector are used for online event selection using the GPU. In the final system, 12 of these farm servers are daisy-chained to handle the total data rate of 100 Gbit/s. However, during the Integration Run, the GPU selection was not used and only some simple read out tests were performed.

To ensure synchronisation among the detectors, which was tested in the previous chapter, the clock and

reset system is employed. This system incorporates another FPGA board that generates a 125 MHz clock with synchronised reset signals for start and stop operations, which are distributed to the detectors.

11.3 Results of the Mu3e Integration Run 2021

To verify that all subsystems and the DAQ system work, various correlations are studied. Since the SciFi detector was not fully working during data collection, only the results for the vertex detector are shown.

Figure 11.3 shows the spatial correlations between the chips in layer zero and the chips of layer one. These correlations are caused by muon-decay electrons passing a working chip on both layers. However, it is evident from Figure 11.4 that a significant number of chips did not function properly. The green chips indicate fully operational ones, while the red chips experienced issues when applying high voltage (HV). Violet chips were excluded due to noise during the Mu3e measurements, and orange chips were excluded during the μ SR measurements. Grey chips were excluded during both sets of measurements. One of the main reasons why many detector chips were not functioning properly was the lack of individual chip testing. This issue was caused by time delays resulting from the COVID-19 pandemic, which affected the testing and preparation of the chips before they were mounted on ladders.




Figure 11.4: Working chips during the Mu3e Integration Run 2021. Ladder ID and z (chip) are used for the spatial chip position, while the chipIDs used for configuration are marked for each chip individually.

Due to these limitations, only a few correlations between the different layers are feasible. For the Mu3e measurements, chip-to-chip correlations can be observed for chip numbers ranging from 10 to 20 on both ladders.

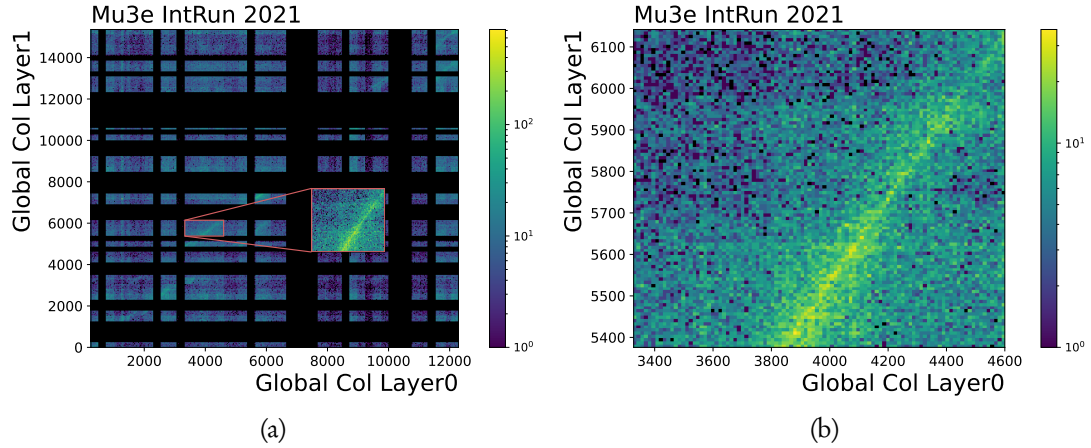


Figure 11.5: Column to column correlations of the two layers of the vertex detector prototype using the Mu3e Target are shown in Figure 11.5a. Figure 11.5b shows the zoomed in picture.

Figure 11.5a presents the column-to-column correlations between the two layers of the vertex detector prototype using the Mu3e Target¹. The positions of the pixel hits are indicated by the global column positions of the individual pixels on each layer² of the vertex detector. The correlation shown in the figure corresponds to the chip-to-chip correlation observed in Figure 11.3.

This correlation analysis provides information on the alignment and spatial relationship between the different layers of the prototype vertex detector. By examining the column-to-column correlations, one can observe the consistency and agreement between the positions of the pixel hits on each layer.

Figure 11.6b displays the time difference between two consecutive pixel hits in both layers of the prototype vertex detector using the Mu3e Target. Since the hits already showed spatial correlations, caused by the muon-decay electrons, the time correlation should be expected. Each pixel hit carries spatial and time information. In Figure 11.6a two more correlations can be seen in the corners of the graph. This is probably caused by an imperfect reset of the different FEBs. Furthermore, the correlation between all hits generates a Gaussian-shaped background in the time difference distribution (see Figure 11.6b).

To account for this background, a double Gaussian fit was performed, which allowed a better representation of the data. The Gaussian model of the "signal" shows a total time resolution of approximately 54 ns (see GaussianTop). It is important to note that due to the absence of a calibrated reference time, no ToT corrections could be applied to improve the time resolution.

Considering the lack of full calibration of the detector, low HV, etc., the measured time resolution is consistent with the results observed in various MuPix test beams [134]. However, to further validate

¹In Figure B.6 of the Append B.5 the row-to-row correlation is shown.

²Global column position is defined as the ladder ID \times 6 plus z (chip) \times 256. One can refer to Figure 11.4 for a visual representation of the definition of ladder ID and z (chip).

the integration of the SciFi detector and perform more comprehensive time resolution studies with the entire system, a second test run was conducted in 2022 and will be presented in Chapter 12. In this test, cosmic rays were utilised to investigate detector correlations and improve the overall understanding of the system's performance.

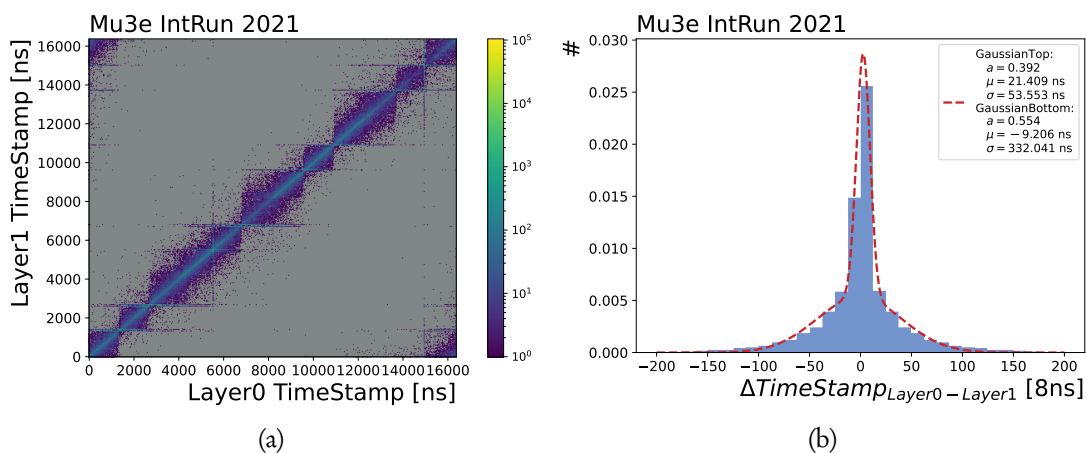


Figure 11.6: Time correlations of the two consecutive pixel hit from different layers using the Mu3e Target.

11.4 Results of the μ SR measurement




Figure 11.7: μ SR target with the first disc housing a scintillator and the second disc containing the actual μ SR target.

During the μ SR measurement, the Mu3e Target was replaced with a spacial μ SR target shown in Figure 11.7. Additionally, similar correlation studies were conducted on the prototype vertex detector. As a result of the experience gained from operating and calibrating the detector during the Mu3e Target measurements, improved results were anticipated for the μ SR measurement. However, it is

important to note that a bug in the readout firmware was discovered after data collection, which led to difficulties in reading the entrance counter. Unfortunately, this issue could not be resolved during the offline analysis, and the time information from the entrance counter could not be utilised for various timing studies or muon lifetime measurements. All of this made any μ SR measurement impossible. However, with the gained experience, a second test was conducted in 2023, which will be described in Chapter 15.




Figure 11.8: Chip to chip correlation between the different layers of the inner vertex detector using the μ SR target.

Figure 11.8 displays the chip-to-chip correlations of the Mu3e inner vertex detector. Compared to the Mu3e measurements, these correlations exhibit an even clearer correlation between the two layers. The strongest correlation is observed among the chip numbers between 10 and 20.




Figure 11.9: Column to column correlations of the two layers of the vertex detector prototype using the μ SR target are shown in Figure 11.9a. A zoomed-in version of the correlations is shown in Figure 11.9b.

Figure 11.9 presents the space correlations of the Mu3e inner vertex detector. With improved understanding of detector handling and various firmware improvements on the time-sorting side, the observed spatial corrections are less "noisy"³ than for the Mu3e Target runs. In contrast to the Mu3e Target, the μ SR target consists of two discs a few centimetres apart from each other, and the space correlation of the two layers of the vertex detector exhibits two lines. These lines correspond to the different discs of the target. For more detailed information about measurement and analysis, you can refer to the work of Thomas Rudzki [132].

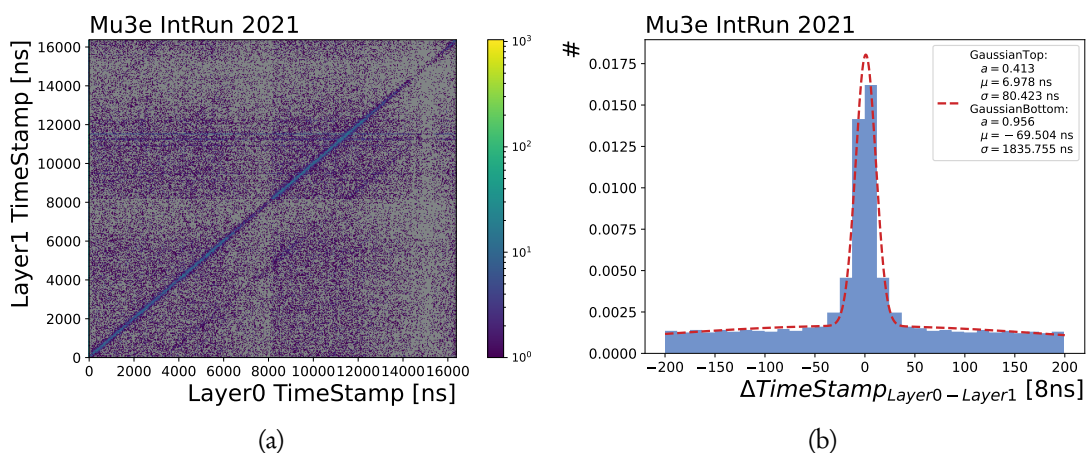


Figure 11.10: Time correlations of the two consecutive pixel hit from different layers using the μ SR target.

Figure 11.10 presents the difference between the hit-times of two consecutive pixel hits in the Mu3e inner vertex detector. The same fit function that was previously used yields a total time resolution of approximately 80 ns. The second correlation in Figure 11.10a is, as for the Mu3e Target runs, again present and most likely caused by a non-functional reset at the start of the run of some FEBs.

Considering these results, it became clear that the chosen detector geometry was not ideal for a first prototype of performing μ SR measurements. Therefore, a redesign of the detector was conducted, specifically for this purpose. The results of this redesign and a first measurement will be presented in Chapter 15.


³Most of the "noise" in the graphs was caused by bugs introduced on the firmware side and not by actual detector noise.

Mu3e Cosmic Run 2022

In the Mu3e Integration Run 2021, the focus was on integrating the inner vertex detector into the Mu3e DAQ system. This intensive system test aimed to ensure proper functioning and integration of the detector components. The details of this integration run are described in Chapter 11. It provided valuable insights and allowed for the refinement of the DAQ system.

The Mu3e Cosmic Run 2022 served as another important system test, where the integration of the inner vertex detector was finalised and the SciFi detector was integrated using cosmic rays. Unlike the Integration Run 2021, the detectors were not operated inside a magnetic field during this test. This allowed for a comprehensive evaluation of the DAQ system under different conditions. Parts of this chapter were published before by the author [90]. Beside the published work, the chapter provides a more detailed analysis of the integration process within the DAQ system. The author actively participated in the Mu3e Cosmic Run 2022, which involved setting up the experiment at PSI, providing the DAQ system, and contributing to the analysis process. A significant contribution by the author during this run was the successful integration of the SciFi detector, a task that was not achieved during the Mu3e Integration Run 2021.

12.1 Mu3e Cosmic Run setup



Scintillator panels

1 x SciFi ribbon

Inner pixel layers

Scintillator panels

Scintillator panels

Figure 12.1: Picture of the detectors used during the Mu3e Cosmic Run 2022.

Figure 12.1 provides an overview of the Mu3e Cosmic Run 2022 setup. On the left side of the figure,

a picture of the detector prototypes used in the cosmic run is shown. The inner vertex detector is enclosed with Kapton foil, which serves to distribute the flow of helium for cooling the detector. Additionally, a layer of SciFi ribbons is mounted directly above the pixel detector.

To identify cosmic ray muons, three scintillating panels are positioned around the detectors. These panels are used to detect the coincidence of signals, which serves as a cross-check for cosmic rays. However, it should be noted that the alignment of the three panels is not perfect, resulting in suboptimal coverage of the detector acceptance.

On the right side of Figure 12.1, a sketch of the setup is provided, illustrating the arrangement of the detectors and scintillating panels.




Figure 12.2: Working chips during the Mu3e Cosmic Run 2021.

In the evaluation of the prototype of the Mu3e inner vertex detector during the Integration Run at PSI, some changes were made to the setup. Specifically, the ladders of the inner vertex detector were re-shuffled to achieve optimal cosmic ray coverage. The best-performing ladders, as determined in the previous integration run (see Figure 11.4), were placed in a geometrically better position to detect cosmic rays.

During and after data collection, the same quality checks were performed as in the Mu3e Integration Run 2021 to assess the performance of the reshuffled chips. The results of this evaluation are shown in Figure 12.2. Chips that functioned properly without any problems are marked in green. Chips that experienced thermal runaway or had a high current when HV was applied are marked in orange. Chips marked in red encountered problems such as broken links, non-working LVDS connections, or noisy pixels. Chips marked in grey exhibited a combination of these issues.

As observed during the Mu3e Integration Run 2021, there were once again issues with some chips related to HV problems. Given that no new ladders were constructed for this run, and a re-shuffling of the existing components was carried out, it is probable that this issue resulted from the lack of individual chip testing prior to ladder production. As a result, even the best-working chips were not operated under design conditions, leading to reduced efficiency and timing resolution.

12.2 Mu3e Cosmic Run DAQ system

In addition to the setup described in Figure 11.2 and explained in Section 11.2, the system for the Mu3e Cosmic Run 2022 incorporated additional scintillator panels. These panels were used to create a cosmic ray coincidence using Nuclear Instrumentation Module (NIM) logic. The signal was digitised by an additional FEB and sent to the SWB.




Figure 12.3: Sketch of the Mu3e Online Analyzer framework used at the Mu3e Cosmic Run 2022.

Since the online selection for the GPUs was not available during this test, online reconstruction was performed using the Central Processing Unit (CPU). The Mu3e Online Analyzer framework, detailed in Section 8.5, was used for this purpose. The readout pipeline employed in this run is illustrated in Figure 12.3. The system was fully integrated into MIDAS and received events from all the different detectors.

In the initial step, the data banks within the events are decoded and sent to a queue, where various cleaning and clustering procedures are performed for each detector type. Online tracking can be carried out using Mu3e Reconstruction software [93]. The resulting tracks are written to a ROOT file [123]. Similarly, raw data are buffered using a second queue to generate data quality plots (DataQuality) for the system. Interactive ROOT-like graphics in web browsers are provided through the use of the THttpServer class of ROOT combined with JavaScript ROOT. For debugging and performance testing, the entire setup can be supplied with simulation data from the Mu3e Reconstruction software.

