

The Mu3e Scintillating Fibre Detector

Light Attenuation, Radiation Damage, and Data
Acquisition

Doctoral Thesis

Author(s):

Gerritzen, Lukas

Publication date:

2022

Permanent link:

<https://doi.org/10.3929/ethz-b-000596779>

Rights / license:

[In Copyright - Non-Commercial Use Permitted](#)

DISS: ETH NO. 28941

The Mu3e Scintillating Fibre Detector

Light Attenuation, Radiation Damage, and Data Acquisition

A thesis submitted to attain the degree of

DOCTOR OF SCIENCES OF ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

Lukas Gerhard Gerritzen

MSc. Physics

born on 5 April 1991

in Offenbach am Main, Germany

accepted on the recommendation of

Prof. Dr. Rainer Wallny

Prof. Dr. Christophorus Grab

Prof. Dr. Angela Papa

2022

ABSTRACT

Mu3e is an experiment designed for the background-free search for the charged lepton-flavour violating decay $\mu^+ \rightarrow e^+e^+e^-$ with a single event sensitivity of $2 \cdot 10^{-15}$ in Phase I and better than 10^{-16} in Phase II, four orders of magnitude improvement over the current upper limit. As the decay is experimentally not accessible in the Standard Model of particle physics with massive neutrinos, an observation would be proof of new physics. A timing detector as part of the Mu3e detector based on long Scintillating Fibres (SciFis) and Silicon Photomultipliers (SiPMs) will suppress accidental backgrounds to a negligible level by providing a time resolution of 250 ps.

This thesis presents light attenuation measurements of scintillating and wavelength-shifting fibres focusing on distances in the scale of the 30 cm long SciFi detector. The decrease of light in short optical fibres differs from the single-exponential decay at distances above 1 m. Radiant flux and spectral measurements are used for a precise characterization of different types of round fibres, showing additional exponential components with shorter attenuation lengths.

In total, $2.5 \cdot 10^{15}$ muons must be stopped in Mu3e to reach the Phase I sensitivity. Due to the small distance to the decay target, the SciFi SiPMs are exposed to a positron fluence of up to $5.5 \cdot 10^{12} \text{ cm}^{-2}$, as shown in detailed simulations. In addition, this thesis presents the irradiation of Hamamatsu S13552 SiPMs with positrons from Michel decays. After irradiation, the SiPMs showed vastly increased leakage currents below and above breakdown, and an increased Dark Count Rate (DCR). The findings led to the incorporation of a hit validation feature in the dedicated readout ASIC MuTRiG to reduce the rate of spurious signals. Furthermore, the need to cool the SciFi SiPMs to lower the DCR was demonstrated, leading to a revision of the mechanical design of the SciFi detector.

The last part presents the integration of the SciFi detector into the Mu3e DAQ system. These efforts allowed for the first joint operation of all Mu3e subdetectors.

ZUSAMMENFASSUNG

Das Mu3e-Experiment ist für die untergrundfreie Suche nach dem geladenen Lepton-Flavor-verletzenden Zerfall $\mu^+ \rightarrow e^+e^+e^-$ mit einer Sensitivität von $2 \cdot 10^{-15}$ in Phase I und besser als 10^{-16} in Phase II konzipiert. Dies stellt eine Verbesserung um vier Größenordnungen gegenüber der aktuellen Obergrenze dar. Ein Zeitdetektor als teil des Mu3e-Detektors, der auf langen szintillierenden Fasern (SciFi) und Silizium-Photomultipliern (SiPMs) basiert, ermöglicht eine zeitliche Auflösung von 250 ps. Dadurch werden zufällige Untergrundsignale auf ein vernachlässigbares Niveau reduziert. Da der Zerfall im Standardmodell der Teilchenphysik mit massiven Neutrinos experimentell nicht beobachtbar ist, wäre eine Entdeckung ein Beweis für Neue Physik.

In dieser Dissertation werden Lichtdämpfungsmessungen an szintillierenden und wellenlängenverschiebenden Fasern präsentiert, wobei der Schwerpunkt auf Entfernungen in der Größenordnung des 30 cm langen Scintillating Fibre (SciFi)-Detektors liegt. Die Abnahme des Lichts in kurzen optischen Fasern unterscheidet sich von demjenigen bei Entfernungen über 1 m, der mit nur einer Zeitkonstante beschrieben werden kann. Messungen von Strahlungsleistung und Wellenlängenspektren werden für eine genaue Charakterisierung verschiedener Arten von runden Fasern verwendet und zeigen zusätzliche exponentielle Komponenten mit kürzeren Dämpfungslängen.

Insgesamt müssen $2.5 \cdot 10^{15}$ Myonen in Mu3e gestoppt werden, um die Sensitivität der Phase I zu erreichen. Aufgrund des geringen Abstands zum Zerfallsziel ist die SciFi Silicon Photomultiplier (SiPMs) einer Positronenfluenz von bis zu $5.5 \cdot 10^{12} \text{ cm}^{-2}$ ausgesetzt, wie in detaillierten Simulationen gezeigt wurde. Darüber hinaus wird in dieser Arbeit die Bestrahlung von Hamamatsu S13552 SiPMs mit Positronen aus Michel-Zerfällen vorgestellt. Nach der Bestrahlung zeigte das SiPMs stark erhöhte Leckströme unterhalb und oberhalb der Durchbruchspannung, sowie eine erhöhte Dunkelrate. Die Ergebnisse führten dazu, dass im Auslese-ASIC MuTRiG eine Treffer-Validierungsfunktion eingebaut wurde, um die Rate der Störsignale zu verringern. Ausserdem wurde die Notwendigkeit der Kühlung der SciFi SiPMs um den Dark Count Rate (DCR) zu senken gezeigt. Dies führte zu einer Überarbeitung des mechanischen Designs des SciFi-Detektors.

Im letzten Teil wird die Integration des SciFi-Detektors in das Mu3e DAQ-System vorgestellt. Diese Aktivitäten ermöglichten den ersten gemeinsamen Betrieb aller Mu3e-Subdetektoren.

Contents

1	THEORY OF LEPTON FLAVOUR VIOLATION	5
1.1	The Standard Model of particle physics	5
1.2	Lepton and quark flavour	6
1.3	Lepton flavour violation	7
1.4	V-A structure of the weak interaction	8
2	MU3E EXPERIMENT AND BEAM INFRASTRUCTURE	11
2.1	The Mu3e experiment	11
2.1.1	Timeline of Mu3e	11
2.1.2	The signal decay	12
2.1.3	Backgrounds	12
2.1.4	The Mu3e detector	14
2.2	Beamlines at PSI	19
2.2.1	PiM1	19
2.2.2	PiE5	20
3	SCINTILLATING FIBRES AND SILICON PHOTODETECTORS	23
3.1	Scintillating fibres	23
3.1.1	Scintillation process	23
3.1.2	Scintillating fibres	24
3.1.3	Light loss in optical fibres	25
3.1.4	Quantifying light loss	27
3.2	Silicon photodetectors	29
3.2.1	Photodiodes	30
3.2.2	Avalanche photodiodes	30
3.2.3	Silicon photomultipliers	31
3.3	Radiation damage in silicon photomultipliers	33
3.3.1	Surface damage	33
3.3.2	Bulk damage	33
3.4	Quantification of radiation damage	34
3.4.1	NIEL hypothesis	35
3.4.2	Other factors	35
3.5	Impact on Mu3e	36
4	ATTENUATION STUDIES OF SCINTILLATING FIBRES	37

4.1	Motivation	37
4.2	Setup	38
4.2.1	Fibres under study	38
4.3	Radiant flux measurements	40
4.3.1	Measurement uncertainty	40
4.3.2	Analysis	41
4.3.3	Results and discussion of the photodiode measurements	42
4.4	Wavelength-dependent attenuation	45
4.4.1	Differential attenuation length	46
4.4.2	Attenuation of integrated wavelength spectra	49
4.5	Discussion of the attenuation studies	50
4.5.1	Comparison to literature	50
4.6	UV Damage in scintillating fibres	55
4.7	Conclusions of the attenuation measurements	57
4.7.1	Experimental aspects and detector construction	58
5	IRRADIATION STUDIES OF SILICON PHOTOMULTIPLIERS	61
5.1	Introduction	61
5.2	Setup	62
5.2.1	Irradiation chamber	62
5.2.2	SiPM carrier board	64
5.2.3	Samples	65
5.2.4	Beamline	66
5.3	Dose simulation	66
5.3.1	Radiation in Mu3e	68
5.3.2	Radiation in PiE5	76
5.3.3	Normalization to the proton current	76
5.3.4	Systematic uncertainties	80
5.4	Fluence per sample	84
5.4.1	Separator Incident	85
5.5	Radiation effects	86
5.5.1	Current during irradiation	86
5.5.2	IV curves	91
5.5.3	Breakdown voltage	94
5.5.4	Waveforms	101
5.5.5	Dark count rate (DCR)	101
5.6	Annealing	106
5.6.1	Discussion and literature	107
5.7	Conclusion of the irradiation and outlook	108
5.7.1	General	109
5.7.2	Impact on Mu3e and possible mitigation	110