12.3 Data flow during the Mu3e Cosmic Run

In order to process the data from the optimised vertex detector, several steps were taken. First, it should be noted that only one side of the SciFi detector was operational due to a broken readout board.

Similarly to the first integration run, the hits from the detectors were collected on the SWB and stored into MIDAS banks. The readout software on the PC then transferred these data directly to MIDAS.

The analysis system connects to MIDAS and checks if the detector events are present. The events from the fibre and vertex detectors are decoded and sent to a queue for further processing. Despite




Figure 12.4: Sketch of the Mu3e Event Display, developed in [124].

threshold tuning and pixel masking, a few chips still had noisy pixels that could not be turned off in hardware. Therefore, these noisy pixels needed to be filtered out of the raw data in the software.

Since the fibre data path lacked a hardware sorter, the data had to be fine-sorted in software to correlate it with the pixel data. Additionally, the scintillator panels were treated as pixel chips with an unused chipID. Since the DAQ system was designed to read all data and write them to disk, these hits could be filtered out during the online and offline analysis. By performing these steps, the data from the vertex detector could be processed and used for further analysis and tracking.

After the cleaning process, a full track reconstruction was carried out using an adapted version of the Mu3e triplet reconstruction software [93]. This software was specifically tailored to reconstruct cosmic tracks in the pixel detector.

The reconstructed hits were then forwarded to an event display, which provided a live view of specific events. The event display allowed for visualising the reconstructed tracks in real-time. Figure 12.4 provides an example of such a display, showing a track reconstructed using the Mu3e triplet reconstruction software.

To enable general time correlations, the scintillator panels were utilised as cosmic tags. These tags were used to correlate fibre and pixel hits within a specific time window around the tags. This allowed for studying the time correlations between different detectors.

In addition to track reconstruction and event display, various quality plots were generated for the different detectors. These plots included hit rate plots, hitmaps, timestamp plots, and more. These simple quality graphs provided valuable information on the performance of the detectors and helped identify any issues or anomalies.

12.4 Timing studies during the Mu3e Cosmic Run

In the analysis of the time correlation between the inner vertex detector and the SciFi detector, only hits that coincided with the three scintillator panels were considered. The plot on the left of Figure 12.5 shows the ToT versus the time difference between the inner vertex detector and the scintillator panels.

The distribution exhibits an asymmetric shape, which is attributed to the time-walk effect. This effect arises from the fact that the comparator of each pixel cell has an absolute threshold on the rising edge

of the analogue signal to determine the hit-time. Consequently, signals with higher amplitude have an earlier time, while signals with lower amplitude have a later time. The difference between the two times is known as time-walk¹.




Figure 12.5: Figure 12.5a shows ToT versus time difference of the inner vertex detector and the scintillating panels. Figure 12.5b shows ToT versus time difference of the inner vertex detector and the scintillating panels with time-walk corrected pixel hit-time.

To mitigate this issue, the MuPix10 chips can be configured to utilize a second timestamp, sampled on the falling edge of the signal, to calculate the ToT (explained in Section 4.3.1) and correct the hit-time. However, during the Cosmic Run, the online correction was not used, necessitating an offline correction using the scintillator panels as a reference. The mean of each ToT value was calculated and subtracted from the sampled hit-time to perform this correction. The corrected time difference is shown in Figure 12.5b. It is important to note that the setting for the readout speed of the second timestamp was not individually optimised for each pixel chip during the test run. Consequently, an imperfect sampling of the second timestamp occurred, particularly visible in Figure 12.5a for hits in the lower left region of the distribution (where $t_{\text{pixel}} - t_{\text{trigger}} < 0$ and $\text{ToT} < 200$). In this region, the ToT calculation was not accurate, resulting in higher variance in the correction for these ToT values (see bottom of Figure 12.5b). Furthermore, the MuPix10 chip does not provide ToT information for columns higher than 236. Therefore, all hits with column addresses greater than 236 were excluded from the analysis. In contrast to the inner vertex detector, the SciFi detector was not configured to provide a second timestamp. Consequently, a time-walk correction was not feasible for this detector.

In Figure 12.6, the time difference between the inner vertex detector (after the time-walk correction) and the SciFi detector is shown. The distribution was fitted using a Gaussian function.

The performed fit yielded a total time resolution of approximately 98.5(27) ns. The errors are primarily determined by the timing resolution of the vertex detector, which is around 100 ns due to the untuned chips. Furthermore, sampling of the scintillator panels was performed at 8 ns on the FEB, limiting the time resolution of this detector. On the other hand, the scintillating timing detector should have a smaller uncertainty on its timing resolution compared to the resolution of the scintillator panels [135].

¹In Section 4.3.1 the 2-comparator threshold method is explained in more detail, which can be used to correct for this time-walk effect.




Figure 12.6: The time correlation between the inner vertex detector and the SciFi detector requiring a coincidence of the scintillator panels.

Despite the limitations and restrictions associated with the detectors, the presented results demonstrate the capability of the DAQ system to process data from different detectors, perform online time synchronisation and provide the hit data necessary to study the time resolutions between the detectors and track reconstruction.

To complete the integration of the scintillating tile detector into the DAQ system and finalise the construction of the complete Mu3e detector, an additional test beam campaign was carried out at DESY. This test beam used the Cosmic Run pixel ladders along with one tile module.

Tile Integration

To complete the integration of the tile detector into the Mu3e DAQ system, a second testbeam campaign was carried out at the DESY testbeam facility. This chapter provides a description of the setup used during the test, followed by the analysis performed to demonstrate the correlation between the tile detector and the pixel detectors.

The specific details of the setup, such as the number of tile modules and their placement, the connection to the FEBs, and the data readout process, are described in this chapter. Furthermore, any modifications or adjustments made to the DAQ system to accommodate the tile detector are explained.

The analysis section outlines the approach taken to establish the correlation between the tile detector and the pixel detectors. This may involve comparing hit patterns, spatial, and timing correlations. The results demonstrate the successful integration of the tile detector into the overall DAQ system and highlight the observed correlations between the tile and the pixel detectors.

The author actively participated in the testbeam, taking responsibility for the pixel detector and the DAQ system, and also contributed to the analysis process.




Figure 13.1: Sketch of the detectors used during the Tile Integration Run 2022.

13.1 Tile Integration Run setup

Figure 13.1 shows the setup used during the Tile Integration Run 2022. For this run, two tile modules, each containing 26 tile matrices, were affixed to a beam pipe replica, closely resembling the final configuration. These modules accommodate a total of 26 MuTRiGs, with each module being read out by a dedicated FEB.

Furthermore, to facilitate beam monitoring and establish correlations between pixel chips and tile modules, three half ladders from the Mu3e Integration / Cosmic Run vertex detector were used. These three half ladders were read out by a single FEB, and all the FEBs were interconnected with a Receiving board via optical cables.

The DAQ system utilised during the test was operated using MIDAS, and subsequent data analysis was performed using the Mu3e Analyzer Framework; see Section 8.5. Unfortunately, it should be noted that only one of the tile modules was fully operational during the test, since the second had an electrical problem and could not be repaired in time.

13.2 Rate monitoring tests

To perform data quality checks in the final Mu3e system, prescaling methods are necessary to handle the large volume of data. However, these methods must be verified during data collection by monitoring the data rate. Given that it is not feasible to read out all the data from a single tile module alone, firmware development was undertaken to measure the data rate. The DESY testbeam facility, with its relatively low electron rate of approximately 10 kHz, offers an ideal environment for studying such methods.




Figure 13.2: Rate map of one tile module. The hitmap shows the beamspot. Each bin in the histogram is one SiPM which is connected to one tile.

A single tile module consists of 13 MuTRiG chips, each with 32 readout channels, resulting in a total of 416 channels. While it may be possible to count the rate for each channel and store the data in registers within the FEB for these 416 channels, the task becomes nearly impossible when considering the pixel detector, which can have up to one million channels per FEB.

One approach is to focus on a specific portion of the detector for a fixed duration, count the hits and then transmit the counters to MIDAS to construct a histogram on the PC. Figure 13.2 illustrates such a histogram for a single tile module, showing a calculated rate that matches the DESY rate of approximately 10 kHz. In the plot, the x-axis represents the direction along the test beampipe, while the y-axis indicates the rotation in ϕ . This straightforward method provides a realistic estimate of the rate observed by each channel in the detector, enabling the verification of prescaling methods and other properties of the system.

13.3 Time correlation studies

Similarly to the Cosmic Run analysis, a time-walk correction was performed for the pixel chips. However, in contrast to the Cosmic Run, the reference time for this correction was obtained from the tile detector. Figure 13.3 illustrates the correction process for the time-walk. Figure 13.3a shows the distribution before the correction, while Figure 13.3b shows the distribution after the correction.




Figure 13.3: Figure 13.3a shows ToT versus time difference of the pixel chips and the tile module. Figure 13.3b shows ToT versus time difference of the pixel chips and the tile module with time-walk corrected pixel hit-time.

Once again, it is important to note that the readout speed of the second timestamp was not individually optimised for each pixel chip during the test run. As a result, an imperfect sampling of the second timestamp occurred, which is particularly evident in the lower time difference part ($t_{\text{pixel}} - t_{\text{tile}} < -100$ and ToT less than 1000) of both graphs in Figure 13.3. In this region, the ToT calculation was not accurate, leading to a higher variance in the correction for these ToT values.

As described in Section 12.4, column addresses greater than 236 were excluded from the analysis, as the second timestamp for the MuPix10 chip does not work for these pixel cells. Furthermore, it is worth noting that a time-walk effect was observed in the tile detector as well. However, for the purposes of this study, its correction was not necessary since the overall timing resolution is primarily limited by the pixel timing resolution.

Figure 13.4 presents the time difference between the three half ladders (after time-walk correction) and the tile module. Unlike the analysis conducted during the Mu3e Cosmic Run (see Section 12.4),

no scintillator reference was used. To investigate the timing correlations between the two detectors, it was necessary to have at least one hit per detector during a pixel sorter frame of $16\ \mu\text{s}$. Note that during this testbeam the tile readout path did not have a hit sorter. Again, a Gaussian fit function was used to obtain the resolution of the time correlation.




Figure 13.4: The time correlation between three half ladders and one tile module.

The performed fit resulted in a total time resolution of approximately $79.5(28)$ ns. This outcome aligns with the timing resolution observed during the Cosmic Run, which was $98.5(27)$ ns. As before, the resolution is mainly influenced by the timing resolution of the three pixel ladders, which is around 100 ns due to the untuned chips and the low HV.

The successful reading of a full tile module and the observation of correlations between the pixel ladders demonstrate the progress made in the integration of the different detectors into the Mu3e DAQ system. These results pave the way for the initial commissioning and engineering tests scheduled for 2024, representing a significant milestone toward the search for $\mu^+ \rightarrow e^+e^-e^+$.

Part IV

Irradiation and μ SR Studies

In this part various applications using the MuPix chips beside the Mu3e experiment are discussed. In the first chapter, the radiation tolerance of 100 μm thin MuPix10 chips [134] was tested to verify potential use in the P2 experiment [136]. The second chapter will show the first simulations, detector designs, and testbeam results for using the MuPix11 chip for a new muon spin rotation (μ SR) detector. Accomplishing these tasks required the collaboration of multiple individuals. The author was actively involved in all the tests outlined in the following chapters. The author was involved in activities such as setting up the readout system, preparing the experimental setup, and analysing the collected data.

MuPix Irradiation Studies

In this chapter, a dive into an irradiation testbeam campaign using the MuPix10 chip is conducted. The primary objective was to assess the chip's performance under conditions that mimic the radiation exposure it would encounter in experiments like Mu3e and the planned P2 experiment. These experiments require the chip to withstand a maximum dose of approximately 30 MRad or 300 kGy [137].




Figure 14.1: Sketch of the setup used for the MAMI irradiation testbeam in September 2021.

Furthermore, High-Voltage Monolithic Active Pixel Sensor (HV-MAPS) radiation hardness of 100 Mrad and $5 \times 10^{15} \text{ n}_{\text{eq}}/\text{cm}^2$ had been demonstrated with very high doses of protons and neutrons [138]. Therefore, bulk radiation damage is not a big problem for the planned chip usage in the Mu3e and P2 experiments. However, in previous studies using instantaneous high electron rates, efficiency losses and increased noise level were observed using MuPix7 and MuPix8 chips [139]. These effects could be partially recovered with annealing and are most likely caused by some oxide charge-up. This test aimed to further investigate the effects of instantaneous high-electron-rate irradiation using MuPix10 the final design of the MuPix chip.

In the first section, the experimental setup used for the testbeam campaign is explained. It outlines the environment and conditions under which the MuPix10 chip was tested. The analysis techniques and methodologies applied to assess the chip's performance are subsequently detailed. This section explains how the data collected during the testbeam campaign was processed and evaluated. The final section of the chapter discusses the results obtained from the analysis.

14.1 Experimental setup

Similarly to the synchronisation testbeam described in Chapter 10, this testbeam was carried out at the Mainz Microtron (MAMI) accelerator facility in Mainz. These tests took place in September 2021, providing a controlled environment to evaluate the chip's performance under specific irradiation conditions. Figure 14.1 presents a schematic representation of the experimental setup employed during the testbeam. In this setup, a $100\ \mu\text{m}$ MuPix10 chip, acting as the device under test (DUT), was positioned within a telescope configuration consisting of three reference layers. The 855 MeV electron beam interacts with the first reference layer before reaching the DUT. Subsequently, it traverses two additional reference layers at the end of the setup.




Figure 14.2: Picture of the setup used for the MAMI irradiation testbeam in September 2021, taken from [140].

The experimental setup included a movable x/y table, as depicted in Picture 14.2, along with the DUT, mounted on a Physik Instrumente (PI) stage. The experiment went through different stages. First, the electron beam passed through the complete telescope configuration to evaluate the DUT's efficiency. Then, the x/y table was adjusted to remove it from the electron beam path. Only the DUT was reintroduced into the beam and exposed to irradiation. Finally, the telescope was re-positioned in its original configuration, allowing another measurement of the efficiency of the DUT. Notably, thanks to the MAMI accelerator's ability to produce a very small electron beam, this procedure could be repeated multiple times to irradiate different spots on the DUT with varying radiation doses.

Figure 14.3 displays the hitmaps for all four layers of the telescope. The diameter of the beam was adjusted to approximately 20 pixels, resulting in an effective size of $1600\ \mu\text{m}$. It is worth noting that the spots below and above the primary beamspot are attributed to cross-talk in the signal transmission,

a known issue in MuPix10 that has been successfully addressed in MuPix11, as detailed in [85]. The




Figure 14.3: Hitmaps of the four detection layers of the MuPix10 telescope.

electron rate can be adjusted by varying the voltage of the Wehnelt cup, which controls the beam current. However, it is important to note that the MAMI accelerator is typically calibrated for higher rates and currents up to $100 \mu\text{A}$, and is not as finely tuned for the lower rate regions used during the testbeam.

In Figure 14.4a, the rate, calculated using the MuPix telescope, was used to establish a relationship between the Wehnelt voltage and the rate. The data obtained from the pixel chips was fitted to the function $f(\text{Wehnelt}) = e^{a+b \cdot x}$, resulting in $a = 63.5(22) \text{ Hz}$ and $b = 4.4(2) \text{ Hz/V}$. For the pixel with the maximum rate, the fitted function yielded $a = 56.9(25) \text{ Hz}$ and $b = 4.3(2) \text{ Hz/V}$. This fitted function was then used to estimate the rate on the DUT during the irradiation tests.

To crosscheck the rate of the pixel chips, an ammeter was used, although with a limited measurement range in the picoampere region. Consequently, the ammeter measurement is expected to have a relatively high margin of error. However, during the irradiation tests, the ammeter measurement served as a safety check to avoid exposing the pixel chips to rate levels that could potentially damage them. To have a safe upper limit of operation, the maximum Wehnelt voltage used during the testbeam was -11 V , which was the lowest point the ammeter could measure.




Figure 14.4: Figure 14.4a shows the measurement of the rate on the chip / single pixel for a given Wehnelt voltage. Figure 14.4b shows the normalised rate map per pixel in the region of the beamspot.

Furthermore a rate map was measured within the beamspot region. Therefore, multiple runs were conducted at different rates, and a normalised rate map was generated (see Figure 14.4b). This map was later used in the analysis to estimate the rate over the whole beamspot on the DUT by multiplying it by the maximum rate of the centre pixel.

14.2 Analysis flow

For analysis and data collection, the MuPix telescope was employed [141]. The MuPix telescope is a specialised setup designed for testing MuPix chips in test beams. It consists of four layers of MuPix chips connected to an field programmable gate array (FPGA) (Stratix IV [142]) housed within the data acquisition (DAQ) system computer.

The coordinate system used within the MuPix telescope is illustrated in Figure 14.5. In this coordinate system, the beam direction aligns with the z -axis, while the x -axis runs along the columns. The global coordinate system, denoted (x, y, z) , is defined by layer 0.

To determine the efficiency of the DUT, a series of analysis steps are performed (see Figure 14.6). Initially, the software uses three reference planes to fit tracks. Tracking algorithms are adapted for MuPix telescope specifics, considering no magnetic field and minimal energy loss. Tracks are fitted based on a simple straight-line model without factoring in multiple Coulomb scattering:

$$f(z) = a + s \cdot z. \quad (14.1)$$

In this equation, z represents the z position in the global coordinate system, a is the initial (x, y) position in the global coordinate system at $z = 0$, and s is a two-dimensional slope. The typical uncertainties for x and y measurements are $\frac{\text{pixel size}}{\sqrt{12}}$, but this error would decrease for particles that trigger more than one pixel. The best-fitting track (with the lowest χ^2) is selected in case of multiple

hits in a layer. This tracking method offers a non-iterative approach for fast implementation, although it yields less precise reconstructed trajectories compared to models accounting for multiple scattering, such as the general broken line fit. However, in settings where the scattering is minimal, this method is sufficient.

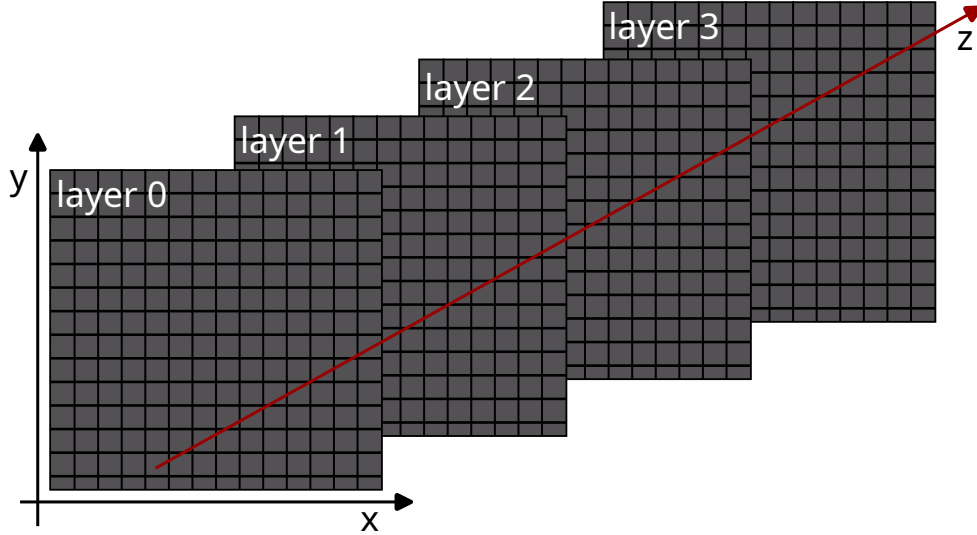


Figure 14.5: Sketch of the coordinate system for the MuPix telescope, the particle beam is shown in red, adopted from [143].

Subsequently, the fitted tracks are used to verify if the DUT also registered a hit simultaneously (see left of Figure 14.6). The hit efficiency ϵ is calculated by counting the total number of tracks N extrapolated to the DUT and those seen by the DUT k (see the middle figure of Figure 14.6), given by:

$$\epsilon = \frac{k}{N}. \quad (14.2)$$

The uncertainty in these measurements can be modeled using a Bernoulli trial with the following posterior distribution:

$$P(\epsilon|k, N) \propto \epsilon^k (1 - \epsilon)^{N-k}. \quad (14.3)$$

The prior in this case is set to 1. Note, that the uncertainty of the efficiency is asymmetric with an upper bound of 1 and a lower bound of 0. The uncertainty in the efficiency, denoted as δx , can be determined to achieve a confidence interval of 1σ :

$$\frac{\int_x^{x+\delta x} \epsilon^k (1 - \epsilon)^{N-k} d\epsilon}{\int_0^1 \epsilon^k (1 - \epsilon)^{N-k} d\epsilon} = \text{C.L.} \quad (14.4)$$

However, it's important to note that the calculated efficiency can be influenced by noise, particularly when the noise rate is high, leading to potentially higher efficiencies. This influence from noise can be quantified as:

$$\epsilon_m = \epsilon_k + (1 - \epsilon_k) \cdot \epsilon_n, \quad (14.5)$$

where $\epsilon_n = \frac{\pi \cdot r_{cut}^2}{A_p} \cdot n_p \cdot t_{window}$ represents the probability of matching a noise hit. In this equation, n_p is the average noise rate per pixel, A_p represents the area of a pixel, r_{cut} is the maximum matching radius and t_{window} is the time window used to search for matching hits [143]. To calculate the influence on the efficiency, the procedure is repeated both before and after irradiating the DUT. The difference between the two histograms is then used to illustrate the loss in efficiency caused by irradiation, as shown in the right of Figure 14.6.




Figure 14.6: Tracks and efficiency plots of the DUT. Left: Tracks position on the DUT. Middle: Efficiency map of the DUT. Right: Efficiency difference before and after the irradiation.

Furthermore, only the inner region of the beamspot is taken into account when calculating the efficiency difference, as indicated by the red circle in the middle figure of Figure 14.6. To enhance hit efficiency, the software also performs an alignment of the different layers. Dedicated runs are conducted to realign the system after each movement of the x/y table.

Threshold Scan

During the testbeam, a threshold scan was conducted to investigate the impact of changing the comparator threshold of the MuPix10 chip on signal efficiency. Figure B.7 in the appendix B.6, shows the threshold scan results obtained during the testbeam. The efficiency was calculated using the `TEfficiency` class from ROOT [123], with error calculations as explained previously.

It's worth noting that the threshold scan during this testbeam was performed with an applied high voltage (HV) of only -10 V for the DUT, as changes in efficiency are more visible at lower HV. However, all irradiation tests were performed with -100 V, which is the HV value where the chip works most efficiently [85]. This initial threshold scan should give an indication of the efficiency at different thresholds to obtain a low and high threshold for the later irradiation study.

Research Questions

The testbeam aimed to address three research questions:

- Does irradiating the chip without HV applied produce different results?
- How does the dose change the signal size?
- Are there any irradiation-related issues at higher doses that could affect efficiency?

To answer these questions, the chip was irradiated at different doses and its hit efficiency was tested at various thresholds. The chosen thresholds ranged from the most efficient at 45 mV to the least efficient at 248 mV. Additionally, during irradiation, HV was alternately switched on and off to assess its impact on chip performance after irradiation. The maximum dose exceeded 250 Gy for a single pixel.

14.3 Irradiation results

Figure 14.7 shows a previous measurement performed using MuPix8 [139]. This measurement demonstrates that the loss in efficiency increases linearly with the dose according to the function $f(x) = a \cdot x + b$, where $a = -0.008(1)$ p.p. and $b = 0.165(133)$ p.p..




Figure 14.7: Efficiency loss of an irradiated MuPix8 chip, adapted and taken from [139].

Figure 14.8 presents the results for the 45 mV threshold, while Figure B.8 of Appendix B.6 pertains to the 248 mV threshold. Both figures are divided into two regions: the upper part shows experiments conducted with HV turned on, while the lower part represents experiments with HV off. The x-axis denotes the column position of the DUT, and the y-axis represents the row position. Each circle in the hitmap represents the difference in efficiency, as shown in Figure 14.6. Additionally, the irradiation time and electron rate for the pixel with the highest rate are displayed above the circles.




Figure 14.8: Different irradiation spots for the 45 mV threshold having HV on and off.

To calculate the dose, the collision stopping power of electrons at 855 MeV in silicon ($3.343 \times 10^{-13} \text{ Jcm}^2/\text{g}$ [144]) is divided by the pixel area ($6.4 \times 10^{-5} \text{ cm}^2$). The total number of hits per pixel is determined using the beam rate, the total irradiation time, and the rate map shown in Figure 14.4b.

Given that the error for efficiency is asymmetric, calculating the error for the difference is not straightforward. Since the likelihood function of the difference is unknown, a parameterisation of the likelihood needs to be performed. However, there is no general solution to this problem; the error calculation utilises an empirical expression of the likelihood function [145]:

$$-\ln L(\alpha) = \frac{1}{2} \frac{\alpha^2}{\sigma_1 \cdot \sigma_2 + (\sigma_1 - \sigma_2) \cdot \alpha}. \quad (14.6)$$

Here, α represents the mean value of a parameter with asymmetric errors, and σ_1 and σ_2 denote the upper and lower errors, respectively. To illustrate, consider two values with asymmetric errors: $A = 5_{-2}^{+1}$ and $B = 3_{-3}^{+1}$. The likelihood function for the difference of A and B becomes:

$$-\ln L(A, B) = \frac{1}{2} \frac{(A-5)^2}{1 \cdot 2 + (1-2) \cdot (A-5)} - \frac{1}{2} \frac{(B-3)^2}{3 \cdot 1 + (3-1) \cdot (B-3)}. \quad (14.7)$$

Note, that the upper and lower errors for B need to be flipped caused by the subtraction. Now, introduce $C = A - B$:

$$-\ln L(C,A) = \frac{1}{2} \frac{(A-5)^2}{1 \cdot 2 + (1-2) \cdot (A-5)} - \frac{1}{2} \frac{(A-C-3)^2}{3 \cdot 1 + (3-1) \cdot (A-C-3)}. \quad (14.8)$$

The objective is to derive the likelihood function for C by minimising with respect to the nuisance parameter A :

$$-\ln L(C) = \min_A -\ln L(C,A). \quad (14.9)$$

The minimum of the solution is then used to extract the errors for C . For calculating the error during the analysis, the software as described in [146] was used. This approach allows to estimate the likelihood function and its associated errors for the difference between two values with asymmetric errors.




Figure 14.9: Efficiency difference vs. dose for different thresholds and HV on or off.

Figure 14.9 illustrates the efficiency difference as a percentage of efficiency before and after irradiation, plotted against dose. This analysis was carried out for two thresholds and with both HV on and off

during irradiation. The small deviations from zero in some of the graphs arise from the calculation method of the difference, as previously explained, combined with the fact that efficiency measurements with a value of 100% do not have an upper error bound. For runs with a low threshold of 45 mV and HV off, a tiny increase in efficiency, in the range of 0.002(1) p.p. per Gy, was observed after irradiation. However, for the run with a high threshold of 248 mV and HV off, a decrease in efficiency of $-0.017(5)$ p.p. per Gy was evident. In the runs where the irradiation was performed with HV on, both thresholds show a constant efficiency value. The reduction in efficiency can be attributed to a decrease in signal amplitude following irradiation.




Figure 14.10: Different noise sources observed after irradiation of around 80 Gy.

The slight increase in efficiency for the low threshold could be attributed to an increase in noise, as indicated in Figure 14.10. Two noise sources became apparent during the testbeam, both recorded after a dose of 80 Gy. It is worth mentioning that in the past no rate effect for the "bumps" was observed [147]. During the testbeam only some noise runs were taken. It cannot be excluded that the effects could also occur at a lower dose.

The first effect, termed "bumps", is depicted in Figure 14.10a (for rows smaller than 125)¹. These are randomly occurring large hit clusters on the pixel chip that are not correlated with the beamspot and persist for some time. The time dependence of this effect is shown in the hit-over-time plot for one noise run over 30 min shown in Appendix B.6, Figure B.9a. The source of this effect is suspected to be within the active pixel cell, as these clusters sometimes produce cross-talk. A possible explanation could involve current flow leading to the formation of these clusters, although this remains a hypothesis that requires further investigation. Another explanation may be some nuclear interaction in the sensor [147].

The second effect, illustrated in Figure 14.10b, is referred to as "after glow". This effect was also observed in tests with MuPix7 [139]. It describes the enhanced noise activity at the beam position for a certain duration. This effect exhibits a clear decay, as seen in Appendix B.6, Figure B.9b. The likely cause is the charge-up between the n-well substrate and the oxide as a result of ionisation in the oxide layer. Over time, these ions are neutralised by electrons from the surrounding material. For most runs, the half-life was approximately 28 sec (see Figure B.9b). It is worth noting that the

¹The hits higher than row 125 are caused by an "after glow".

half-life time changed with the temperature of the chip during older studies using MuPix7 [139]. This process is well-documented in the semiconductor industry, where silicon wafers undergo annealing to rectify atomic-level disorder resulting from production steps such as ion implantation. In this testbeam, the MuPix10 chip was not actively cooled, which could have accelerated the decay of the "after-glow" effect compared to the conditions in the final Mu3e experiment. Unfortunately, detailed Time-over-Threshold (ToT) studies of this effect could not be applied, since some of the high rate runs were conducted in columns higher than 236 which have no ToT information.

In summary, the testbeam results indicate that the MuPix10 chip can withstand high rates up to 200 Gy without encountering significant issues. With the HV turned off during irradiation, two distinct effects were observed. First, for the low threshold, a slight increase in noise attributed to the "after-glow" effect led to a seemingly higher efficiency after irradiation. Detecting this effect in the final experiment might be possible through the implementation of various cuts on the ToT, timestamp, or position. However, due to the lack of statistical data, such detailed studies were not feasible in the recorded testbeam data. Second, a decrease in efficiency was observed for rates up to 250 Gy with a higher threshold, due to a reduction in signal amplitude. Furthermore, "bumps" were observed during noise runs, likely caused by nuclear interactions in the sensor [147], requiring further investigation.

15

Advanced muon spin spectroscopy using MuPix chips

Muon-spin spectroscopy at continuous sources has been constrained to a stopped muon rate of 40 kHz for several decades. The primary impediment arises from the necessity of having only one muon present in the sample during the 10 μ s data gate window. To overcome this limitation and enable μ SR measurements on submillimeter-sized samples, Si-Pixel-based spectrometers with vertex reconstruction schemes for incoming muons and emitted positrons are under development. This advance is particularly crucial with the planned High-Intensity Muon Beams (HIMB) project [148], with the objective of taking advantage of the increase in muon rate for materials science studies at Paul Scherrer Institute (PSI).