6	DATA ACQUISITION AND SCIFI INTEGRATION	115
6.1	General overview	115
6.2	Data producers and consumers	115
6.3	Data flow and filtering	116
6.4	DAQ hardware	118
6.4.1	MuTRiG	119
6.4.2	SciFi module board	123
6.4.3	Micro-twisted pair cables	124
6.4.4	Detector adapter board	124
6.4.5	Front-end board	125
6.4.6	Switching board	127
6.4.7	Receiver board	128
6.5	DAQ software	128
6.5.1	MIDAS	128
6.6	Slow control	128
6.6.1	Functionality	129
6.6.2	SciFi interface with MIDAS	130
6.6.3	SciFi counters and resets	130
6.7	Integration run 2021	131
6.8	Conclusions and outlook	132
7	CONCLUSIONS AND OUTLOOK	135
	APPENDICES	138
	ACKNOWLEDGMENTS	197
	ACRONYMS	199
	LIST OF FIGURES	201
	LIST OF TABLES	209
	BIBLIOGRAPHY	213

Introduction

Over the past century, humankind has come closer to the answer to Dr Faust's famous question on the workings of the world at an unprecedented pace. In less than a century, from what we call *classical physics* today, the Standard Model of particle physics (SM) of particle physics was formed that describes matter in its quantum nature and forces that act upon it. The last predicted particle was discovered in 2012 after a decades-long search. One might feel tempted to repeat the advice given to Max Planck that theoretical physics is more or less complete, and there is little point in pursuing physics. However, as good as the SM was in its prediction, it fails to describe the force that, among other things, holds our planet together: gravity. Cosmology tells us that matter as we know it makes up less than one-fifth of the matter in the universe. There must be *physics beyond the Standard Model*. Another phenomenon that SM does not describe is Lepton Flavour Violation (LFV). Leptons (electrons, muons, taus and their respective neutrinos) were assumed to have a conserved family number. The discovery of neutrino oscillation has proven that lepton flavour is not conserved in the neutral sector. Many theories that extend the SM predict the same for their charged counterparts. The discovery of Charged Lepton Flavour Violation (cLFV) would be proof of physics beyond the SM, while a fruitless search would put tighter constraints on new theories.

One of the experiments dedicated to the search for charged lepton flavour violation is Mu3e [1]. The decay of a muon to three electrons has previously been excluded up to a branching ratio of one in 10^{12} muon decays. Increasing the sensitivity by Mu3e's aim of four orders of magnitude requires the development of novel particle detectors and Data Acquisition System (DAQ) systems. Inside Mu3e, muons are stopped and decay at rest. The charged decay products are tracked with a silicon pixel tracker that allows precise vertex finding and momentum reconstruction via the curvature of the trajectory in a 1 T magnetic field.

Throughout my doctorate, I worked on the Mu3e SciFi subdetector of Mu3e. The purpose of the SciFi detector is to provide sub-nanosecond time information to the reconstructed particle trajectories. When charged particles pass through the fibres, they stimulate light emission. The scintillation light is detected with SiPMs with single-photon sensitivity.

The first part of this thesis consists of an overview of the motivation for cLFV searches (chapter 1), a description of the Mu3e detector and the beam infrastructure at Paul Scherrer Institute (PSI) (chapter 2), and more in-depth information on scintillating fibres and SiPMs (chapter 3). In the following, three different chapters present my research results and contributions to the development of the experiment.

Chapter 4 revolves around scintillating fibres and the attenuation of light therein. Most literature on scintillating fibres considers attenuation only over long distances where a simple exponential decrease can describe it. It is known, however, that at shorter distances, more complex attenuation phenomena play a role (e.g. [2]). As the SciFi detector with a length of 30 cm falls into that range, its attenuation is poorly modelled by a single exponential function.

I performed detailed measurements of the attenuation patterns in the fibres used in Mu3e and six other round fibre types with a diameter of 250 μm . The fibres are stimulated with a Light Emitting Diode (LED) at varying distances up to 5 m from one fibre end. The attenuation lengths at long and short scales are deduced from the light output detected at the fibre end. Knowledge of these attenuation lengths helps model the performance of the Mu3e detector and allows comparison of the different fibre types under study.

The particle detectors must withstand the harsh environment of ionizing radiation inside the experiment. The main worry for the SciFi detector is the radiation hardness of its SiPMs. During Mu3e Phase I, a positron fluence of more than $5 \cdot 10^{12} \text{ cm}^{-2}$ is expected (see appendix B.1).

Like other silicon-based electronics, SiPMs suffer from two types of radiation damage caused by different types of energy transfer [3]. The first type is caused by the total ionizing dose deposited in the silicon that changes the charge distribution on the silicon interfaces. The second type describes modifications to the crystal structure of the silicon itself due to energy transfer to silicon nuclei large enough to displace silicon atoms in the silicon bulk. Ionizing dose is deposited by the irradiation with photons, while bulk damage is induced by the irradiation with neutrons. These two types are often studied separately, even though most radiation inflicts both types of damage.

A common standard comparing radiation damage from different particles and energies is the equivalent fluence of neutrons with an energy of 1 MeV. This standard is based on the hypothesis that radiation damage grows linearly with the deposited non-ionizing energy, also known as the non-ionizing energy loss (NIEL) hypothesis. Studies have shown that the type and distribution of defects in silicon depend on the particle type that inflicts the damage, violating the NIEL hypothesis [4]. The NIEL hypothesis is particularly flawed when describing bulk damage induced by electrons. Therefore, modifications, such as “effective NIEL”, have been proposed to describe the damage better [5, 6]. Additionally, the modelled damage depends on the minimum energy needed to displace a silicon atom in the crystal. A wide range of thresholds is used in models, partially accounting for thermal annealing effects [7].

The Hamamatsu S13552 SiPM column array used in Mu3e was originally developed for the LHCb collaboration. In the scope of LHCb, samples of this SiPM model were irradiated with neutrons [8]. No studies of high-fluence irradiation with Michel positrons of any SiPM are found in the literature. With a technology under active development,

such as SiPMs, it is not straightforward to make precise predictions on the behaviour after irradiation based on a different model.

Due to the difference in particles and SiPM models in most studies, one can only make rough predictions for the expected radiation hardness in Mu3e. Chapter 5 presents a radiation hardness study with the SiPM array used in Mu3e and Michel positrons obtained at the PiE5 beamline at PSI.

Finally, chapter 6 reflects my contributions to the Mu3e DAQ framework. An experiment as large as Mu3e requires joint development efforts. Within the team of DAQ developers, I worked on the integration of the SciFi detector. My contributions are presented as part of the current status of the DAQ system. A strong emphasis lies on the SciFi detector for which I designed hardware, wrote software and participated in extensive tests of the DAQ system. The 2021 integration run at PSI marks the interim culmination of the DAQ integration efforts [9]. For the first time, a vertical slice of the Mu3e DAQ was operated at PSI. The integration run paved the way for subsequent efforts.

CONTENTS

1

Theory of lepton flavour violation

This chapter briefly introduces the theoretical motivation behind the search for the decay $\mu^+ \rightarrow e^+e^+e^-$, which is the main objective of the Mu3e experiment. It also contains the relevant theory used in later chapters.

1.1 THE STANDARD MODEL OF PARTICLE PHYSICS

The Standard Model of particle physics (SM) describes all known elementary matter particles (fermions) and three of the four fundamental interactions [10, 11]. Each interaction is described by a quantum field theory, in which force-carrying bosons couple to particles with the corresponding charge. The electromagnetic interaction between electrically charged particles is described by quantum electrodynamics (QED), by the exchange of virtual photons. The theory describing the strong interaction via gluons is quantum chromodynamics (QCD). Only quarks and gluons carry the corresponding colour charge. All fundamental fermions in the SM participate in the weak interaction. The corresponding charge is the weak isospin, and the force is mediated by electrically charged W^\pm and neutral Z^0 bosons. Gravitation is the only known fundamental interaction the SM cannot describe.

Fermions come in two groups, quarks and leptons. Each of these groups can be sorted into three generations, containing two fermion flavours each (see table 1.1). Down-type quarks (d, s, b) have an electric charge of $-\frac{1}{3}$, while up-type quarks (u, c, t) have a charge of $+\frac{2}{3}$, here and in the following as multiples of the (positive) electron charge $\approx 1.6 \cdot 10^{-19}$ C.

Table 1.1: Quark and lepton flavours

Generation	Quarks			Leptons		
	I	II	III	I	II	III
	d	s	b	e^-	μ^-	τ^-
	u	c	t	ν_e	ν_μ	ν_τ

1.2 LEPTON AND QUARK FLAVOUR

Only interactions with W^\pm bosons change the flavour. Conserving the charge, the W^\pm boson couples an up-type quark to a down-type quark. The coupling of quarks from different generations is described by the Cabibbo-Kobayashi-Maskawa (CKM) matrix. Leptons, however, are only coupled to their respective partner within a generation. The number of leptons within one generation is therefore conserved. For every generation, the *lepton family number* L_e , L_μ , and L_τ denotes the number of leptons. Leptons are assigned a value of $L = +1$, and antileptons are assigned $L = -1$ for their respective family number (see table 1.2).

Table 1.2: Left-chiral leptons (weak eigenstates) and selected properties [12]

symbol	electric charge Q	mass	weak isospin T_3	L_e	L_μ	L_τ
Generation I						
e	-1	$-\frac{1}{2}$	1	0	0	0
ν_e^-	< 1 eV	511 keV	0	$+\frac{1}{2}$	1	0
Generation II						
μ^-	106 MeV	-1	$-\frac{1}{2}$	0	1	0
ν_μ	< 1 eV	0	$+\frac{1}{2}$	0	1	0
Generation III						
τ^-	1.78 MeV	-1	$-\frac{1}{2}$	0	0	1
ν_τ	< 1 eV	0	$+\frac{1}{2}$	0	0	1




Figure 1.1: A Feynman diagram showing the decay $\mu \rightarrow eee$ via neutrino oscillation. This process is forbidden in the SM and strongly suppressed in the SM with massive neutrinos.