In this chapter, simulations and preliminary testbeam results are presented for the first prototype μ SR spectrometer that employs MuPix11 chips. The chapter is based on an approved grant proposal from the Swiss National Science Foundation [149] and a testbeam campaign scheduled at PSI in September 2023. The author contributed to the testbeam proposal, conducted detector simulations, adapted the Mu3e DAQ system for reading out the μ SR spectrometer, and participated in the testbeam.

As the final testbeam results were not available at the time of writing this thesis, only preliminary work is presented. However, there are plans to publish the results obtained during the testbeam.

15.1 Background and introduction

The ability to perform μ SR measurements on submillimeter samples is a long-held aspiration of material scientists, particularly in the realm of novel quantum materials that are challenging to produce in large quantities. Moreover, data collection rates at continuous muon sources, which provide high-resolution measurements in time, have remained static at a stopped muon rate of 40 kHz since the technique was invented. The MuPix11 chip [85], has demonstrated significant progress in silicon pixel detector chips used in particle physics experiments. This advancement has resulted in a spatial resolution of approximately 23 μ m for a single hit and a time resolution of less than 15 ns. Using this technology has immense potential to improve the μ SR technique in various ways:

- Precisely characterise the magnetic and electronic properties of materials at submillimeter lateral resolutions.
- Simultaneously measure multiple samples.
- Conduct pump-probe and transient measurements, capitalising on the high data-collection rate.

- Acquire μ SR signals with a reduced uncorrelated background, allowing longer time window measurements beyond the current 10 μ s limit without compromising time resolution or data rate.

To achieve these objectives, thin MuPix11 chips (approximately $0.5 \times 10^{-3} X_0$) are used to construct a μ SR spectrometer capable of continuous operation with a muon beam at PSI.

15.1.1 Muon beam at PSI

A polarised-muon beam is generated by collecting muons produced through the two-body decay of positive pions: $\pi^+ \rightarrow \mu^+ + \nu_\mu$. These pions are created in production targets located inside the beam line that enters the high-energy proton accelerator, which operates at 590 MeV. Since pions have zero spin and exclusively produce left-handed neutrinos (ν_μ) during their decays, the resulting muons possess spins that are antiparallel to their momentum in the rest frame of the pion.

At the Swiss Muon Source (S μ S) located at the PSI, various muon beams are available with different energy ranges. Most muon beams at PSI are produced from muons generated near the surface of the production target, known as the "surface" or "Arizona" beam [150]. The muons in this beam are 100% polarised and ideally monochromatic, exhibiting very low momenta at approximately 29.8 MeV/c, corresponding to a kinetic energy of 4.1 MeV.

15.1.2 Introduction to the μ SR technique

This section provides a brief introduction to the μ SR technique, which closely follows the introduction in [151]. For more in-depth information, the readers are encouraged to consult comprehensive textbooks and review articles [152, 153, 154, 155, 156].

In μ SR, polarised positive muons (μ^+) implanted in a sample are subjected to magnetic interactions, resulting in a time-dependent polarisation denoted $\mathbf{P}_\mu(t)$. The muons have a finite lifetime, described by the exponential decay function $e^{-(t/\tau_\mu)}$, with $\tau_\mu \approx 2.197 \times 10^{-6}$ s representing the muon lifetime.

The μ SR technique relies on the weak decay of muons ($\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu$), which produces an asymmetric distribution of emitted positrons relative to the muon spin direction at the time of decay. By measuring the spatial distribution of the positrons as a function of the muon decay time using detectors placed around the sample, it is possible to deduce the time evolution of muon polarisation $\mathbf{P}_\mu(t)$. For a positron detector orientated in the direction \hat{n} with respect to initial polarisation of muons $\mathbf{P}_\mu(0)$, the time histogram of the time intervals collected between muon implantation and positron detection is described by the equation:

$$N(t) = N_0 \cdot \exp\left(-\frac{t}{\tau_\mu}\right) \left[1 + A_{sym} \cdot \mathbf{P}_\mu(t) \cdot \hat{n}\right] + N_{bg}. \quad (15.1)$$

Here, N_{bg} represents a time-independent background, N_0 is the number of counts at $t = 0$, and the exponential term accounts for the muon decay. The detector's asymmetry is given by $A(t) = A_{sym} \cdot \mathbf{P}_\mu(t) \cdot \hat{n}$, with $\mathbf{P}_\mu(0) = |\mathbf{P}_\mu(0)|$, corresponding to the beam polarisation, typically on the order of ~ 1 . The parameter a depends on the weak decay mechanism's intrinsic asymmetry, the detector's solid angle, efficiency, and positron absorption and scattering in materials along their path.

The muon polarisation function $P_{\mu}(t)$ describes the time dependence and contains all pertinent information on the magnetic interactions of muons within the sample. In μ SR measurements, the primary result reported is usually the μ SR signal $A(t)$, obtained by fitting individual histograms using Equation 15.1 or measured directly from two positron counters placed on opposite sides of the sample.

15.2 Conceptual design

To build an effective μ SR spectrometer, it is crucial to achieve precise tracking of incoming muons and emitted positrons with exceptional timing and spatial precision. Figure 15.1 presents the conceptual design of this detector, which employs two layers of Si-Pixel detectors surrounding the test sample.




Figure 15.1: Sketch of the Si-Pixel based μ SR spectrometer.

First an incoming muon (μ^+) enters the spectrometer from the left, registering hits in both detector layers. Subsequently, the muon is stopped in the sample (S) and decays while in a stationary state. The emitted positron (e^+) exits the sample, leaving hits in both detector layers.

This design provides complete coverage in all directions. The spatial distribution of the decay positrons is governed by the properties of the sample under examination and any externally applied magnetic fields. All positrons that follow a trajectory that passes through both detector layers will be accurately recorded.

The efficiency of the detector is primarily based on the hit efficiency of the pixel detectors used. Continual developments for the MuPix chip [134] have shown high efficiency exceeding 99% and time resolution of less than 20 ns. The primary source of inaccuracy in the determination of the

track is caused by multiple Coulomb scatterings, such as the incoming surface muons at PSI possess a momentum of 28 MeV/c. To mitigate this effect, the proposed MuPix chips are ultra-thin, with a thickness of 50 μm , resulting in a minimal material budget of less than $1 \times 10^{-3} X_0$ per layer.

15.2.1 Current status

Previously conducted beam tests have demonstrated the feasibility of operating multiple MuPix chips within a prototype of the Mu3e vertex detector. Furthermore, these tests successfully reconstructed the tracks of decay positrons using a μSR sample, as documented in references [132, 131, 90].

For a complete understanding of these test runs, additional details can be found in Chapters 11 and 12 of this thesis. In particular, Figure 11.9b in Chapter 11 illustrates the spatial correlations between two pixel layers of the vertex detector prototype using an μSR insert. This insert comprises two layers separated by a distance of 2.5 cm. The first layer houses a scintillator that serves as an entrance counter, while the second layer accommodates the actual sample under investigation, a silver disc positioned between two permanent magnets. The spatial correlation plot features two lines that correspond to the individual layers of the insert.

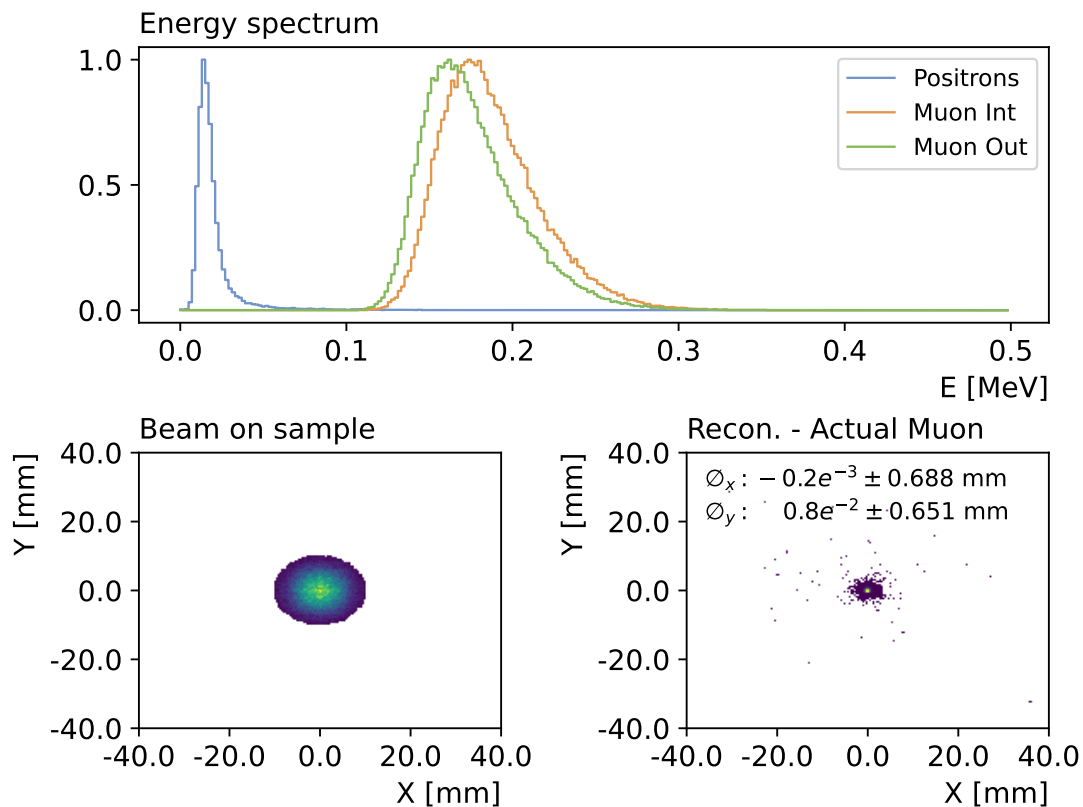


Figure 15.2: Top: Deposited energy in the Si chip for positrons (blue), muons in the inner Quad-Modules (orange) and muons in the outer Quad-Modules (green). Left: Muon beam on the sample. Right: Difference between the reconstructed and actual Muon position from the Monte-Carlo simulation.

15.3 Monte-Carlo simulation

The primary source of uncertainty in the vertex reconstruction of an incoming muon trajectory, determined from the outer and inner hits, is attributed to multiple Coulomb scattering occurring within the inner Si chip. To provide an initial estimate of the uncertainty in the vertex reconstruction, MusrSim simulations [157, 158], which are based on GEANT4 [97], are employed to model the scattering behaviour of muons traversing $50\ \mu\text{m}$ silicon. In Figure 15.2, different results of simulating the setup using a distance of $r = 10\ \text{mm}$ between the sample and inner MuPix chips are shown.

The cross section of the muon beam that impinges on the sample is presented in the lower left part of Figure 15.2. In the lower right part of Figure 15.2 the difference between the reconstructed and actual muons at the target is shown. The standard deviation of this difference is used as a measure of the uncertainty in determining the muon's lateral position on the sample via vertex reconstruction. In particular, the uncertainty in the proposed configuration is less than $0.7\ \text{mm}$ ($0.688\ \text{mm}$ in x and $0.651\ \text{mm}$ in y), which falls below the target resolution of approximately $\sim 1\ \text{mm}$. It is important to note that further enhancement of the resolution and reduction of uncertainty can be achieved by combining the extrapolated tracks of the muon and the coincident positron to obtain a match. The key to efficient track extrapolation for both muons and positrons lies in the ability to differentiate between these two particle species.

One effective approach to achieve this differentiation is to tune the ToT measurement implemented on the MuPix chip based on the particle momenta present in the μSR setup. The top of Figure 15.2 illustrates the simulated deposited energy spectra for positrons (blue), muons in the inner Quad-Modules (orange) and muons in the outer Quad-Modules (green), providing a basis for such differentiation.




Figure 15.3: Figure 15.3a shows one Quad-Module with four MuPix11 chips. Figure 15.3b shows the prototype detector with four layers of Quad-Modules.

15.4 Prototype and first beam tests

Figure 15.3a shows a Quad-Module, which is made up of four MuPix11 chips mounted on a Kapton foil. The Quad-Module was specifically developed in Heidelberg [159]; the author provided feedback during the development phase for a smooth integration into the DAQ system. The MuPix11 chips are wire bonded onto a printed circuit board (PCB), with more comprehensive information available in [159].

In Figure 15.3b, you can observe the prototype detector inspired by the design of the MuPix telescope, as described in [141]. This design incorporates four layers of Quad-Modules, each holding four MuPix11 chips. The primary focus of this design is to optimise submillimeter lateral resolution to determine the positions of incoming muons in the sample.




Figure 15.4: First results taken during the testbeam in September 2023. The hitmaps of the four Quad-Modules are shown on the left part of the plot, while the right part shows normalised ToT distributions for the 15 working chips.

The estimation of the stopping position of the muon in the sample involves reconstructing the muon trajectory from pixel hits in two detector layers. Timing information from the pixel detectors is used to measure the arrival time of the muon. Similarly, the outgoing positron's trajectory is detected and time-tagged, providing data on the observed lifetime of the muon and the direction of the decay positron. This setup allows one to measure the time evolution of the asymmetry and muon-spin polarisation for muons in a specific region within the sample.

Furthermore, the setup includes two scintillators to verify the time measurements of the pixel detectors for the incoming and outgoing times of the muons and positrons, respectively. In September 2023, an initial test was conducted at the π E3 beam line at PSI, where fully polarised muons were directed at various targets to assess the detector's performance.

Figure 15.4 presents the initial results of the performance test beam. The left side of the figure displays different hitmaps from the four Quad-Modules used in the testbeam. During the test run presented, the target consisted of a silver dot placed between two permanent magnets. Silver was used to stop the incoming muons at a specific location, while permanent magnets are necessary for the spin rotation of the muons.

The hitmaps in the first two layers clearly show a beam spot, but this spot decreases in the last two layers. In layers 0 and 2, some noise was observed at the edges of the chips, leading to masking of these pixels during data collection. Furthermore, a chip on layer 2 displayed noisy columns, which were also excluded during data collection. A layer 3 chip was non-functional and remained turned off.

The right side of the figure depicts the normalised ToT distributions for all chips. While individual ToT measurements were not made during the test beam, two distinct peaks can be observed in the distributions. The peak with higher ToT values, exceeding 1000, is probably associated with muons, while the peak with lower ToT values corresponds to decay positrons. However, further tests are necessary to obtain more unified ToT spectra for all chips.

The Corryvreckan software framework [160, 161] was used to fit the tracks, as described in Section 14.2. These tracks were separated into tracks from layers 0 and 1 (upstream (US) tracks) and layers 2 and 3 (downstream (DS) tracks). To distinguish tracks from muons and tracks from positrons, various time and spatial cuts were applied. In the first step, tracks that are located in the region of the muon beam (see hitmap of layer 0) and at the same time can be extrapolated to the z position of the target are tagged as muons. The corresponding positron track is defined as the next track in time. If more than one track was recorded in a time window of $13 \mu\text{s}$ after the muon track, the whole event is discarded. Furthermore, the distance between the incoming muon and outgoing positron tracks should be less than 1 mm on the sample. When both a muon track and a positron track met these requirements, the time difference between them was recorded.