1.3 LEPTON FLAVOUR VIOLATION

In violation of SM predictions, it was experimentally observed that neutrinos possess the ability to change flavour [13–15]. Modifications to the SM explaining this so-called neutrino oscillation require a non-zero neutrino rest mass. In the following, SM refers to the Standard Model with massless neutrinos and no neutrino oscillation, while models with neutrino oscillation will be denoted as ν SM.

The violation of lepton family number conservation is referred to as Lepton Flavour Violation (LFV). Experimentally, LFV has only been observed for neutrinos, the (electrically) neutral sector. Charged Lepton Flavour Violation (cLFV), i.e. LFV by charged leptons, is technically possible in the ν SM, although the branching ratios for such processes are below 10^{-54} and are therefore experimentally inaccessible [16]. Observing cLFV experimentally would be proof of physics beyond the ν SM. An example cLFV process is shown in figure 1.1.

The earliest searches for cLFV predate the SM [17]. Many extensions of the SM predict branching fractions of cLFV decays that are large enough to be experimentally accessible. Examples for such extensions are a second Higgs doublet or supersymmetric models [18].

This chapter highlights searches in (rare) muon decays. Other searches for cLFV include the study of B -meson decays [19] and τ decays [20] or the spontaneous conversion of muonium to antimuonium $\mu^+e^- \rightarrow \mu^-e^+$. There are three *golden channels* of cLFV in muon decays:

$\mu^+ \rightarrow e^+\gamma$: The search for the cLFV decay $\mu^+ \rightarrow e^+\gamma$ was one of the first studies of muon decays [21]. If the muon decays at rest, both daughter particles must carry the same momentum and a total energy corresponding to the muon’s rest mass. As of 2022, the MEG collaboration has set the best limit on the branching ratio of less than $4.2 \cdot 10^{-13}$ (90 % C.L.) [22]. In 2021, the successor MEG II started taking data to increase the sensitivity even further. No other experiment will have a comparable sensitivity in this decay channel in the foreseeable future.

$\mu^- \rightarrow e^-$ CONVERSION: The spontaneous conversion of negative muons to electrons in the vicinity of nuclei $\mu^- N \rightarrow e^- N$ has been searched for in several different materials, e.g. gold [23]. No event was observed, and an upper limit (90 %) on the branching ratio of $7 \cdot 10^{-13}$ was set. Currently, two experiments are under construction, aiming to surpass the sensitivities of past searches by several orders of magnitude: Mu2e at Fermilab [24] and COMET at J-PARC [25]. An electron originating from this process would carry most of the energy from the muon $E_e \lesssim m_\mu c^2$, distinguishing it very clearly from all known processes in which $E_e \lesssim \frac{m_\mu}{2} c^2$.

$\mu^+ \rightarrow e^+ e^+ e^-$ The third channel is the neutrinoless decay of a muon to three electrons, currently excluded to $1 \cdot 10^{-12}$ (90 % C.L.) [26]. This thesis is part of the effort to develop a modern experiment with a target Single Event Sensitivity (SES)¹ of $2 \cdot 10^{-15}$ on the branching fraction, Mu3e [1]. Phase II of Mu3e, not pursued further herein, aims at an SES of less than 10^{-16} [1]. The Mu3e experiment is described in detail in chapter 2. As of 2022, Mu3e is the only experiment to make high-sensitivity measurements of this decay in the foreseeable future.

The current upper limits of these three golden channels and the predicted sensitivities of future searches are summarized in table 1.3. A graphical summary of historical searches and the planned improvements is shown in figure 1.2.

Table 1.3: Current and planned upper branching ratio limits of the three golden channels (all at 90 % C.L.)

Channel	Present limit (90 % C.L.)	Planned SES
$\mu^+ \rightarrow e^+ e^+ e^-$	$1.0 \cdot 10^{-12}$ [26]	$\leq 10^{-16}$ [1]
$\mu^+ \rightarrow e^+ \gamma$	$4.2 \cdot 10^{-13}$ [22]	$2.6 \cdot 10^{-14}$ [27] ⁶
$\mu^- N \rightarrow e^- N$	$7 \cdot 10^{-13}$ [23]	$2.6 \cdot 10^{-17}$ [25], $2.87 \cdot 10^{-17}$ [24]

1.4 V-A STRUCTURE OF THE WEAK INTERACTION

The probability of particle interactions depends on a so-called vertex factor. Without going too much into detail, the vertex factor of the weak interaction has the property that it is zero for particles with right-handed chirality. A more comprehensive explanation of why the weak interaction has this property can be found in textbooks such as [11]. The relevant property is the so-called V-A (vector minus axial vector) structure of the exchange of W bosons in the weak interaction.

Chirality is a Lorentz-invariant quantum property. A particle has one of two possible chirality states: either left-handed or right-handed. It is closely connected to the

¹SES = $1/(N_\mu \cdot \epsilon)$ with the number of muon stops N_μ and the detection efficiency $\epsilon \approx 20\%$




Figure 1.2: Historical and projected upper limits for the three golden cLFV channels, adapted from [28], originally from [29]. The predictions are upper limits at 90 % if the experiments do not observe an event with the SES listed in table 1.3.

Lorentz-non-invariant helicity. A particle has a right-handed helicity state if its spin is parallel to its momentum and a left-handed helicity state if it is antiparallel. For massless particles, helicity and chirality are equal. In the SM, where neutrinos are massless, only left-handed neutrinos and right-handed antineutrinos interact with other particles, as they possess no electrical or colour charge.

In the two-body decay of a positive pion at rest

$$\pi^+ \rightarrow \mu^+ \nu_\mu$$

both daughter particles will be emitted back-to-back. Pions have a spin of zero, which is conserved in the decay. With the neutrino being left-handed (both in chirality and helicity), the antimuon's spin must be opposite of the spin of the neutrino. Therefore, it is in a left-handed helicity state. A beam of antimuons from pion decays is, therefore, polarized—the spins of all antimuons are aligned.

In the Michel decay [30] of an antimuon

$$\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu,$$

the muon antineutrino is right-handed, and the electron neutrino is left-handed. In combination with the V-A structure of the weak interaction, the spin of the antimuon defines a preferred direction under which the positron is emitted. The differential decay width is a function of the angle θ_e between the muon spin and the positron momentum. It is typically parametrized in terms of the energy fraction of the electron $x = E_e \cdot \frac{2m_\mu}{(m_\mu^2 + m_e^2)} \approx 2E_e/m_\mu$ [31] as:

$$\frac{d^2\Gamma}{dx d\cos\theta_e} \propto x^2 [(3 - 2x) + \cos\theta_e(2x - 1)] \quad (1.1)$$

For large E_e , the highest probability corresponds to $\theta_e = 0$, meaning that positrons with higher energies are more likely to be emitted in the direction of the muon spin. Positrons with $x < 0.5$, i.e. $E_e < m_\mu/4$ are preferably emitted in the opposite direction of the muon spin. Integrating equation 1.1 over x yields

$$\frac{d\Gamma}{d\cos\theta_e} \propto 1 + \frac{1}{3} \cos\theta_e. \quad (1.2)$$

Overall more Michel positrons are emitted in the direction parallel to the muon spin.

The V-A structure of the weak interaction leads to polarized muon beams if the generating pion decays at rest. This is the case for surface muon beams as they are used at Paul Scherrer Institute (PSI) and in the Mu3e experiment (see section 2.2).

2

Mu3e experiment and beam infrastructure

Mu3e [1] is an experiment designed for the search for the lepton-flavour violating decay $\mu^+ \rightarrow e^+e^+e^-$. The aim is to either observe this decay or exclude a branching fraction of $1 \cdot 10^{16}$ or greater at the 90 % confidence level. The experiment is planned in two phases. In Phase I, an SES of $2 \cdot 10^{-15}$ is pursued using the currently available PiE5 beamline at PSI. If the decay is not discovered in Phase I, the SES translates to an upper limit of $4.6 \cdot 10^{-15}$ (90 % C.L.). In total, $2.5 \cdot 10^{15}$ muons must be stopped inside the Mu3e detector to achieve this sensitivity. Phase II requires a muon rate only a future beamline can offer to achieve the goal sensitivity. The High-Intensity Muon Beams (HIMB) project at PSI will offer such a muon rate [28, 32]. This work revolves around the Phase I detector.

2.1 THE MU3E EXPERIMENT

The experimental concept is to stop surface antimuons (see section 2.2) on a stopping target and observe charged decay products. Charged particles produce signals in a four-layer pixel tracker, called hits. Combining hits to tracks allows the reconstruction of the three-dimensional particle trajectory. The particle momentum is reconstructed from the curvature in a homogeneous 1 T magnetic field parallel to the beam.

In the following, the word “muon” refers to positively charged antimuons. The word “electron” refers to electrons and positrons unless the charge is explicitly stated.

2.1.1 TIMELINE OF MU3E

The timeline of the Mu3e experiment as of December 2021 is shown in table 2.1. Commissioning of all detector components is planned for 2023 and early 2024. In 2024,

engineering runs and first data taking are planned, after which stable operation is expected. In parallel to Mu3e Phase I, upgrades to the detector (Phase II) and a new beamline (HiMB) are planned, necessary for reaching the goal sensitivity of $1 \cdot 10^{-16}$.

Table 2.1: Time schedule of Mu3e as of December 2021 [33]

	2021	2022	2023	2024	2025	2026	2027	2028	≥ 2029
Mu3e Phase I 1st 4 years	construction & commissioning first data								
Mu3e Phase I 2nd 4 years					operation & high sensitivity preparation HiMB				
Mu3e Phase II	R&D				R&D				upgraded and extended experiment at HiMB

2.1.2 THE SIGNAL DECAY

The $\mu^+ \rightarrow e^+e^+e^-$ decay signal consists of three electrons originating from a common vertex in space and coincident in time. As the muon decays at rest, the electron momenta add up to zero:

$$\sum_{eee} \vec{p} = \vec{0}$$

The energies of the three electrons add up to the muon mass:

$$\sum_{eee} E = m_\mu$$

Here, and in the following, natural units are implied, i.e. the speed of light and the reduced Planck constant are defined as $c = \hbar = 1$.

The kinematics are highly model-dependent and a priori unknown. The acceptance, i.e. the fraction of detectable signal decays, depends on the minimum transverse electron momentum $p_{T,\min}$ that can be reconstructed. For $p_{T,\min} > \frac{m_\mu}{3}$, no signal can be detected due to the conservation of energy and momentum. Based on [31], different effective LFV models were compared. An acceptance of more than 50 % requires a $p_{T,\min} \gtrsim 10$ MeV for the models presented in [1].

2.1.3 BACKGROUNDS

The experiment is designed to run background free. Only the vSM decay $\mu^+ \rightarrow e^+e^+e^-$ in figure 1.1 mimics the exact signal signature and would therefore constitute an

irreducible background. Due to its branching fraction of $< 10^{-54}$, this decay is so rare that it does not constitute a potential background. In the following, potential sources of reducible background are portrayed, together with means of suppressing them below the target sensitivity.

INTERNAL PHOTON CONVERSION

The radiative muon decay with internal conversion of a virtual photon $\mu^+ \rightarrow e^+e^+e^-\nu_e\bar{\nu}_\mu$ has a branching fraction of $(3.4 \pm 0.4) \cdot 10^{-5}$ [34]. This decay differs from signal decays in the reconstructed rest mass $m_{eee} < m_\mu$ and total momentum $\sum_{eee} \vec{p} \neq 0$. However, if the neutrinos carry little momentum, such a decay poses a potential background. To suppress this background below $2 \cdot 10^{-15}$, the resolution of the reconstructed rest mass of the muon has to be better than 1.0 MeV. The kinematics of this decay have been studied in detail [35, 36]. When the neutrinos carry little energy, the kinematics resemble a back-to-back topology in which one positron carries half of the energy and, in the opposite direction, the remaining energy is split between the other positron and the electron.

High-precision measurements of internal conversion decays allow for tests of SM predictions. Therefore, this type of decay is not only a background source but also interesting on its own. An experiment optimized for the search for $\mu^+ \rightarrow e^+e^+e^-$ is also suited to study the radiative decay with internal conversion.

COMBINATORIAL BACKGROUNDS

Combinatorial backgrounds consist of particles from independent processes that, in combination, mimic signal decays. By far, the most common muon decay is the Michel decay $\mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu$ with a branching ratio of almost 100 % [12]. High-precision measurements of the kinematics¹ of this decay allow searches for physics beyond the vSM. In addition to positive positrons, which exist in abundance from Michel decays, a negative electron is necessary to mimic a signal decay. Apart from radiative decays with internal conversion, no SM decay of μ^+ produces particles with a negative charge. Photons from the radiative muon decay $\mu^+ \rightarrow e^+\nu_e\bar{\nu}_\mu\gamma^*$, $\gamma^* \rightarrow e^+e^-$ are a source of e^- contributing to combinatorial backgrounds. When a positron scatters on electrons in detector material (Bhabha scattering $e^+e^- \rightarrow e^+e^-$), electrons with sufficient transverse momentum ($\gtrsim 10$ MeV) become candidates for combinatorial backgrounds. Both photon conversion and Bhabha scattering are suppressed by minimizing the detector and target material.

As particle trajectories are reconstructed by connecting individual hits in a tracker, a positron could be misreconstructed as a negative electron by connecting the hits in

¹Ultimately, the complete form of equation 1.1 is tested.

reverse order. Reliable charge identification of such tracks is realised by measuring the time of flight which requires a time resolution of 500 ps per track. Backwards reconstruction is additionally suppressed by measuring the energy loss along a particle's trajectory. Good time resolution suppresses all types of combinatorial backgrounds by identifying non-simultaneous muon decays. A time resolution of better than 100 ps per track allows for suppressing combinatorial backgrounds to a negligible level [1]. The better the track and vertex resolutions, the lower the probability of falsely combining particles from spatially separated processes. If the decay vertex is additionally resolved in space, $\mu^+ \rightarrow e^+e^+e^-$ candidates that do not lie on the stopping target can be rejected.




Figure 2.1: Schematic of the Mu3e Phase I detector in lateral and transverse view. Due to a homogeneous magnetic field, positrons and electrons (red and blue) describe helical trajectories that are reconstructed from hits in the four layers of the pixel tracker.

2.1.4 THE MU3E DETECTOR

The design of the Mu3e detector [1] follows the requirements imposed by the signal reconstruction and the different backgrounds. The detector consists of one central and two recurl stations (see figure 2.1). All subdetectors are barrel-shaped and concentric around the beam axis. The coordinate system of Mu3e originates in the centre of the stopping target. The z -axis lies along the beam direction, the y -axis points straight up, and the x -direction follows a right-handed coordinate system convention.

All subsystems are designed to contain as little material as possible to minimize multiple scattering. The stronger the scattering of positrons, the worse the spatial and momentum resolution. In the worst case, strong scattering leads to inefficiency of track reconstruction since the hits along the trajectory cannot be combined meaningfully. The experiment is embedded in a gaseous helium atmosphere at atmospheric pressure to reduce scattering further and provide cooling for the detectors.

When a particle is deflected by an angle Ω , multiple scattering by an angle θ_{MS} limits the momentum resolution, as shown in figure 2.2a. In the first order, the relation is

$$\frac{\sigma_p}{p} \propto \frac{\theta_{\text{MS}}}{\Omega}. \quad (2.1)$$

However, if the particle makes half a turn, the effect of multiple scattering in the last layer cancels out in the first order (see figure 2.2b). The Mu3e detector exploits this fact, and its size is chosen to optimize the average momentum resolution.

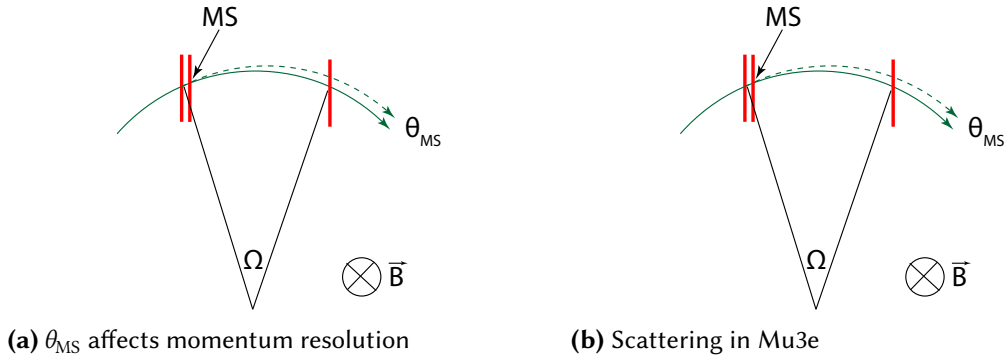


Figure 2.2: When the particle momentum is deduced by the deflection Ω , multiple scattering by θ_{MS} limits the momentum resolution. For half a turn, the effect of θ_{MS} vanishes in the first order. Figures from [1].

PIXEL TRACKER

The pixel tracker [37] is responsible for detecting the trajectories of charged particles coming from the stopping target. The deflection in the 1 T magnetic field parallel to the beam allows for momentum reconstruction. As the momentum resolution deteriorates with multiple scattering, the pixel tracker is designed to have a low material budget of 0.115 % of a radiation length per layer. This is possible thanks to the HV-CMOS technology [38] allowing thin monolithic sensors with a much lower material budget than hybrid pixel chips that typically have a thickness of 1 % of a radiation length.

The pixel sensors are bonded to a thin, high-density interconnect and glued on a polyimide support structure. The maximum power dissipation of the individual sensors is 350 mW/cm^2 . Cooling of the sensors is realized via helium flows between the detector layers [39]. The central pixel tracker comprises four layers with radii between 23.3 and 86.3 mm. Each recurl station has two pixel detector layers, identical to the outer central layers.

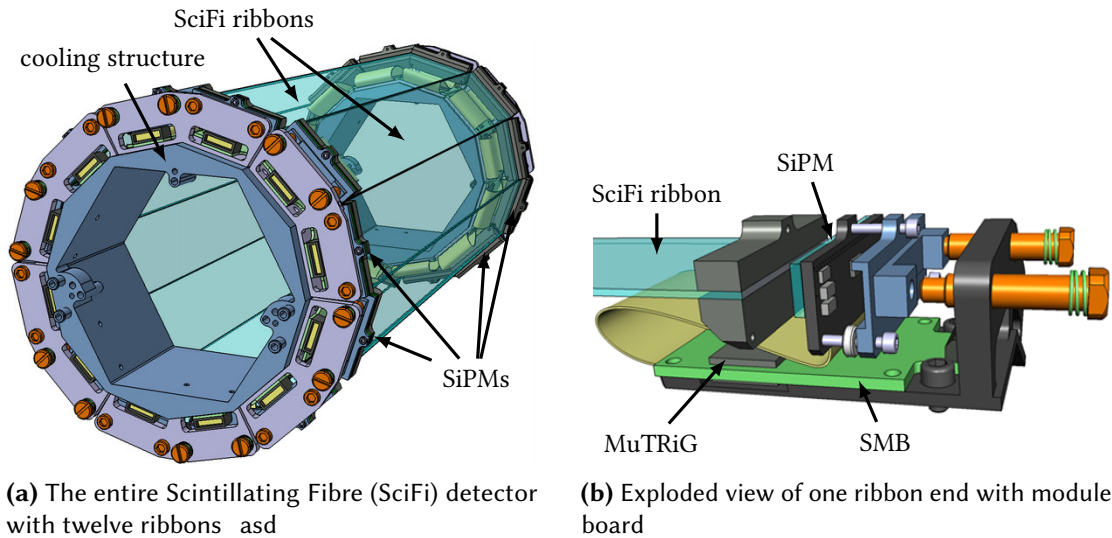


Figure 2.3: CAD renderings of the SciFi detector [1]. Each of the twelve ribbons in (a) is connected to two Silicon Photomultipliers (SiPMs) and SciFi Module Board (SMB) as shown in (b), one per side.