Figure 15.5 presents the difference of the US tracks for a test with the silver dot sample using the Quad-Module setup, and also includes a reference measurement conducted with the General Purpose Surface-Muon Instrument (GPS) [151]. Data fitting was carried out using Musrfit [162] for the GPS data. For the Quad-Module data, Equation 15.1 was adapted by replacing $A(t)$ with a damped cosine function resulting in:

$$N(t) = N_0 \cdot \exp\left(-\frac{t}{\tau_\mu}\right) \cdot (1 + A_{sym} \cdot e^{-t \cdot lam} \cdot \cos(2 \cdot \pi \cdot f \cdot t + \varphi)) + N_{bkg} \quad (15.2)$$

Here, f is the Larmor frequency, lam is the width of the field distribution, and φ the phase. For both measurements, the Larmor frequency (Quad-Module: 0.868(57) MHz, GPS: 0.857(2) MHz) and the width of the distribution (Quad-Module: 0.180(55) $1/\mu\text{s}$, GPS: 0.176(10) $1/\mu\text{s}$) were found to be nearly identical within the margin of error. The uncorrelated background in the spectrum is nearly zero, which is an improvement over GPS, and the muon lifetime is fixed at 2.2 μs , convergent to this value during the fit. The difference in phase can be explained with by the different orientation of the detector relative to the polarisation of the muon beam.




Figure 15.5: Comparison of the muon oscillation of the silver sample using the Quad-Modules and the GPS apparatus [151].

These are promising preliminary results that indicate that the proposed technique can be used for advanced muon spin spectroscopy. However, most of the testbeam data had not been analysed at the time of writing this thesis. This section is intended to provide an initial glimpse, and a paper with a more detailed analysis of the other samples used during the testbeam is currently being prepared. In addition, a follow-up testbeam for 2024 is planned to integrate the system into a cryostat to operate at cold temperatures.

Part V
Conclusion and Outlook

Conclusion and outlook

The pursuit of physics beyond the Standard Model of particle physics (SM) is a central theme in particle physics, driving a multitude of high-intensity experiments. Among these, the Mu3e experiment stands out with its aim of detecting the charged lepton-flavour violation (CLFV) decay $\mu^+ \rightarrow e^+ e^- e^+$. This decay, forbidden within the confines of the SM and suppressed even in extensions that incorporate neutrino oscillations, has an expected branching fraction (BF) of less than 10^{-54} . The potential observation of $\mu^+ \rightarrow e^+ e^- e^+$ would be a clear sign of new physics, while its absence would impose stringent constraints on models that extend beyond the SM.

In its Phase I, Mu3e aims to improve sensitivity to the branching ratio of $\mu^+ \rightarrow e^+ e^- e^+$ to 2×10^{-12} , surpassing the current limit set by SINDRUM in 1988 by nearly three orders of magnitude. Achieving this requires a detector that is not only fast and granular to handle 10^8 muon decays per second but also maintains an extremely low material budget - around 0.1 % of a radiation length per detection layer. The complexities introduced by multiple Coulomb scattering in the low-energy decay products of muons at rest are addressed using four pixel layers equipped with High-Voltage Monolithic Active Pixel Sensor (HV-MAPS) as an ultra-thin tracking detector. These layers are made of MuPix sensors, each with an active area of $2 \times 2 \text{ cm}^2$, specifically developed for Mu3e.

Furthermore, precision timing measurements are facilitated by a scintillating fibre detector and a scintillating tile detector, both read out via the Muon Timing Resolver including Gigabit-link (MuTRiG) application specific integrated circuit (ASIC), sharing similar link speeds of 1.25 Gbits/s with the MuPix chip. To manage the total data rate of 100 Gbit/s, a field programmable gate array (FPGA) based data acquisition (DAQ) system is employed, capable of processing, time-aligning and selecting data in real time, a necessity given the spatial randomness of decay electrons from muons at rest.

This thesis has made significant contributions to the Mu3e experiment, particularly in integrating subdetectors into the DAQ system and optimising the data flow within the FPGA based system. It details the development and testing of the data path for the fibre detector on the Front-end board (FEB), the time-alignment of different Front-end boards (FEBs) within the FPGA system, and elucidates the data journey through the Switching Board (SWB) to the farm FPGA. Furthermore, the thesis introduces an online data quality analysis framework, which has been tested in various testbeam campaigns, ensuring reliable data integrity during operation.

A critical challenge addressed was the synchronisation of the scintillating and pixel detectors. The MuPix operates at 125 MHz, while MuTRiG functions at a 625 MHz clock, counting to $2^{15} - 1$. Aligning the timestamps of these application specific integrated circuits (ASICs) to an 8 ns depths ensures synchronisation with the global 125 MHz clock.

Moreover, the work analysed the irradiation testbeams of the MuPix10 chip, which are crucial in assessing the long-term durability and reliability of the detector under high radiation conditions. These testbeams simulate the intense radiation environment of the Mu3e experiment, ensuring that all components maintain their performance over extended periods. The testbeam results indicate that the MuPix10 chip can withstand high rates up to 200 Gy without encountering significant problems.

Based on the work presented in this thesis, the final commissioning of the Mu3e detectors and the readout system can start in 2024. More work needs to be done to fully integrate the graphics processing unit (GPU) selection, while the developed algorithms need to be tuned to operate at full performance. Optimising all of these parameters is achievable in 2024, leading to a first physics run in 2025.

In addition to the work for the Mu3e experiment, the thesis investigated the possibility of using the MuPix chip to build an advanced muon spin rotation (μ SR) spectrometer. The first testbeams conducted showed that μ SR measurements are possible using pixel detectors, opening new possibilities for material scientists. A follow-up testbeam for 2024 is planned to further tune the setup and integrate the system into a cryostat in order to operate at cold temperatures together with cooled samples.

Part VI

Appendices

A

Acronyms

ALICE	A Large Ion Collider Experiment	105
APS	Active Pixel Sensors	16
ASIC	application specific integrated circuit	145
ATX	Advanced Transmit PLL	27
Avalon-MM	Avalon Memory Mapped Interface Main	29
BAR	Base Address Register	35
BE	bit error	32
BE	Byte Enable	34
BERT	Bit Error Rate Test	70
BF	branching fraction	145
CAD	computer-aided design	41
CDR	Clock and Data Recovery	25
CGB	clock generation block	27
CKM	Cabibbo-Kobayashi-Maskawa	5
CL	confidence level	9
CLFV	charged lepton-flavour violation	145
CMBL	Compact Muon Beam Line	41
CMOS	complementary metal-oxide-semiconductor	16
CMS	Compact Muon Solenoid Experiment	62
CPU	Central Processing Unit	113
CRC	cyclic redundancy check	33
d	down quark	3
DAB	Detector Adaptor Board	71
DAC	digital-to-analog converter	47
DAQ	data acquisition	145
DC	direct current	64
DDR SDRAM	Double Data Rate Synchronous Dynamic Random-Access Memory	58
DESY	Deutsches Elektronen-Synchrotron	91
DMA	Direct Memory Access	92
DS	downstream	141
DUT	device under test	158

Acronyms

DW	double word	33
e	electron	5
EOP	end of package	66
ESR	Enhanced Specular Reflector	53
FE	front-end	78
FEB	Front-end board	145
FIFO	first in first out memory	67
FMC	FPGA Mezzanine Card	59
Fmt	Format	33
FPGA	field programmable gate array	145
fPLL	fractional PLL	27
g	gluon	
γ	photon	4
GPS	General Purpose Surface-Muon Instrument	141
GPU	graphics processing unit	146
HDD	hard disk drive	34
HDI	High-Density Interconnect	60
HIMB	High-Intensity Muon Beams	135
HIPA	High Intensity Proton Accelerator	40
HV	high voltage	128
HV-MAPS	High-Voltage Monolithic Active Pixel Sensor	145
I/O	input/output	20
IP	Hard Intellectual Property Core	20
JTAG	Joint Test Action Group	64
KamLand	Kamioka Liquid Scintillator Antineutrino Detector	37
LDO	Low Drop Out	45
LE	Logic Element	163
LED	light-emitting diode	31
LF	loop filter	25
LFV	lepton-flavour violation	1
LHC	Large Hadron Collider	62
LHCb	Large Hadron Collider beauty	105
LSB	last significant bit	82
LUT	Look-Up Table	19
LVDS	low-voltage differential signaling	92
MAC	Media Access Controller	33
MAMI	Mainz Microtron	156
MAPS	Monolithic Active Pixel Sensors	16
MEG	Mu to E Gamma	62
MIDAS	Maximum Integrated Data Acquisition System	91

MMIO	Memory-mapped input/output	35
μ	muon	3
μ SR	muon spin rotation	146
MuTRiG	Muon Timing Resolver including Gigabit-link	145
NIM	Nuclear Instrumentation Module	113
NRZ	non-return-to-zero	25
ν_e	electron neutrino	5
ν_μ	muon neutrino	5
ν_τ	tau neutrino	5
$\bar{\nu}_\mu$	anti muon neutrino	7
PC	personal computer	92
PCB	printed circuit board	140
PCI	Peripheral Component Interconnect	32
PCI-SIG	PCIe Special Interest Group	32
PCIe	Peripheral Component Interconnect Express	73
PCML	pseudo current mode logic	25
PCS	Physical Coding Sublayer	27
PFD	phase frequency detector	25
PHY	physical layer	24
PI	Physik Instrumente	124
PLL	phase locked loop	25
PMA	Physical Medium Attachment	27
PMA	Physical Medium Attachment	27
PMNS	Pontecorvo–Maki–Nakagawa–Sakata	6
PMOS	p-channel metal-oxide-semiconductor	46
PMT	photomultiplier tube	11
ppm	part-per-million	31
PSI	Paul Scherrer Institute	135
QFT	quantum field theory	3
QSFP	quad small-form-factor pluggable ports	79
R&D	research and development	91
RAM	Random-Access Memory	92
RD	running disparity	26
Receiving board	PC interface board	92
RMS	root mean square	15
RPC	Remote Procedure Call	78
SERDES	serializer or deserializer	28
SFP+	Enhanced Small Form-factor Pluggable	73
SiPM	silicon photomultiplier	59
SK	Super-Kamioka Neutrino Detection Experiment	37

Acronyms

SM	Standard Model of particle physics	145
SMB	SciFi module board	60
S μ S	Swiss Muon Source	136
SNO	Sudbury Neutrino Observatory	37
SOP	start of package	66
SPI	Serial Peripheral Interface	49
SRAM	Static Random Access Memory	20
SSW	Service Support Wheel	60
SUB	sub-header	66
SUSY	supersymmetry	8
SWB	Switching Board	145
t	top quark	5
τ	tau	5
TC	Traffic Class	
TDC	time-to-digital converter	49
TLP	Transaction Layer Packet	33
TMB	Tile module board	78
ToA	time-of-arrival	47
ToT	Time-over-Threshold	133
TX	Transceiver	27
u	up quark	3
US	upstream	141
VCO	voltage controlled oscillator	25
VEV	vacuum expectation value	3
VHDL	Very High Speed Integrated Circuit Hardware Description Language	19

B

Additional Material

B.1 The Mu3e Experiment




Figure B.1: Rate dependence at Mu3e centre versus spectrometer field, taken from [82].




Figure B.2: Sketch of the functionality of the MuTRiG chip, taken from [51].

B.2 Switching Board




Figure B.3: State machine of the hit-time-alignment firmware.

B.3 First Detector Integration

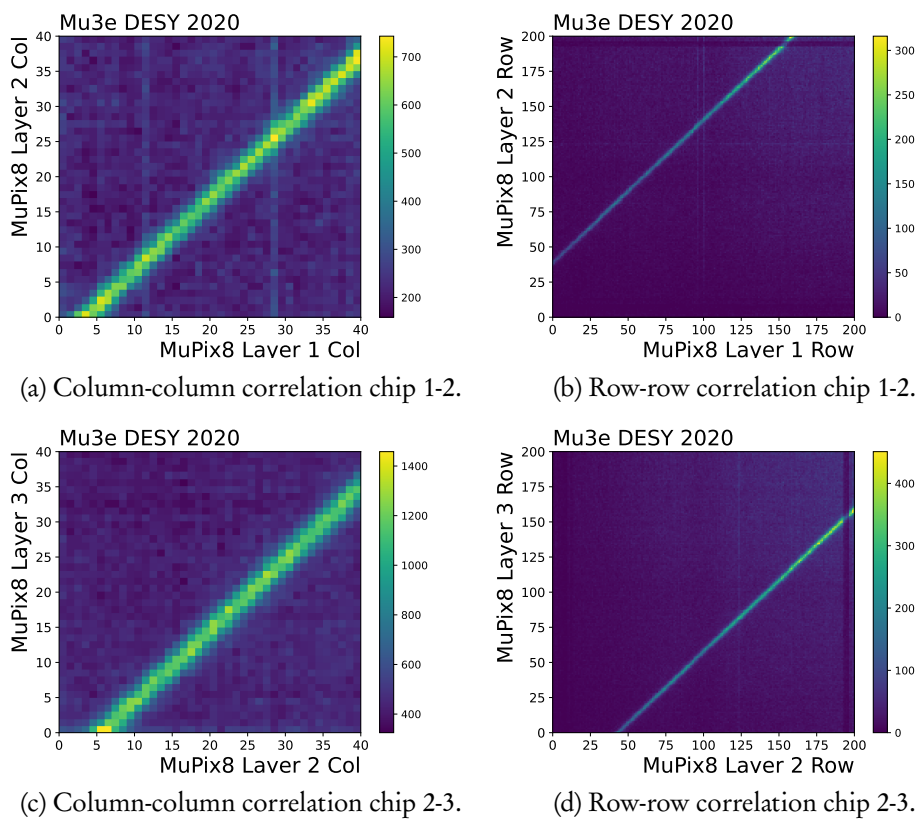


Figure B.4: Space correlations of layer 1, layer 2 and layer 3 of the MuPix8 chips.

B.4 Front-end Board Synchronisation

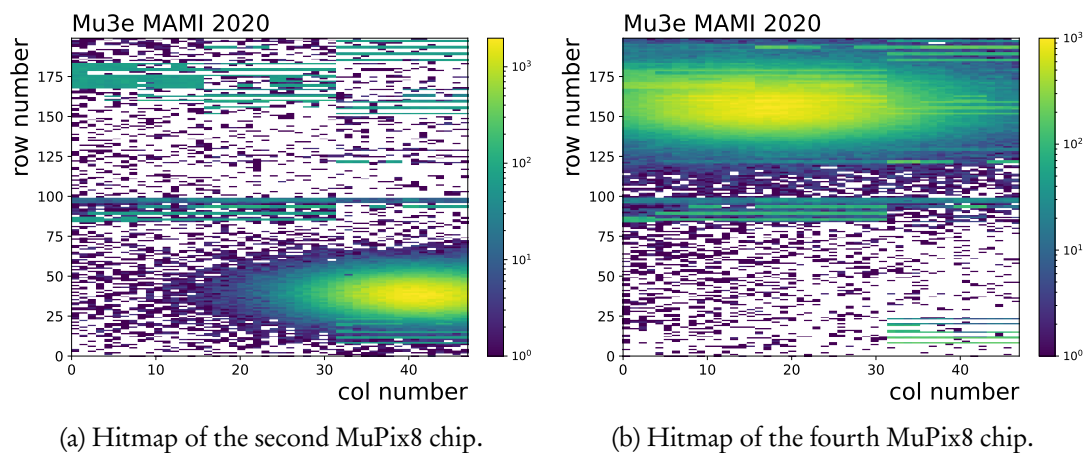


Figure B.5: Hitmaps of the second MuPix8 chips of the two FEBs in the Mainz Microtron (MAMI) synchronisation setup.

B.5 Mu3e Integration Run 2021




Figure B.6: Row to row correlations of the two layers of the vertex detector prototype using the Mu3e Target.

B.6 MuPix Irradiation Studies




Figure B.7: Threshold scan of the 100 μm device under test (DUT) at -10 V . With the low and high thresholds marked with red lines.