SCINTILLATING FIBRE DETECTOR

Reliable charge identification and background suppression require a time resolution the pixel sensors cannot offer. Therefore, a dedicated timing detector is located between the second and third pixel layer of the central station, at a radius of 61 mm: the Scintillating Fibre (SciFi) detector. The active part consists of scintillating fibres coupled to Silicon Photomultiplier (SiPM) arrays. With a low material budget of less than 0.2% of a radiation length, a time resolution better than 500 ps has been achieved, thus fulfilling the Mu3e specifications [40]. This thesis is focused on the development of the SciFi detector. The following presents an overview of the SciFi detector and its components. Details of the components are explained in chapter 3.

SCINTILLATING FIBRES The SciFi detector consists of fibre ribbons of three layers of round fibres with a diameter of $250\ \mu\text{m}$ (see also section 3.1). The fibre type chosen is the multi-clad, high-purity Kuraray SCSF-78 plastic scintillating fibre [41]. Each ribbon is 32.5 mm wide and 303 mm long and consists of three fibre layers held together with black epoxy. A slightly shorter ribbon prototype is shown in figure 2.4a.

SILICON PHOTOMULTIPLIERS One Hamamatsu S13552 [42] SiPM column array at each ribbon end detects scintillation light from the fibres (see section 3.2.3). The 128 channels have a pitch of $250\ \mu\text{m}$ and are not matched to the individual fibres. Each channel has an effective photosensitive area of $230\ \mu\text{m} \times 1625\ \mu\text{m}$. Figure 2.4c shows a sample of the final SiPM type used in the experiment.

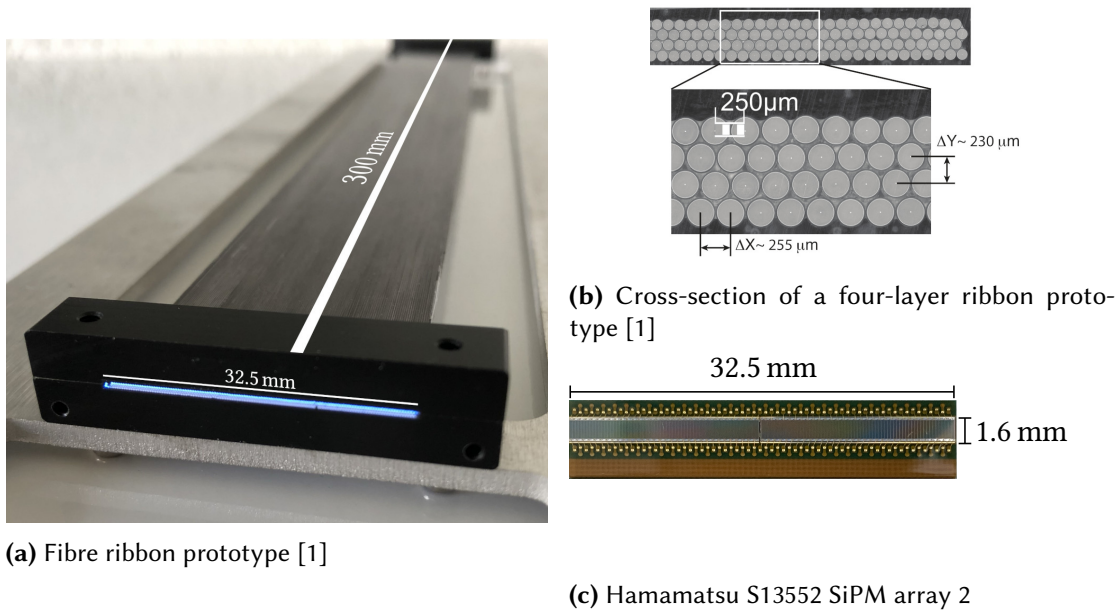


Figure 2.4: Prototypes of the main SciFi detector components


The SiPM arrays are bonded to the so-called SciFi Module Boards (SMBs) containing four readout chips (MuTRiG [43]) that digitize the analog signals. Those readout chips and the SMBs are described in detail in sections 6.4.1 and 6.4.2.

In the context of this thesis, the upstream and downstream SiPMs are indexed as shown in figure 2.5. That way, both SiPMs attached to the same fibre ribbon have the same index number.

MECHANICS AND COOLING The twelve fibre ribbons are arranged in a prism around the target and inner pixel layers (see also [1]). The ribbons are staggered by 12.5 mm along the z -axis and fixed to the SMB. An aluminium structure provides both cooling for the SMBs and mechanical support (see figure 2.3a). Each SMB produces approximately 5 W of heat dissipated via the support structure. Additionally, the aluminium structure cools the SiPMs that do not produce heat locally but are exposed to heat from the pixel detector. Cooling reduces the rate of spurious signals produced by the SiPMs.

SCINTILLATING TILE DETECTOR

The Scintillating Tile (SciTile) detector [44] provides precise timing information for particles reaching the upstream and downstream recurl layers (see figure 2.1). It is located between the recurl pixel layers and the beam pipe. A total of 5824 EJ-228 [45] plastic scintillator tiles with dimensions of 6.3 mm × 6.2 mm × 5.0 mm are coupled



(a) Upstream silicon photomultiplier indices, as seen from upstream (b) Downstream silicon photomultiplier indices, as seen from downstream

Figure 2.5: Numbering scheme in the scintillating fibre detector

individually to $3 \text{ mm} \times 3 \text{ mm}$ SiPMs [46]. The detector provides a time resolution of better than 100 ps and fulfils the requirements for background suppression.

Each tile is wrapped in reflective foil, increasing the light yield and preventing optical crosstalk between channels. The tiles are arranged in individual 4×4 matrices, which are read out in pairs by the same 32-channel MuTRiG Application-Specific Integrated Circuit (ASIC) used in the SciFi detector. The tile matrices are mounted on a cooling support structure that dissipates the heat produced by the readout ASICs. Using similar readout hardware allows for joint development of the hardware, firmware and software (see also chapter 6).

COMBINED TIME INFORMATION

In Mu3e, two types of tracks consist of more than four pixel hits. If the particle has a low momentum in the z -direction p_z , it will pass the four pixel layers in the central detector and the SciFi more than once as it recurls in the magnetic field. On its re-entry, it causes additional hits, which are added to the reconstructed track. If the momentum in $\pm z$ is large enough for a particle to reach one of the recurl stations (see figure 2.1), it will pass the SciFi detector only once. However, it will create a signal in the SciTile detector. The detector geometry makes both the scintillating tiles and fibres necessary for precise timing information. Previous simulations have shown that the background suppression power of either detector is insufficient, and only the combined time information allows suppression of Bhabha scattering by a factor of approximately 70 [40].

2.2 BEAMLINES AT PSI

PSI in Villigen, Switzerland, is host to the High-Intensity Proton Accelerator (HIPA) [47]. HIPA offers a proton beam with a high beam current of up to $I_p = 2.4$ mA. The protons are accelerated to a kinetic energy of 590 MeV to produce secondary particles by collision with fixed targets. Two graphite targets *Target M*, from French *mince* (“thin”, as it is 5 mm thin) and *Target E*, from French *épais* (“thick” for its thickness of 40 mm²) are used to produce charged pions, muons and electrons. Inside the graphite, the beam protons react with nucleons, producing positive pions π^+ and other final-state particles. Those pions can be grouped into three categories: pions whose energy is too small to escape the target material, pions that escape but decay in-flight close to the target, and pions that escape with energies high enough that they can be captured and steered in a beamline.

The last category is the source for *pion beams*. As charged pions predominantly decay into muons [12], pions that decay inside or near the target are potential sources for muons. The decay $\pi^+ \rightarrow \mu^+ \nu_\mu$ is a two-body decay. In the pion rest frame, all muons have the same momentum of $p_\mu = 29.8$ MeV/c. The most important class of particles for the scope of this thesis is called *surface muons* [48]. Muons stemming from pion decays near the surface of the target can escape the target with only very little energy loss. These nearly monochromatic muons with $p \lesssim 29.8$ MeV/c are highly polarized due to the chiral structure of the weak interaction (see section 1.4). Decays in the *pion cloud* around the target yield muons with a broader range of momenta than pion decays at rest. As the decay does not happen at rest, the emerging *cloud muons* are less polarized. A muon beam at the momentum of surface muons has a polarization of less than 100 % due to contamination with cloud muons. Since the proton beam consists of positively-charged particles, the rates of secondary particles with a positive charge are significantly higher than for their negative counterparts. Due to the nuclear capture of π^- in the proton target, virtually no negative surface muons are produced. Therefore, experiments utilizing surface muons such as Mu3e and MEG (II) reach their sensitivity only for μ^+ decays.

Users tune the currents in the beamline elements to select particles of specific momenta and adjust the beam position and focus. The following describes two beamlines of particular interest for this thesis in further detail.

2.2.1 PiM1

PiM1 is a beamline located at Target M. Users can utilize a high-intensity particle beam consisting mainly of electrons, pions and muons. Most testbeam campaigns of the Mu3e SciFi group, many of which I contributed to, were conducted here. The typical

²In the past, also a 60 mm Target E was used

momenta of the particles are around 200 MeV/c to 300 MeV/c, allowing beam intensities in the order of 10^8 s^{-1} . These particles are almost minimum ionizing ($\beta\gamma \approx 2-3$), which is a valuable property for detector characterization. Unfortunately, at these momenta, multiple Coulomb scattering is noticeable with detectors as thick as ours. The projected scattering angle distribution is a typical quantity to measure multiple Coulomb scattering [49]. The root mean square of that distribution for pions with a momentum of 200 MeV/c passing through a 500 μm thick plastic detector is 3.7 mrad. Our largest test setup covered a length of roughly 1.5 m along the beam direction, which makes corrections for scattering necessary.

2.2.2 PiE5

The PiE5 beamline provides the world's highest-intensity muon beam available for particle physics experiments. In the past, present and future, it serves the search for cLFV by different experiments. The MEG experiment was conducted at PiE5, as is the MEG II experiment as of 2022. In the future, Mu3e will run there as well.

Different beamline configurations were developed to accommodate the needs of Mu3e and MEG (II). The new Compact Muon Beam Line (CMBL) [50] is designed to fit the needs of Mu3e and delivers the necessary muon rates. The muon rate at the injection to the Mu3e solenoid is $1.1 \cdot 10^8 \text{ s}^{-1}$, normalized to a proton current of 2.4 mA and the 40 mm baseline Target E [51, 52]. Recent studies for different target geometries have been performed at PSI [32, 51, 52]. A slanted Target E increases the rate of surface muons in PiE5 by 42 % [32] compared to the baseline design.

The following is a brief overview of this beamline's elements, shown in figure 2.6. The nomenclature of electromagnets encodes the order of the magnetic multipole: A for dipole magnets (German: Ablenkmagnet), Q for Quadrupole magnets and H for sextupole/Hexapole magnets. A quadrupole field focuses a particle beam in one dimension and defocuses it in the other. Therefore, quadrupole magnets are placed in sets of two or three to combine focusing properties in both dimensions. A steering magnet (SML) allows for fine-tuning.

A user can set currents in electromagnets, adjustable slit systems to tune the beam rate and momentum by, and the high voltage of the separator. The magnet currents and high voltage are regulated with feedback loops that are considered very stable. The magnets are cycled to the maximum current in either direction before setting the desired current. Cycling the magnets prevents inconsistencies due to magnetic hysteresis. By construction, the slits do not move once they are set.

The trajectory of a charged particle with momentum p is deflected by the angle

$$\theta = 0.2997 \cdot \frac{Bl}{Tm} \frac{\text{GeV}}{pc} \quad (2.2)$$

when it travels a distance l in a homogeneous field with a flux density B . The beamline

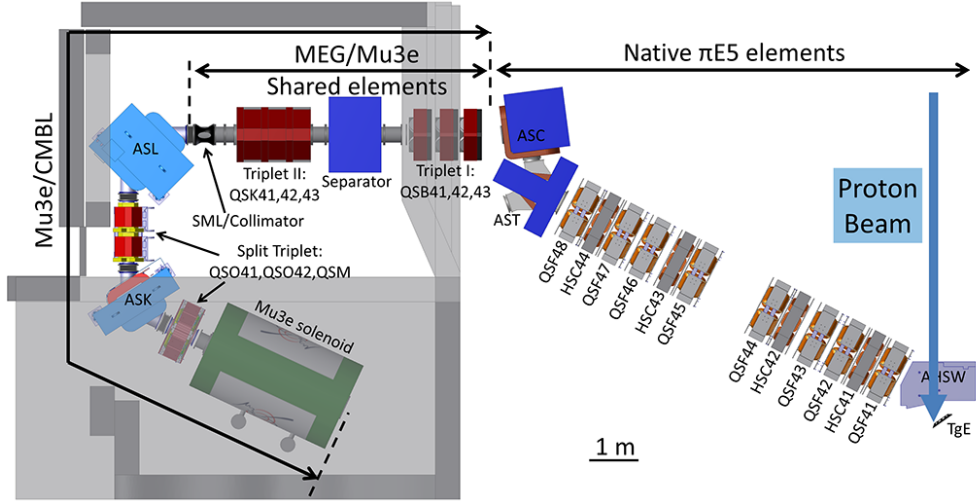


Figure 2.6: Rendering of the PiE5 secondary beamline in the Compact Muon Beam Line (CMBL) configuration used for the Mu3e experiment [1]

transports positrons, muons and pions with the same momentum. A Wien filter (SEP41) is installed to separate the particle types [50]. By generating an electric field E and a magnetic field B perpendicular to each other, particles are filtered by their velocity. Only particles with a velocity

$$v = \frac{E}{B} \quad (2.3)$$

pass the separator undeflected. With the geometry in PiE5, particles that are too fast (i.e. positrons) are deflected upwards.

Charged pions with a momentum of 28 MeV/c have an average decay length of

$$v\gamma\tau \approx 1.6 \text{ m} \quad (2.4)$$

with the velocity v , the Lorentz factor γ and the decay time $\tau \approx 2.6 \cdot 10^{-8}$ s [12]. The beam contains virtually no pions since the distance between Target E and the experimental area is more than 11 m, approximately seven decay lengths.

In 2020, this beamline hosted the irradiation campaign described in chapter 5 of this thesis. The irradiation setup was mounted downstream of SML. In 2021, the Mu3e collaboration had its first joint beamtime at PiE5 in the CMBL configuration. The efforts leading to this beamtime are described in chapter 6.

3

Scintillating fibres and silicon photodetectors

3.1 SCINTILLATING FIBRES

This section explains the fundamentals of scintillation in general and scintillating fibres in particular, which are a core component of the Mu3e SciFi detector. A focus lies on light-loss mechanisms in scintillating fibres, laying the groundwork for chapter 4.

3.1.1 SCINTILLATION PROCESS

Scintillation is a phenomenon that causes a substance to emit light as a reaction to energy deposition by ionizing radiation. The light emission results from the relaxation of electrons from excited states. Therefore, the probability of said relaxation dictates the emission over time which follows an exponential decrease. This pattern is characterized by the decay time τ , after which the light emission is reduced to $\frac{1}{e} \approx 37\%$. Scintillation can occur in inorganic materials (e.g. NaI crystals) and organic materials.

Insulators and semiconductors have discrete energy bands, the so-called valence band, in which electrons are bound to individual atoms, and the conduction band, where electrons can move freely in the material. Scintillation in *inorganic scintillators* is caused by the relaxation of electrons from the conduction band to the valence band. The process of excited electrons falling from the conduction band back to the valence band has a low probability and is, therefore, slow. Photons emitted this way have energies in the Ultraviolet (UV) range, making this process unsuitable for particle detectors. Adding impurities to the scintillator, so-called *activators*, generates additional energy states in the band structure. The probability of relaxation from an activator excited state to an activator ground state is much higher, thus reducing the decay time to 30 to 500 ns [53]. In cases where the gap between the activator states is lower than in

the intrinsic scintillator, the emitted photons are in the optical spectrum, simplifying photodetection.

Organic scintillators are based on organic compounds. *Pure organic scintillators*, such as anthracene [54], consist of a crystal that intrinsically shows scintillation. *Plastic scintillators* consist of fluorescent molecules, so-called *fluors*, embedded in a *base*, a solid polymer, e.g. polystyrene (PS). When ionizing particles deposit kinetic energy in a plastic scintillator, a significant fraction of the energy is dissipated nonradiatively, either as heat or lattice vibrations. Via the Förster mechanism [49], energy is transferred to the fluor, where it induces photon emission through fluorescence. Since the absorption and emission spectra of the (typically UV-emitting) fluors overlap, reabsorption limits light propagation. To combat this, additional *wavelength shifters*, fluorescent molecules that emit light at larger wavelengths, are added to the plastic scintillator.

The typical decay time of plastic scintillators is several orders of magnitude shorter than that of inorganic scintillators [49]. A novel type of plastic scintillator consists of single Nanostructured Organosilicon Luminophores (NOLs) that combine initial activation and wavelength-shifting in a single molecule [55]. NOL-based plastic scintillators offer shorter decay times than conventional plastic scintillators, where fluor and wavelength-shifter are separate molecules. Research and development are in an early state, and the viability for particle detectors is subject to current research activities.

3.1.2 SCINTILLATING FIBRES

Plastic scintillators are produced in various shapes, notably plastic optical fibres, most of which have either a square or round cross-section. Scintillating Fibres (SciFis) combine light emission and guiding of the scintillation light.

Scintillating fibres have been used historically [56] and in modern particle detectors, such as in the LHCb tracker upgrade [57] and Mu3e [1]. Scintillating fibres allow for better position resolution of the detected particles than scintillator blocks. NOL-based scintillating fibres have been produced [58] and are being evaluated for future upgrades to fibre detectors, although they are not yet commercially available in sizeable quantity.

Fibres containing no initial fluor, but only wavelength-shifting agents, are used to guide light from scintillator blocks for detection over long distances. A larger wavelength minimizes light loss as shorter wavelengths suffer from stronger attenuation (see BULK SCATTERING in section 3.1.3).