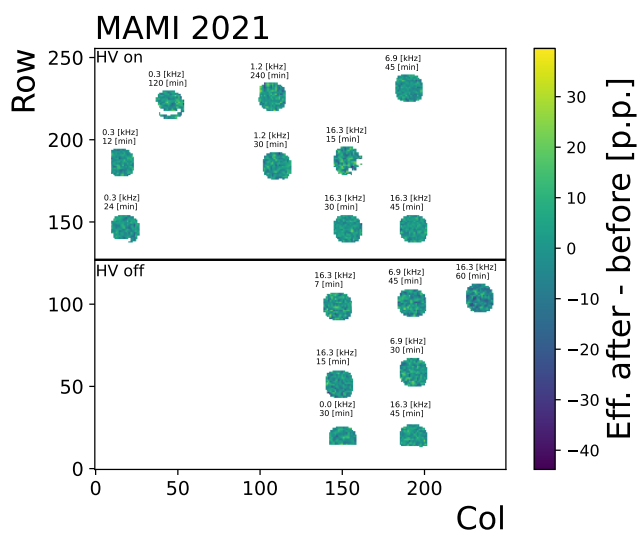


Figure B.8: Different irradiation spots for the 248 mV threshold having HV on and off.




Figure B.9: Decays from different noise sources observed after irradiation of around 80 Gy.




Figure B.10: No noise for a run with a dose less than 80 Gy and another "after glow" of a run with a dose higher than 80 Gy.



Publications

Parts of the ideas and work discussed in this thesis have been previously published in the following journal articles and conference proceedings.

Data Flow in the Mu3e DAQ

M. Köppel
IEEE Transactions on Nuclear Science

Mu3e Integration Run 2021

M. Köppel
The 22nd International Workshop on Neutrinos from Accelerators

The Mu3e Data Acquisition

H. Augustin et al.
IEEE Transactions on Nuclear Science

Technical design of the phase I Mu3e experiment

K. Arndt et al.
Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment

Performance of the large scale HV-CMOS pixel sensor MuPix8

H. Augustin et al.
Journal of Instrumentation

Data Flow in the Mu3e Filter Farm

M. Köppel
Master Thesis

Furthermore, some articles are still unpublished but were part of the thesis.

Advanced Muon Spin Spectroscopy using MuPix chips

to be published

Multiple Scattering in thin materials

to be published

The following publications are related to research projects in addition to the PhD study.

Google Topics as a way out of the cookie dilemma?

J.-P. Muttach, M. Köppel and G. Hornung
Computer & Recht 2023

The Case for Correctability in Fair Machine Learning

M. Cerrato, A. V. Coronel and Marius Köppel
European Workshop on Algorithmic Fairness 2023

Can machine learning solve the challenge of adaptive learning and the individualization of learning paths? A field experiment in an online learning platform

T. Klausmann et al.

AAAI 2023 Artificial Intelligence for Education

Invariant Representations with Stochastically Quantized Neural Networks

M. Cerrato et al.

AAAI Conference on Artificial Intelligence 2023

Learning to Rank Higgs Boson Candidates

M. Köppel et al.

Nature Scientific Reports

Ranking Creative Language Characteristics in Small Data Scenarios

J. Siekiera et al.

Proceedings of 13th International Conference on Computational Creativity

Fair Interpretable Representation Learning with Correction Vectors

M. Cerrato et al.

arXiv:2202.03078

Fair Interpretable Learning via Correction Vectors

M. Cerrato et al.

ICLR-21 Workshop on Responsible AI

Fair Group-Shared Representations with Normalizing Flows

M. Cerrato et al.

ICLR-21 Workshop on Responsible AI

Fair pairwise learning to rank

M. Cerrato et al.

2020 IEEE 7th International Conference on Data Science and Advanced Analytics

Pairwise Learning to Rank by Neural Networks Revisited: Reconstruction, Theoretical Analysis and Practical Performance

M. Köppel et al.

Joint European Conference on Machine Learning and Knowledge Discovery in Databases

List of Figures

1.1	Overview of the SM.	4
1.2	Muon decays in the SM.	7
1.3	μ decay via SUSY particles.	8
1.4	Production channels of $\mu^+ \rightarrow e^+e^-e^+$ in the Scotogenic Model.	8
1.5	Sketch of the SINDRUM experiment.	10
1.6	Sketch of a simulated event in the MEG experiment.	11
2.1	Particle travelling through a silicon detector.	13
2.2	Sketch of a particle travelling through a tracking detector.	14
2.3	Particle passing through matter.	15
2.4	Tracking in the scattering dominated regime.	16
2.5	Sketch of a HV-MAPS sensor.	17
2.6	Sketch of light travelling inside a fibre scintillator.	17
3.1	Structure of an FPGA.	19
3.2	Illustration of a simple LE.	20
3.3	Sketch of connection and routing of a Logic Element (LE).	21
3.4	Infinite Cascade model.	24
3.5	Illustration of the NRZ code.	25
3.6	Illustration of the principle of a PLL.	26
3.7	Overview of the TX bank architecture.	28
3.8	Overview of a TX channel.	29
3.9	Building blocks of an Intel Arria 10 TX.	30
3.10	TX channel datapath and clocking.	31
4.1	History of searches for CLFV.	37
4.2	Overview of the signal process and the main background processes.	38
4.3	Sketch of the Mu3e detector concept.	39
4.4	Multiple scattering in a magnetic field.	40
4.5	CAD model of the entire π E5 channel and CMBL.	41
4.6	Measured beam profile at the collimator.	42
4.7	Hollow double-cone muon stopping target made of aluminised Mylar foil.	43
4.8	Picture of the delivery of the Mu3e magnet.	44
4.9	Layouts and size comparison of selected MuPix prototypes.	45
4.10	MuPix10 block diagram.	46
4.11	Analogue electronics within a pixel cell of the MuPix chips.	47
4.12	Illustration of the 2-comparator threshold method.	48
4.13	Illustration of the MuTRiG channel's components and signal flow.	49

4.14	Working principle of the TDC.	49
4.15	Schematic of the MuTRiG TDC.	50
4.16	Figure of a full size SciFi ribbon prototype.	51
4.17	CAD rendering of the SciFi detector.	51
4.18	Position of the tile detector inside the Mu3e experiment.	52
4.19	Tile scintillator geometry: (left) edge tile, (right) central tile.	52
4.20	Picture of individual tiles with ESR foils.	53
5.1	Sketch of the final Mu3e DAQ system.	58
5.2	Sketch and picture of the clock and reset distribution system.	59
5.3	Active detector elements and the readout cabling.	60
6.1	Picture of the FEB.	63
6.2	Sketch of the hit sorter.	64
6.3	Sketch of the FEB datapath of the SciFi detector.	71
6.4	Waveform of the lapse correction with a two bit width to visualise the principle.	72
7.1	Picture of the PCIe40 Board.	73
7.2	Sketch of the SWB datapath of the central pixel detector.	74
7.3	Example of a four-to-one time-alignment firmware on the SWB.	76
7.4	Design blocks for the slow control communication from the SWB to the FEB.	77
7.5	Design blocks for writing a slow control command from MIDAS to the FEB.	78
8.1	Picture of the Receiving board.	79
8.2	Sketch of the Receiving board datapath.	80
8.3	Sketch of the transformed hits located in the PC RAM.	81
8.4	State machine for the MIDAS event builder firmware on the Receiving board.	83
8.5	State machine for the DMA readout firmware on the Receiving board.	84
8.6	Sketch of the Mu3e Online Analyzer framework.	85
8.7	Flow diagram of the online reconstruction software.	86
9.1	Sketch of the setup used for the DESY integration testbeam.	92
9.2	Hitmaps of the all four MuPix8 chips used in the setup.	93
9.3	Space correlations of layer 0 and layer 1 of the MuPix8 chips.	94
9.4	Time correction between the Mupix8 and the SciFi detector.	94
9.5	Space correlation between SciFi and MuPix.	95
9.6	Substraction of full space correlation and background.	95
10.1	Floorplan of the MAMI facility.	97
10.2	Sketch of the setup used for the MAMI integration testbeam.	98
10.3	Hitmaps of the first MuPix8 chips of the two FEBs.	99
10.4	Time correlations of chip 0 and chip1.	100
10.5	Time synchronisation with reconfiguration.	100
10.6	Time synchronisation without reconfiguration.	101
11.1	Picture of the Mu3e Integration Run 2021 detector prototypes.	103
11.2	Sketch of the Mu3e DAQ system used in the Integration and Cosmic Run.	104
11.3	Chip to chip correlation of the inner vertex detector using the Mu3e Target.	105
11.4	Working chips during the Mu3e Integration Run 2021.	106
11.5	Column to column correlations of the vertex detector using the Mu3e Target.	107

LIST OF FIGURES

11.6	Time correlations of inner vertex detector using the Mu3e Target.	108
11.7	Picture of the μ SR target.	108
11.8	Chip to chip correlations of the vertex detector using the μ SR target.	109
11.9	Column to column correlations of the vertex detector using the μ SR target.	109
11.10	Time correlations of inner vertex detector using the μ SR target.	110
12.1	Picture of the detectors used during the Mu3e Cosmic Run 2022.	111
12.2	Working chips during the Mu3e Cosmic Run 2021.	112
12.3	Sketch of the Mu3e Online Analyzer framework used at the Mu3e Cosmic Run 2022.	113
12.4	Sketch of the Mu3e Event Display.	114
12.5	ToT versus time difference of the vertex detector and the scintillating panels.	115
12.6	The time correlation between the inner vertex detector and the SciFi detector.	116
13.1	Sketch of the detectors used during the Tile Integration Run 2022.	117
13.2	Rate map of one tile module.	118
13.3	ToT versus time difference of the pixel chips and the tile module.	119
13.4	The time correlation between three half ladders and one tile module.	120
14.1	Sketch of the MAMI irradiation testbeam.	123
14.2	Picture of the setup used for the MAMI irradiation testbeam.	124
14.3	Hitmaps of the four detection layers of the MuPix10 telescope.	125
14.4	Rate versus Wehnelt voltage and normalised rate map.	126
14.5	Sketch of the coordinate system for the MuPix telescope.	127
14.6	Tracks and efficiency plots of the DUT.	128
14.7	Efficiency loss of an irradiated MuPix8 chip.	129
14.8	Different irradiation spots for the 45 mV threshold having HV on and off.	130
14.9	Efficiency difference vs. dose for different thresholds and HV on or off.	131
14.10	Different noise sources observed after irradiation of around 80 Gy.	132
15.1	Sketch of the Si-Pixel based μ SR spectrometer.	137
15.2	Monte-Carlo simulations for the μ SR detector prototype.	138
15.3	Picture of the μ SR detector prototype.	139
15.4	Hitmaps and normalised ToT distributions of the four Quad-Modules.	140
15.5	Comparison of the Quad-Modules and the GPS apparatus.	142
B.1	Beam rate tests.	152
B.2	Sketch MuTRiG chip.	153
B.3	State machine of the hit-time-alignment firmware.	154
B.4	Space correlations of layer 1, layer 2 and layer 3 of the MuPix8 chips.	155
B.5	Hitmaps second MuPix8 chips.	156
B.6	Row to row correlation Mu3e Target.	157
B.7	Threshold scan of the 100 μ m DUT at -10 V.	158
B.8	Different irradiation spots for the 248 mV threshold having HV on and off.	158
B.9	Decays from different noise sources.	159
B.10	Noise scans with less and higher than 80 Gy.	159

List of Tables

1.1	List of the 19 free parameters of the SM.	6
3.1	Choice of disparity.	26
3.2	Comma words in 8b/10b data stream.	27
3.3	Supported PCS types.	29
3.4	PCIe link performance.	33
3.5	PCIe write request.	33
5.1	Data rate estimation from the detector simulation.	61
6.1	Structure of the MuPix data protocol.	66
6.2	Structure of the MuTRiG hit data send from the FEB to the SWB.	67
6.3	Structure of the Slow Control packet.	68
6.4	Structure of run control signals.	69
6.5	Reset link protocol.	70
7.1	Optical fibre cabling inside the Mu3e DAQ system.	74
7.2	Input mapping of the different detector regions of the central pixel SWB.	75
8.1	Structure of the farm MIDAS event header.	82
8.2	Structure of the farm MIDAS Bank Header.	82
8.3	Structure of the MIDAS BANK32A.	83

Bibliography

- [1] 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,” *Physics Letters*, vol. B716, pp. 1–29, 2012.
- [2] S. Chatrchyan *et al.*, “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC,” *Physics Letters*, vol. B716, pp. 30–61, 2012.
- [3] L. Evans and P. Bryant, “LHC Machine,” *Journal of Instrumentation*, vol. 3, no. 08, pp. S08 001–S08 001, 2008.
- [4] D. J. Griffiths, “Introduction to elementary particles; 2nd rev. version,” <https://cds.cern.ch/record/111880>, New York, NY, 2008.
- [5] D. Galbraith and C. Burgard, “Standard Model of the Standard Model,” <http://davidgalbraith.org/portfolio/ux-standard-model-of-the-standard-model/>, [Online; accessed 07-January-2024].
- [6] R. L. Workman and Others, “Review of Particle Physics,” *Progress of Theoretical and Experimental Physics*, vol. 2022, p. 083C01, 2022.
- [7] S. L. Glashow, “Partial Symmetries of Weak Interactions,” *Nuclear Physics*, vol. 22, pp. 579–588, 1961.
- [8] P. W. Higgs, “Broken Symmetries and the Masses of Gauge Bosons,” *Physical Review Letters*, vol. 13, pp. 508–509, 1964, [160(1964)].
- [9] M. Kobayashi and T. Maskawa, “CP-Violation in the Renormalizable Theory of Weak Interaction,” *Progress of Theoretical Physics*, vol. 49, no. 2, pp. 652–657, 02 1973. [Online]. Available: <https://doi.org/10.1143/PTP.49.652>
- [10] I. Aitchison and A. Hey, “Gauge Theories in Particle Physics: A Practical Introduction: From Relativistic Quantum Mechanics to QED, Fourth Edition,” 2012. [Online]. Available: <https://library.oapen.org/handle/20.500.12657/50885>
- [11] E. Fermi, E. Teller, and V. Weisskopf, “The decay of negative mesotrons in matter,” *Physical Review*, vol. 71, no. 5, p. 314, 1947.
- [12] F. Zwicky, “Die Rotverschiebung von extragalaktischen Nebeln,” *Helvetica Physica Acta*, vol. 6, pp. 110–127, 1933.
- [13] T. Kajita, E. Kearns, and M. Shiozawa, “Establishing atmospheric neutrino oscillations with Super-Kamiokande,” *Nuclear Physics*, vol. B908, pp. 14–29, 2016.
- [14] Y. Fukuda *et al.*, “Evidence for oscillation of atmospheric neutrinos,” *Physical Review Letters*, vol. 81, pp. 1562–1567, 1998.
- [15] S. N. Ahmed *et al.*, “Measurement of the total active B-8 solar neutrino flux at the Sudbury Neutrino Observatory with enhanced neutral current sensitivity,” *Physical Review Letters*, vol. 92, p. 181301, 2004.