Light propagates in scintillating fibres by total internal reflection. The maximum angle of total internal reflection (critical angle) θ_c at the interface of two materials depends on their refractive indices n_1 and n_2 . Total internal reflection is only possible in the medium with the larger refractive index. If $n_1 > n_2$, then the largest angle under which total internal reflection occurs is

$$\sin \theta_c = \frac{n_1}{n_2}, \quad (3.1)$$

also called the critical angle (see figure 3.1a). The larger θ_c in a scintillating fibre, the more light is collected.

The base material is optionally clad with one or more layers of other plastics, protecting the material from mechanical damage and dust, and increasing the light capture efficiency. The layers have to be added in decreasing order of refractive index to allow total internal reflection, a common material being acrylic (poly(methyl methacrylate) (PMMA)) with a refractive index of $n = 1.49$. Considering total internal reflection only on the inner core-cladding and cladding-cladding interfaces yields a lower limit for the capture efficiency, as reflection is also possible at the outermost interface if the fibre is in air. A second cladding improves the light yield of fibres by up to 50 % [41] (see figure 3.1).

Round fibres have a trapping efficiency that depends on the distance of the light origin to the central axis. Saint-Gobain Crystals states the minimum value for their single-clad fibres as 3.4 % at the central axis and approximately 7 % at the interface between core and cladding [59]. The trapping efficiency of square fibres is independent of the position of the light event. For square single-clad fibres, Saint-Gobain Crystals states a uniform trapping efficiency of 4 % [60].



(a) Light propagation in a single-clad PS fibre ($n = 1.59$). The cladding consists of PMMA ($n = 1.49$) (b) Light propagation in a double-clad fibre. In addition to (a), the second cladding consists of a fluoropolymer ($n = 1.42$)

Figure 3.1: Comparison of the capture efficiency between single- and double-clad scintillating fibres. Only total internal reflection on the inner interfaces are considered. The second cladding increases the critical angle θ_c and thus the capture efficiency. Adapted from [41].

3.1.3 LIGHT LOSS IN OPTICAL FIBRES

Several mechanisms lead to a loss of light in optical fibres. In the following, the most important mechanisms are summarized.

ABSORPTION

Although it looks transparent to the human eye, the fibre bulk material absorbs light with a certain probability. The attenuation spectrum of PS, a commonly used base for plastic scintillating fibres, is shown in figure 3.2.




Figure 3.2: Attenuation spectrum of PS optical fibres. Figure adapted from figure 2.176 in [61]. The two curves show two independent measurements, both of which show a peak in attenuation for wavelengths around 610 nm.

In addition to the pure bulk material, photons can be absorbed by impurities in the material. Such impurities can be either in the bulk itself or at the surface between the bulk and cladding. Macroscopic impurities (e.g. a speck of dust) lead to the loss of a significant percentage of light which causes a sudden drop in intensity along the fibre. Homogeneously distributed impurities lead to an exponential decrease of light along the path.

BULK SCATTERING

Similar to absorption processes, scattering targets can be situated in the bulk material or at the interface of the fibre surface. If the new direction exceeds θ_c , the condition of total internal reflection is no longer fulfilled: The probability of light scattering on microscopic impurities and density variations depends on the wavelength of the light. Inside the fibre core, Rayleigh scattering is the dominant process, while Mie scattering occurs at the core-cladding interface [62]. If the wavelength is comparable to the size of a spherical scattering target, the scattering is described by Mie scattering [63]. For scattering targets much smaller than the wavelength, that process is described using the Rayleigh approximation [64]. While for Mie scattering, all wavelengths λ are scattered approximately equally, the scattering probability under the Rayleigh approximation is proportional to $\frac{1}{\lambda^4}$. Therefore, shorter wavelengths experience stronger scattering and attenuation. A similar phenomenon is known as the Tyndall effect.

CLADDING CONTAMINATION

Impurities on the cladding can decrease the critical angle at the cladding-air interface, either by absorption or by modifying θ_c . The light that is transmitted through this undesired cladding-contaminant interface is then lost.

SIMPLE ESCAPE

Light rays whose incident angle at the cladding exceeds the critical angle are only *partially* reflected. As the diameter of scintillating fibres is much smaller than the length, photons undergo several reflections. Only light rays with $\theta > \theta_c$ survive long distances along the fibre. For short paths, the number of reflections is small enough for a fraction of the light to reach the fibre end, even by partial reflection. The partially reflected light leads to a non-exponential decrease for short lengths.

Uneven or mechanically damaged cladding will lead to locally different incidence angles, leading to light escape in the damaged area.

BENDING LOSSES

For ideal, straight fibres, the angle of incidence at the cladding is the same for every reflection. Bending the fibre, however, can decrease the angle of incidence to a value less than the angle of total internal reflection, thus allowing light to escape the fibre. Smaller bending radii lead to more light loss.

SUMMARY

Light loss in optical fibres has many reasons, some intrinsic to the fibre, some extrinsic, caused by external influences, such as contaminants or damage. Rayleigh scattering is more pronounced at shorter wavelengths, and the transparency of the fibre is wavelength-dependent. Therefore, light attenuation generally is wavelength-dependent and different attenuation lengths can be associated with different wavelengths.

3.1.4 QUANTIFYING LIGHT LOSS

The radiant flux Φ_e denotes the radiant energy Q_e per unit of time:

$$\Phi_e = \frac{\partial Q_e}{\partial t} \quad (3.2)$$

The radiant flux per unit of wavelength λ is called the spectral flux $\Phi_{e\lambda}$ and is defined as

$$\Phi_{e\lambda} = \frac{\partial \Phi_e}{\partial \lambda}, \quad (3.3)$$

allowing the quantification of spectral attenuation phenomena.

As light is guided along a fibre, Φ_e decreases with the covered fibre length z . If the probability of light loss is constant along the fibre, the transmitted radiant power is described by the differential equation

$$\frac{d\Phi_e}{dz} = -\frac{1}{\Lambda} \Phi_e(z) \quad (3.4)$$

with the *attenuation length* Λ . The solution of this differential equation is an exponential function:

$$\Phi_e(z) = \Phi_e(0) \exp^{-z/\Lambda} \quad (3.5)$$

The same applies to spectral flux, except the attenuation length is a function of the wavelength λ .

$$\Phi_{e,\lambda}(z) = \Phi_{e,\lambda}(0) \exp^{-z/\Lambda(\lambda)} \quad (3.6)$$

In industry, light loss in fibre optics is often quoted in decibels (dB) to express the deterioration of optical signals with an incident radiant flux $\Phi_e(0) = \Phi_e^{\text{in}}$ and an outgoing radiant flux $\Phi_e(d) = \Phi_e^{\text{out}}$ after a defined distance d . Therefore, one defines the attenuation A as

$$A = 10 \log_{10} \left(\frac{\Phi_e^{\text{in}}}{\Phi_e^{\text{out}}} \right). \quad (3.7)$$

Cumulative losses from coupled devices can be calculated by adding the attenuation in each device.

The conversion from attenuation in decibels after one kilometre to *attenuation length* is straightforward as long as equation 3.5 is fulfilled.

$$A \text{ [in dB/km]} = 10 \log_{10} \left(\frac{\Phi_e(0)}{\Phi_e(1 \text{ km})} \right) \quad (3.8)$$

$$= 10 \log_{10} \left(\frac{1}{\exp(-1 \text{ km}/\Lambda)} \right) \quad (3.9)$$


$$= \frac{10}{\ln(10)} \cdot \frac{1 \text{ km}}{\Lambda} \quad (3.10)$$

$$\Rightarrow \Lambda = \frac{10}{\ln(10)} \cdot \frac{1 \text{ km}}{A} \quad (3.11)$$


However, if the attenuation is more complex, equation 3.5 needs to be modified. If two groups of photons experience attenuation from two different mechanisms, one with a short attenuation length Λ_{short} and the other with a long attenuation length Λ_{long} , equation 3.5 becomes

$$\Phi_e(z) = \Phi_e^{\text{short}} \exp^{-z/\Lambda_{\text{short}}} + \Phi_e^{\text{long}} \exp^{-z/\Lambda_{\text{long}}}, \quad (3.12)$$

where Φ_e^{short} and Φ_e^{long} describe the radiant flux from both groups at $z = 0$.



(a) Meridional rays cross the central region of the fibre during reflection



(b) Helical rays are propagated under many reflections and stay close to the cladding

Figure 3.3: Two different ways of propagation light under total internal reflection. The left part shows a cross-section of a round fibre, while the right part shows a lateral view. Helical light rays have greater path lengths and undergo more reflections than meridional rays.

Such multiple-exponential behaviour can arise from different propagation paths inside a fibre. The first one concerns the path of light rays inside the fibre (see figure 3.3). Between reflections, a *meridional* ray crosses the centre of the fibre during propagation. A *helical* or *skew* ray is reflected such that the path is close to that of a helix. For the same distance along the fibre, a skew ray has a greater path length and undergoes more reflections. Therefore, helical rays have a shorter attenuation length than meridional rays.

In clad fibres, only reflections at internal interfaces usually are considered (see figure 3.1). However, if the optical density of the environment is lower than that of the outer cladding, total reflection is also possible at this outer interface. As listed above, such rays are exposed to loss mechanisms the internal ones are not. As the probability of light loss of such *cladding rays* per reflection is higher, such rays have a smaller attenuation length than rays reflected at internal interfaces.

If the attenuation length is wavelength-dependent, then different parts of the wavelength spectrum die off at different rates. This also causes a non-exponential decrease in light intensity.