- [16] K. Eguchi *et al.*, “First results from KamLAND: Evidence for reactor anti-neutrino disappearance,” *Physical Review Letters*, vol. 90, p. 021802, 2003.
- [17] B. Pontecorvo, “Inverse beta processes and nonconservation of lepton charge,” 1957. [Online]. Available: <https://api.semanticscholar.org/CorpusID:117294049>
- [18] Z. Maki, M. Nakagawa, and S. Sakata, “Remarks on the Unified Model of Elementary Particles,” *Progress of Theoretical Physics*, vol. 28, no. 5, pp. 870–880, 11 1962. [Online]. Available: <https://doi.org/10.1143/PTP.28.870>
- [19] Wikipedia contributors, “Standard Model — Wikipedia, The Free Encyclopedia,” https://en.wikipedia.org/w/index.php?title=Standard_Model&oldid=1179748656, 2023, [Online; accessed 07-January-2024].
- [20] L. Calibbi and G. Signorelli, “Charged lepton flavour violation: An experimental and theoretical introduction,” *La Rivista del Nuovo Cimento*, vol. 41, no. 2, pp. 71–174, Feb 2018. [Online]. Available: <https://doi.org/10.1393/ncr/i2018-10144-0>
- [21] R. R. Crittenden, W. D. Walker, and J. Ballam, “Radiative Decay Modes of the Muon,” *Physical Review*, vol. 121, pp. 1823–1832, Mar 1961. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRev.121.1823>
- [22] G. Hernández-Tomé *et al.*, “Flavor violating leptonic decays of τ and μ leptons in the Standard Model with massive neutrinos,” *The European Physical Journal C*, vol. 79, no. 1, p. 84, Jan 2019.
- [23] J. Bernabeu, E. Nardi, and D. Tommasini, “ $\mu - e$ conversion in nuclei and Z' physics,” *Nuclear Physics*, vol. B409, pp. 69–86, 1993.
- [24] E. Ma, “Verifiable radiative seesaw mechanism of neutrino mass and dark matter,” *Physical Review*, vol. D73, p. 077301, 2006.
- [25] T. Toma and A. Vicente, “Lepton Flavor Violation in the Scotogenic Model,” *Journal of High Energy Physics*, vol. 01, p. 160, 2014.
- [26] J. C. Pati and A. Salam, “Lepton number as the fourth “color”,” *Physical Review*, vol. D10, pp. 275–289, 1974. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.10.275>
- [27] R. N. Mohapatra and J. C. Pati, “Left-right gauge symmetry and an “isoconjugate” model of CP violation,” *Physical Review*, vol. D11, pp. 566–571, 1975. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.11.566>
- [28] G. Senjanovic and R. N. Mohapatra, “Exact left-right symmetry and spontaneous violation of parity,” *Physical Review*, vol. D12, no. 5, pp. 1502–1505, 1975.
- [29] R. E. Marshak and R. N. Mohapatra, “Quark-lepton symmetry and B - L as the U(1) generator of the electroweak symmetry group,” *Physics Letters*, vol. B91, no. 2, pp. 222–224, 1980.
- [30] A. Davidson, “ $B - L$ as the fourth color within an $SU(2)_L \times U(1)_R \times U(1)$ model,” *Physical Review*, vol. D20, pp. 776–783, 1979. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.20.776>
- [31] U. Chattopadhyay and P. B. Pal, “Radiative neutrino decay in left-right models,” *Physical Review*, vol. D34, no. 11, pp. 3444–3448, 1986.
- [32] H. M. Georgi, S. L. Glashow, and S. Nussinov, “Unconventional model of neutrino masses,” *Nuclear Physics*, vol. B193, no. 2, pp. 297–316, 1981.
- [33] R. Mohapatra and P. Pal, “Massive Neutrinos in Physics and Astrophysics,” 2004. [Online]. Available: <https://worldscientific.com/worldscibooks/10.1142/5024#t=aboutBook>
- [34] E. J. Chun, K. Y. Lee, and S. C. Park, “Testing Higgs triplet model and neutrino mass patterns,” *Physics Letters*, vol. B566, no. 1, pp. 142–151, 2003. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0370269303007706>

BIBLIOGRAPHY

- [35] M. Senami and K. Yamamoto, “Lepton flavor violation with supersymmetric Higgs triplets in the TeV region for neutrino masses and leptogenesis,” *Physical Review*, vol. D69, p. 035004, 2004. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.69.035004>
- [36] J. Bernabéu, E. Nardi, and D. Tommasini, “ $\mu - e$ conversion in nuclei and Z' physics,” *Nuclear Physics*, vol. B409, no. 1, pp. 69–86, 1993. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/055032139390446V>
- [37] L. Randall and R. Sundrum, “Large Mass Hierarchy from a Small Extra Dimension,” *Physical Review Letters*, vol. 83, pp. 3370–3373, 1999. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.83.3370>
- [38] N. Arkani-Hamed and M. Schmaltz, “Hierarchies without symmetries from extra dimensions,” *Physical Review*, vol. D61, p. 033005, 2000. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevD.61.033005>
- [39] T. Mori, “Final Results of the MEG Experiment,” *Nuovo Cimento*, vol. C39, no. 4, p. 325, 2017.
- [40] U. Bellgardt *et al.*, “Search for the decay $\mu^+ \rightarrow e^- e^+ e^+$,” *Nuclear Physics*, vol. B299, no. 1, pp. 1–6, 1988.
- [41] P. Wintz, “Results of the SINDRUM-II experiment,” *Conference Proceedings*, vol. C980420, pp. 534–546, 1998.
- [42] M. Ardu and G. Pezzullo, “Introduction to charged lepton flavor violation,” *Universe*, vol. 8, no. 6, p. 299, 2022.
- [43] A. M. Baldini *et al.*, “Search for the lepton flavour violating decay $\mu^+ \rightarrow e^+ \gamma$ with the full dataset of the MEG experiment: MEG Collaboration,” *The European Physical Journal C*, vol. 76, pp. 1–30, 2016.
- [44] —, “The design of the MEG II experiment: MEG II Collaboration,” *The European Physical Journal C*, vol. 78, pp. 1–60, 2018.
- [45] Y. Kuno, “A search for muon-to-electron conversion at J-PARC: The COMET experiment,” *Progress of Theoretical and Experimental Physics*, vol. 2013, no. 2, p. 022C01, 2013.
- [46] L. Bartoszek *et al.*, “Mu2e technical design report,” eScholarship, University of California, Tech. Rep., 2015.
- [47] M. Hildebrandt, “The drift chamber system of the MEG experiment,” *Nuclear Instruments and Methods*, vol. A623, no. 1, pp. 111–113, 2010.
- [48] MEG II collaboration, “A search for $\mu^+ \rightarrow e^+ \gamma$ with the first dataset of the MEG II experiment,” 2023. [Online]. Available: <https://arxiv.org/abs/2310.12614>
- [49] M. Thomson, “Modern particle physics,” <http://www-spires.fnal.gov/spires/find/miscs/www?cl=QC793.2.T46::2013>, New York, 2013.
- [50] G. F. Knoll, “Radiation detection and measurement; 4th ed.” <https://cds.cern.ch/record/1300754>, New York, NY, 2010.
- [51] K. Arndt *et al.*, “Technical design of the phase I Mu3e experiment,” *Nucl. Instrum. Methods Phys. Res.*, vol. 1014, p. 165679, 2021.
- [52] I. Peric, “A novel monolithic pixelated particle detector implemented in high-voltage CMOS technology,” *Nuclear Instruments and Methods*, vol. A582, pp. 876–885, 2007.
- [53] M. Köppel, “Data Flow in the Mu3e Filter Farm,” Master’s thesis, Johannes Gutenberg-Universität Mainz, 2019.
- [54] Intel Corporation, “Intel Quartus Prime Standard Edition User Guide,” <https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug-qps-getting-started.pdf>.

- [55] A. P. Clark, "Principles of Digital Data Transmission," New York, NY, USA, 1976.
- [56] S. J. Orfanidis, "Electromagnetic waves and antennas," <http://eceweb1.rutgers.edu/~orfanidi/ewa/>, 2016, [Online; accessed 07-January-2024].
- [57] W. Demtröder, "Experimentalphysik 2: Elektrizität und Optik," ISBN: 978-3-540-20210-3, 01 2004.
- [58] H. Johnson and M. Graham, "High-speed Signal Propagation: Advanced Black Magic," Upper Saddle River, NJ, USA, 2003. [Online]. Available: <https://dl.acm.org/doi/10.5555/1405687>
- [59] Texas Instruments, "Interface Circuits for TIA/EIA-644 (LVDS)," https://www.ti.com/lit/an/slla038b/slla038b.pdf?ts=1704740175409&ref_url=https%253A%252F%252Fwww.google.com%252F, [Online; accessed 08-January-2024].
- [60] ALTERA, "The Evolution of High-Speed Transceiver Technology," https://www.altera.com/content/dam/altera-www/global/en_us/pdfs/literature/wp/wp_hs_transceiver.pdf, [Online; accessed 08-January-2024].
- [61] H. W. Johnson and M. Graham, "High-speed Digital Design: A Handbook of Black Magic," Upper Saddle River, NJ, USA, 1993. [Online]. Available: <https://vdoc.pub/documents/high-speed-digital-design-a-handbook-of-black-magic-2siqof9089m0>
- [62] A. X. Widmer and P. A. Franzaszek, "A DC-balanced, Partitioned-block, 8B/10B Transmission Code," *IBM Journal of Research and Development*, vol. 27, no. 5, pp. 440–451, 1983.
- [63] Wikipedia contributors, "8b/10b encoding," https://en.wikipedia.org/w/index.php?title=8b/10b_encoding&oldid=910321009, 2019, [Online; accessed 07-January-2024].
- [64] Intel Corporation, "Intel Arria 10 Transceiver PHY User Guide," https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/hb/arria-10/ug_arria10_xcvr_phy.pdf, [Online; accessed 07-January-2024].
- [65] PCI-SIG, "PCI Express Base Specification Revision 3.0," <https://pcisig.com/specifications>, 2010, [Online; accessed 07-January-2024].
- [66] National Instruments, "Guide To Fiber Optics and Premises Cabling," <https://www.thefoa.org/tech/ref/appln/transceiver.html>, 2012, [Online; accessed 07-January-2024].
- [67] I. V. Narsky, "Estimation of upper limits using a Poisson statistic," *Nuclear Instruments and Methods*, vol. A450, pp. 444–455, 2000.
- [68] D. vom Bruch, "Pixel Sensor Evaluation and Online Event Selection for the Mu3e Experiment," Ph.D. dissertation, Ruprecht-Karls-Universität Heidelberg, 10 2017.
- [69] Xillybus Ltd., "How PCI express devices talk (Part 1)," <http://xillybus.com/tutorials/pci-express-tlp-pcie-primer-tutorial-guide-1>, [Online; accessed 07-January-2024].
- [70] PCI-SIG, "PCI Express Base Specification Revision 1.0," <https://pcisig.com/specifications>, 2002, [Online; accessed 07-January-2024].
- [71] —, "PCI Express Base Specification Revision 2.0," <https://pcisig.com/specifications>, 2006, [Online; accessed 07-January-2024].
- [72] —, "PCI Express Base Specification Revision 4.0, Version 1.0," <https://pcisig.com/specifications>, 2017, [Online; accessed 07-January-2024].
- [73] PCIe Special Interest Group, "PCIe Special Interest Group," <https://pcisig.com/>, [Online; accessed 07-January-2024].
- [74] Wikipedia contributors, "PCI Express," https://en.wikipedia.org/w/index.php?title=PCI_Express&oldid=912780692, 2019, [Online; accessed 07-January-2024].

BIBLIOGRAPHY

- [75] Intel Corporation, “Intel Arria 10 and Intel Cyclone 10 GX Avalon-MM Interface for PCI Express User Guide,” https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug_a10_pcie_avmm.pdf, [Online; accessed 07-January-2024].
- [76] J. N. Bahcall and R. Davis, “Solar neutrinos: A scientific puzzle,” *Science*, vol. 191, no. 4224, pp. 264–267, 1976. [Online]. Available: <https://www.science.org/doi/abs/10.1126/science.191.4224.264>
- [77] W. J. Marciano *et al.*, “Charged Lepton Flavor Violation Experiments,” *Annual Review of Nuclear and Particle Science*, vol. 58, no. 1, pp. 315–341, 2008.
- [78] H. P. Eckert, “The Mu3e Tile Detector,” Ph.D. dissertation, Ruprecht-Karls-Universität Heidelberg, 2015.
- [79] H. Chen *et al.*, “MuTRiG: a mixed signal Silicon Photomultiplier readout ASIC with high timing resolution and gigabit data link,” *Journal of Instrumentation*, vol. 12, no. 01, pp. C01 043–C01 043, jan 2017.
- [80] A. Pifer, T. Bowen, and K. Kendall, “A High Stopping Density μ^+ Beam,” *Nuclear Instruments and Methods*, vol. 135, pp. 39–46, 1976.
- [81] A. M. Baldini *et al.*, “MEG Upgrade Proposal,” *ArXiv e-prints*, 2013. [Online]. Available: <http://adsabs.harvard.edu/abs/2013arXiv1301.7225B>
- [82] G. Dal Maso, “Optimization of the High-Intensity Muon Beamlines for MEG II, Mu3e and HIMB,” Ph.D. dissertation, ETH Zürich, 2024, [unpublished].
- [83] T. Roberts, “G4Beamline.” [Online]. Available: <http://g4beamline.muonsinc.com>
- [84] F. Berg, “CMBL - A High-intensity Muon Beam Line & Scintillation Target with Monitoring System for Next-generation Charged Lepton Flavour Violation Experiments,” Ph.D. dissertation, ETH Zürich, 2017.
- [85] H. Augustin, “Development of a novel slow control interface and suppression of signal line crosstalk enabling HV-MAPS as sensor technology for Mu3e,” Ph.D. dissertation, Ruprecht-Karls-Universität Heidelberg, 2021.
- [86] W. Shen *et al.*, “A Silicon Photomultiplier Readout ASIC for Time-of-Flight Applications Using a New Time-of-Recovery Method,” *IEEE Transactions on Nuclear Science*, vol. 65, no. 5, May 2018.
- [87] E.-U. Proposal, “Novel Multimodal Endoscopic Probes for Simultaneous PET/Ultrasound Imaging for Image-Guided Interventions,” *European Union 7th Framework Programme*, vol. 186, pp. 2007–2013, 2011.
- [88] H. Chen, “A Silicon Photomultiplier Readout ASIC for the Mu3e Experiment,” Ph.D. dissertation, Heidelberg University, 2018.
- [89] B. A. *et al.*, “Development of the scintillating fiber timing detector for the Mu3e experiment,” *Nuclear Instruments and Methods*, vol. A1058, p. 168766, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S016890022300757X>
- [90] M. Köppel, “Data Flow in the Mu3e DAQ,” *IEEE Transactions on Nuclear Science*, pp. 1–1, 2023.
- [91] H. Augustin *et al.*, “The Mu3e Data Acquisition,” *IEEE Transactions on Nuclear Science*, vol. 68, no. 8, pp. 1833–1840, 2021.
- [92] K. Olchanski, “MIDAS analyzer,” [Online; accessed 07-January-2024]. [Online]. Available: <https://bitbucket.org/tmidas/manalyzer/>
- [93] N. Berger *et al.*, “A New Three-Dimensional Track Fit with Multiple Scattering,” *Nucl. Instr. Meth A.*, vol. 844, p. 135–140, 2017.

-
- [94] DIGILENT, “Genesys 2 FPGA Board Reference Manual,” https://reference.digilentinc.com/_media/reference/programmable-logic/genesys-2/genesys2_rm.pdf, [Online; accessed 07-January-2024].
- [95] H. Augustin *et al.*, “The Mu3e Data Acquisition,” *IEEE Transactions on Nuclear Science*, vol. 68, pp. 1833–1840, 10 2020.
- [96] F. M. Aeschbacher, M. Deflorinc, and L. O. S. Noehdeb, “Mechanics, readout and cooling systems of the Mu3e experiment,” *Proceedings of Science*, vol. Vertex2019, 2020.
- [97] S. Agostinelli *et al.*, “Geant4—a simulation toolkit,” *Nuclear Instruments and Methods*, vol. A506, no. 3, pp. 250–303, 2003. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900203013688>
- [98] S. Bachmann *et al.*, “The proposed trigger-less TBit/s readout for the Mu3e experiment,” *Journal of Instrumentation*, vol. 9, pp. C01 011 – C01 011, 2014.
- [99] W. P. Vazquez *et al.*, “The ATLAS Data Acquisition System: from Run 1 to Run 2,” *Nuclear and Particle Physics Proceedings*, pp. 273–275, 2016.
- [100] CMS collaboration, “CMS Technical Design Report for the Level-1 Trigger Upgrade,” 2013, [Online; accessed 07-January-2024]. [Online]. Available: <http://cds.cern.ch/record/1556311>
- [101] R. Aaij *et al.*, “Allen: A High-Level Trigger on GPUs for LHCb,” *Computing and Software for Big Scienc*, vol. 4, 2020.
- [102] ALICE collaboration, “Real-time data processing in the ALICE High Level Trigger at the LHC,” *Computer Physics Communications*, vol. 242, pp. 25–48, 2019.
- [103] Intel Corporation, “Arria V GX Starter Kit User Guide,” https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/ug/ug_avgx_starter_dev_kit.pdf, [Online; accessed 07-January-2024].
- [104] Samtec, “FireFly,” <https://www.samtec.com/de/products/ecuo>, [Online; accessed 07-January-2024].
- [105] Intel Corporation, “Intel MAX 10 FPGA,” <https://www.intel.de/content/www/de/de/products/details/fpga/max/10/docs.html>, [Online; accessed 07-January-2024].
- [106] A.-K. Perrevoort, “Sensitivity Studies on New Physics in the Mu3e Experiment and Development of Firmware for the Front-End of the Mu3e Pixel Detector,” Ph.D. dissertation, Ruprecht-Karls-Universität Heidelberg, 05 2018.
- [107] N. Berger, “Pixel Hit Time Sorting,” 2020, [Mu3e internal Wengen Meeting, March 2020]. [Online]. Available: <https://indico.psi.ch/event/8068/contributions/23584/attachments/16294/23234/Sorter.pdf>
- [108] M. Müller, “A Control System for the Mu3e DAQ,” Master’s thesis, Johannes Gutenberg-Universität Mainz, 2019.
- [109] S. Kilani, “Clock and Reset System Specification,” 2017, [Mu3e internal note 0043].
- [110] —, “Mu3e Clock & Reset Distribution,” 2019, [Mu3e internal talk 25.07.2019].
- [111] N. Berger, “Run Start and Reset Protocol,” 2018, [Mu3e internal note 0046].
- [112] G. Hesketh, “Clock & Reset Protocols and MIDAS,” 2018, [Mu3e internal talk 29.11.2018].
- [113] M. Müller, “A Control System for the Mu3e DAQ,” Ph.D. dissertation, Johannes Gutenberg-Universität Mainz, 2023, [unpublished].
- [114] P. Durante *et al.*, “100 Gbps PCI-Express readout for the LHCb upgrade,” *Journal of Instrumentation*, vol. 10, no. 04, pp. C04 018–C04 018, apr 2015.

BIBLIOGRAPHY

- [115] LHCb collaboration, “The LHCb PCIe Readout,” https://indico.cern.ch/event/468486/contributions/1144380/attachments/1241282/1825481/LHCb_PCIe_Readout.pdf, 2016, [Online; accessed 07-January-2024].
- [116] Intel Corporation, “Intel Arria 10 FPGA and SoC FPGA,” <https://www.intel.com/content/www/us/en/products/details/fpga/arria/10/docs.html>, [Online; accessed 07-January-2024].
- [117] Broadcom, “MiniPOD™12x10G Transmitter Module,” https://docs.broadcom.com/doc/AV02-4467EN_PB_AFBR-811Xx3Z_2014-03-06, [Online; accessed 07-January-2024].
- [118] —, “Optical Transceiver SFP+,” <https://www.broadcom.com/products/fiber-optic-modules-components/networking/optical-transceivers/sfpplus>, [Online; accessed 07-January-2024].
- [119] Intel Corporation, “Intel MAX V FPGA,” <https://www.intel.de/content/www/de/de/products/sku/210261/max-v-5m2210z-cpld/specifications.html?wapkw=max%20v>, [Online; accessed 07-January-2024].
- [120] Terasic Technologies Inc., “DE5a-Net FPGA Development Kit User Manual,” https://www.terasic.com.tw/cgi-bin/page/archive_download.pl?Language=English&No=970&FID=0bc2c05d074b8a05252d9b8e363d69d1, [Online; accessed 07-January-2024].
- [121] V. Henkys, B. Schmidt, and N. Berger, “Online Event Selection for Mu3e using GPUs,” *2022 21st International Symposium on Parallel and Distributed Computing (ISPDC)*, pp. 17–24, 2022.
- [122] S. Ritt, P. Amaudruz, and K. Olchanski, “Maximum Integration Data Acquisition System,” 2001, [Online; accessed 07-January-2024]. [Online]. Available: <https://midas.triumf.ca/>
- [123] R. Brun and F. Rademakers, “ROOT - An Object Oriented Data Analysis Framework,” *Nucl. Inst. & Meth. in Phys. Res. A*, vol. 389, pp. 81–86, 1997.
- [124] B. Gayther, “Preparations for Phase I of the Mu3e experiment,” Ph.D. dissertation, University College London, 2023.
- [125] V. Henkys, “Online Event Selection using GPUs for the Mu3e Experiment,” Master’s thesis, Johannes Gutenberg-Universität Mainz, 2022.
- [126] H. Murugan, “Online Track Reconstruction for the Mu3e Experiment,” <https://www.psi.ch/en/media/82327/download>, 2023, [Online; accessed 07-January-2024].
- [127] A.-K. Perrevoort, “Sensitivity Studies on New Physics in the Mu3e Experiment and Development of Firmware for the Front-End of the Mu3e Pixel Detector,” Ph.D. dissertation, Ruprecht-Karls-Universität Heidelberg, 01 2018.
- [128] R. Diener *et al.*, “The DESY II test beam facility,” *Nuclear Instruments and Methods*, vol. A922, pp. 265–286, 2019. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168900218317868>
- [129] H. Augustin *et al.*, “Performance of the large scale HV-CMOS pixel sensor MuPix8,” *Journal of Instrumentation*, vol. 14, no. 10, p. C10011, 2019.
- [130] T. Rudzki *et al.*, “Successful cooling of a pixel tracker using gaseous helium: Studies with a mock-up and a detector prototype,” *Nuclear Instruments and Methods*, vol. A1054, p. 168405, 2023.
- [131] M. Köppel, “Mu3e Integration Run 2021,” *Proceeding of science NuFact2021*, p. 233, 2022.
- [132] T. T. Rudzki, “The Mu3e vertex detector - construction, cooling, and first prototype operation,” Ph.D. dissertation, Ruprecht-Karls-Universität Heidelberg, 04 2022.

-
- [133] J. P. Cachemiche, P. Y. Duval, F. Hachon, R. Le Gac, and F. Réthoré, “The PCIe-based readout system for the LHCb experiment,” *Journal of Instrumentation*, vol. 11, no. 02, p. P02013, 2016.
- [134] H. Augustin *et al.*, “MuPix10: First Results from the Final Design,” *Proceedings of VERTEX2020*, 12 2020.
- [135] S. Bravar *et al.*, “Scintillating fibre detector for the Mu3e experiment,” *Journal of Instrumentation*, vol. 12, no. 07, p. C07011, 2017.
- [136] D. Becker *et al.*, “The P2 experiment,” *The European Physical Journal A*, vol. 54, no. 11, p. 208, Nov 2018. [Online]. Available: <https://doi.org/10.1140/epja/i2018-12611-6>
- [137] M. Zimmermann, “Particle Rate Studies and Technical Design Development for the P2 Silicon Pixel Tracking Detector,” Ph.D. dissertation, Johannes Gutenberg-Universität Mainz, 2019.
- [138] M. Benoit *et al.*, “Testbeam results of irradiated ams H18 HV-CMOS pixel sensor prototypes,” *Journal of Instrumentation*, vol. 13, no. 02, p. P02011, feb 2018. [Online]. Available: <https://dx.doi.org/10.1088/1748-0221/13/02/P02011>
- [139] C. Grzesik, “HV-MAPS for the P2 Tracking Detector,” Ph.D. dissertation, Johannes Gutenberg-Universität Mainz, 2023, [unpublished].
- [140] Niklaus Berger, “Mainz Testbeam September 2021,” <https://www.flickr.com/photos/nberger/51641890446/in/album-72157720086226423/>, 2021, [Online; accessed 07-January-2024].
- [141] H. Augustin *et al.*, “The MuPix Telescope: A Thin, high Rate Tracking Telescope,” *Journal of Instrumentation*, vol. 12, no. 01, p. C01087, 2017.
- [142] Intel Corporation, “Stratix IV Device Handbook,” <https://www.intel.com/content/www/us/en/content-details/654799/stratix-iv-device-handbook.html>, [Online; accessed 07-January-2024].
- [143] L. Huth, “A High Rate Testbeam Data Acquisition System and Characterization of High Voltage Monolithic Active Pixel Sensors,” Ph.D. dissertation, Ruprecht-Karls-Universität Heidelberg, 2018.
- [144] National Institute of Standards and Technology, “ESTAR,” <https://physics.nist.gov/PhysRefData/Star/Text/ESTAR.html>, [Online; accessed 07-January-2024].
- [145] R. Barlow, “Asymmetric Statistical Errors,” 2004. [Online]. Available: <https://arxiv.org/abs/physics/0406120>
- [146] P. Laursen *et al.*, “Lyman α -emitting galaxies in the epoch of reionization,” *A&A*, vol. 627, p. A84, 2019. [Online]. Available: <https://doi.org/10.1051/0004-6361/201833645>
- [147] L. Dittmann, “Test of an HV-CMOS Prototype for the LHCb Mighty Tracker,” Master’s thesis, Heidelberg University, 2022.
- [148] M. Aiba *et al.*, “Science Case for the new High-Intensity Muon Beams HIMB at PSI,” 2021. [Online]. Available: <https://arxiv.org/abs/2111.05788>
- [149] Z. Salman and T. Prokscha, “Advanced muon spin spectroscopy with high lateral resolution using Si pixel detectors,” <https://data.snf.ch/grants/grant/215167>, 2023, [Online; accessed 07-January-2024].
- [150] T. Bowen, “The Surface Muon Beam,” *Physics Today*, vol. 38, no. 7, pp. 22–34, 07 1985. [Online]. Available: <https://doi.org/10.1063/1.881018>
- [151] A. Amato *et al.*, “The new versatile general purpose surface-muon instrument (GPS) based on silicon photomultipliers for μ SR measurements on a continuous-wave beam,” *Review of Scientific Instruments*, vol. 88, no. 9, sep 2017. [Online]. Available: <https://doi.org/10.1063%2F1.4986045>

BIBLIOGRAPHY

- [152] F. H. Combley, “Muon Spin Rotation Spectroscopy: Principles and Applications in Solid State Physics,” *Physics Bulletin*, vol. 36, no. 10, p. 430, oct 1985. [Online]. Available: <https://dx.doi.org/10.1088/0031-9112/36/10/029>
- [153] A. Amato, “Heavy-fermion systems studied by μ SR technique,” *Reviews of Modern Physics*, vol. 69, pp. 1119–1180, Oct 1997. [Online]. Available: <https://link.aps.org/doi/10.1103/RevModPhys.69.1119>
- [154] S. J. Blundell, “Spin-polarized muons in condensed matter physics,” *Contemporary Physics*, vol. 40, no. 3, pp. 175–192, 1999.
- [155] A. Yaouanc *et al.*, “Muon Spin Rotation, Relaxation, and Resonance: Applications to Condensed Matter,” 2011. [Online]. Available: <https://api.semanticscholar.org/CorpusID:118874468>
- [156] G. Bassani, G. Liedl, and P. Wyder, “Encyclopedia of Condensed Matter Physics,” 2005. [Online]. Available: <https://www.sciencedirect.com/referencework/9780323914086/encyclopedia-of-condensed-matter-physics>
- [157] K. Sedlak *et al.*, “Geant4 Simulation of the New ALC μ SR Spectrometer,” *IEEE Transactions on Nuclear Science*, vol. 57, no. 4, pp. 2187–2195, 2010.
- [158] Z. Salman *et al.*, “HiFi—A new high field muon spectrometer at ISIS,” *Physica B: Condensed Matter*, vol. 404, no. 5, pp. 978–981, 2009. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0921452608007230>
- [159] L. Mandok, “no title,” [unpublished], Heidelberg University, 2023.
- [160] D. Dannheim *et al.*, “Corryvreckan: a modular 4D track reconstruction and analysis software for test beam data,” *Journal of Instrumentation*, vol. 16, no. 03, p. P03008, mar 2021. [Online]. Available: <https://dx.doi.org/10.1088/1748-0221/16/03/P03008>
- [161] W. Morag *et al.*, “Corryvreckan - A Modular 4D Track Reconstruction and Analysis Software for Test Beam Data,” Dec. 2020. [Online]. Available: <https://doi.org/10.5281/zenodo.4384186>
- [162] A. Suter and B. Wojek, “Musrfit: A Free Platform-Independent Framework for μ SR Data Analysis,” *Physics Procedia*, vol. 30, pp. 69–73, 2012. [Online]. Available: <https://doi.org/10.1016%2Fj.phpro.2012.04.042>

Acknowledgement

I express my gratitude to Niklaus Berger, my supervisor, for providing me the opportunity to conduct my thesis in his research group. He has been a constant source of support throughout my academic journey, starting from my master's studies and extending until the completion of this thesis. I am particularly grateful for his introduction to the Mu3e Experiment and for carefully reviewing my work. I am deeply appreciative of the excellent collaboration with Alexandr Kozlinskiy and Martin Müller, who provided valuable assistance in comprehending the challenging process of hardware development. I would also like to thank all the members of the AG Berger, who contributed to the great atmosphere that allowed me to accomplish a lot over the last years while also having a lot of fun.

I would like to express my gratitude to all individuals involved in the Mu3e collaboration for their exceptional efforts during the Mu3e Integration Run, the Cosmic Run held at PSI and the Tile Integration Run at DESY. Without their help, this thesis would not have been possible.

Special gratitude is extended to the individuals who provided proofreading assistance, namely Heiko Augustin, Konrad Briggel, Sophie Gagneur, Carsten Grzesik, Alexandr Kozlinskiy, Thomas Rudzki, Zaher Salman, and Frederik Wauters. Additionally, I would like to express my appreciation to Tim Klausmann and Janik Köppel for their support during my recuperation from my eye surgery.

During the last months of writing this thesis, I was deeply saddened by the sudden death of my very good friend Martin W. in a tragic accident with whom I shared not only my passion for climbing, but also for physics. Martin provided support throughout my studies and was a constant companion that I am grateful to have had in my life.

In conclusion, I express my gratitude to my parents and family members for their unwavering support throughout my academic journey. Finally, I extend my thanks to Pia Snella for her continuous support and love, which has been invaluable not only during the process of writing this work.