3.2 SILICON PHOTODETECTORS

From the first light-sensitive silicon devices in the 1940s to single-photon detection in high-speed applications, silicon photodetectors have seen huge efforts in research and development. This section explains the working principle of silicon photodetectors, starting from photodiodes and leading to Silicon Photomultipliers (SiPMs).

3.2.1 PHOTODIODES

Photodiodes are, as the name suggests, diodes designed for light detection. The most simple silicon diode consists of a junction of a p- and an n-doped region of silicon. The p-doped region is called the anode, and the n-doped region is the cathode. Electrons diffuse from the anode to the cathode, and holes diffuse in the opposite direction. At the equilibrium of this diffusion process, an electric field builds up at the junction. That electric field creates a built-in potential difference V_{bi} .

Applying an external voltage from the anode to the cathode is called *forward bias*. If the external voltage exceeds V_{bi} , current flows from the anode to the cathode—the diode becomes conductive. A *reverse bias* amplifies the internal electric field, and almost no current flows from the cathode to the anode. This happens only if the bias voltage does not exceed a specific limit called breakdown voltage V_{BD} . If the bias voltage exceeds V_{BD} , the reverse current increases exponentially.

If a photon creates an electron-hole pair in a reverse-biased diode, those charge carriers are accelerated to the anode and cathode, and a small current is detectable. Therefore, a reverse-biased diode is sensitive to light and the current increases with illumination. The sensitivity can be improved by a layer of intrinsic silicon between the two doped regions, then named PIN diode. An exemplary use case is the light measurement described in chapter 4, where such a diode was used extensively.

3.2.2 AVALANCHE PHOTODIODES

Charge carriers are accelerated in electric fields. In a strong enough field, an electron gains enough kinetic energy to create new electron-hole pairs through impact ionization. The new free electron is also accelerated, creating more electron-hole pairs and an electron avalanche forms. An electron avalanche produces a current gain. This effect is induced in Avalanche Photodiodes (APDs) by a large doping gradient, which creates a strong electric field. The avalanche region created this way can either be near the cathode (called reach-through configuration) [53] or near the anode (reverse reach-through configuration), as shown in figure 3.4.

If the bias voltage exceeds V_{BD} , a single electron-hole pair is enough to bring the photodiode to a conducting state, known as Geiger discharge. The current in the conducting state is self-sustaining. An APD operated in this Geiger mode is also called Geiger-mode Avalanche Photodiode (GAPD). A quenching resistor R_q of several hundred kilohms is added in series to the diode to stop a GAPD from conducting. When a current starts flowing, the voltage drop across R_q decreases the bias voltage at the diode. Once the voltage is below V_{BD} , the current is no longer self-sustaining and stops. Without current, there is no voltage drop across R_q anymore, and the diode is back in Geiger mode, ready for the next photon. The time it takes for the avalanche to be quenched and the voltage across the diode to be restored is referred to as recovery

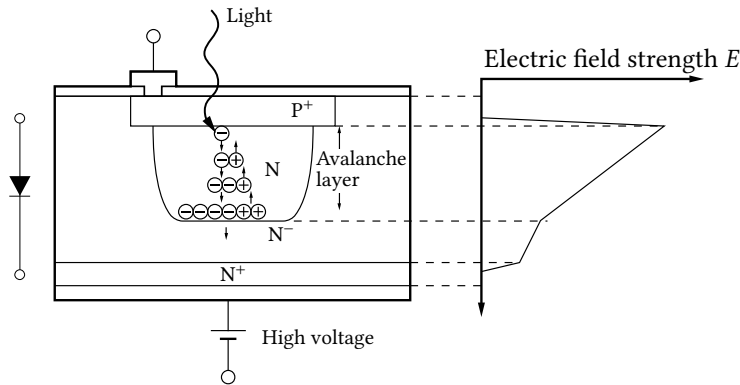


Figure 3.4: Layout of a silicon APD in reverse reach-through configuration. Adapted from [65]

time. The recovery time depends on R_q and the capacitance of the GAPD.

A single photon can create detectable currents in a GAPD, but any number of concurrent photons creates the same output signal. That is why GAPDs are also referred to as Single-Photon Avalanche Diode (SPAD).

3.2.3 SILICON PHOTOMULTIPLIERS

The lack of distinction between one and many photons by GAPDs can be compensated by a parallel circuit of multiple GAPD pixels. Each pixel creates the same signal when a photon creates an electron-hole pair. The pitch of such a pixel array is typically tens of micrometres. If two or more pixels are triggered simultaneously, the output signal is larger by the corresponding number of hit pixels. Arrays of GAPD pixels are referred to as Multi-Pixel Photon Counters (MPPCs) or Silicon Photomultipliers (SiPMs) [65].

In contrast to a Photomultiplier Tube (PMT), a SiPM works in strong magnetic fields and is significantly smaller (millimetres compared to centimetres). PMTs require a bias voltage in the order of kilovolts, while the typical bias voltage is below 100 V. The ability to detect and resolve single photons makes SiPMs a viable alternative to PMTs.

Similar to PMTs, SiPMs are prone to afterpulses. During an electron avalanche, charge carriers can get trapped in energy levels below the conduction band due to impurities. The time until such a charge carrier is released lies in the order of nanoseconds. If the release happens during recovery time, a second, smaller avalanche is triggered. After recovery, the triggered avalanche is indistinguishable from a second, independent avalanche.

SiPMs have a typical rise time in the order of 1 ns allowing their use in applications requiring high-precision time measurements. Among the applications that use SiPMs are laser ranging with time-of-flight measurements and particle detectors, for example, in the Mu3e SciFi detector.

During an avalanche, charges are accelerated and decelerated during impact ioniz-

ation, which can generate Bremsstrahlung. If a Bremsstrahlung photon crosses the bounds to a neighbouring pixel, it can generate electron-hole pairs. If the electron-hole pair is generated in the amplification region of the neighbouring pixel, it triggers a simultaneous avalanche. This is known as *prompt crosstalk*. The signal of prompt crosstalk is indistinguishable from a two-photon event. Such a photon can also interact with a neighbouring cell outside the amplification region and generate an electron-hole pair there. When the electron reaches the amplification region by diffusion, it triggers an avalanche with a time delay. This is referred to as *delayed crosstalk*. Optical trenches between the pixels significantly decrease the probability of both types of optical crosstalk.

DARK PULSES AND DARK COUNT RATE

As for the GAPD, every electron-hole pair in the avalanche region triggers an electron avalanche. When an electron in the avalanche region is thermally excited to the conduction band, it triggers a spurious signal indistinguishable from a real single-photon signal. The mean rate at which such signals occur in the dark is called Dark Count Rate (DCR) ν . As it is a thermal effect, it is strongly dependent on temperature. An increase of 10°C typically doubles the DCR [65]. At the same time, the DCR increases with a higher bias voltage.

In settings where single-photon signals are sought, a high DCR generates large amounts of spurious data, potentially shadowing actual data during the recovery time. A high DCR makes it impossible to resolve individual photoelectrons if the output of a SiPM is so noisy that the individual signals overlap in time. Applications that rely on the time resolution of SiPMs are affected by DCR as a source of electric jitter, limiting the time resolution.




Figure 3.5: Hamamatsu S13552 SiPM column array [42]. Each of the 128 channels has an effective photosensitive area of $230\ \mu\text{m} \times 1625\ \mu\text{m}$ and consists of 4×26 individual Geiger-mode APDs (see blow-up region). One APD has an area of $57.5\ \mu\text{m} \times 62.5\ \mu\text{m}$.

3.3 RADIATION DAMAGE IN SILICON PHOTOMULTIPLIERS

In Mu3e, most radiation comes from positrons from muon decays at rest, thus with energies up to 53 MeV. Positrons can damage silicon detectors by several mechanisms that can be grouped into modifications to the *surface* and the *bulk* of the silicon. This section contains explanations and theoretical backgrounds that help understand the radiation damage in silicon photomultipliers that is expected in the Mu3e experiment.

3.3.1 SURFACE DAMAGE

The first group of defects, *surface damage*, is caused by ionizing energy loss quantified by the Total Ionizing Dose (TID). Ionization radiation builds up space charge and interface traps at the silicon oxide and the Si-SiO₂ interface [3]. Space charge modifies the electrical field inside a SiPM, while interface traps generate a surface current I_{surf} under external electric fields. If a part of I_{surf} reaches the amplification region of a SiPM operated above the breakdown voltage, the charge carriers induce avalanches and, thus, an increase in DCR. Below the breakdown voltage, I_{surf} is not amplified.

3.3.2 BULK DAMAGE

Bulk damage is inflicted by high-energy particles depositing energy inside the bulk material of silicon through non-ionizing energy loss (NIEL). If the energy transfer to a bulk atom is high enough, the atom is displaced to an interstitial state. This interstitial-vacancy pair is called a *Frenkel pair* [3]. If that Primary Knock-on Atom (PKA) has sufficient energy, it can displace other atoms, leading to cluster defects. These types of defects add new energy levels to the band structure. Defects with energies in the middle of the band gap lead to the generation of free charge carriers [3]. Those charge carriers are amplified in a SiPM, causing an increase in DCR. Other types of bulk defects include charge traps that decrease the signal amplitude. Non-uniform signals of the different pixels in a SiPM affect the single-photon resolution.

Furthermore, radiation can change the doping of silicon. Modified doping will change the electric field inside the SiPM, affecting the breakdown voltage of that particular pixel. Again, the single-photon resolution is worsened. Bulk damage is quantified by NIEL or by particle fluence (see section 3.4).

DISPLACEMENT THRESHOLD ENERGY

The energy needed to create a Frenkel pair is several times higher than the energy for adiabatic displacements [66]. The experimentally determined *displacement threshold energy* E_d accounts for the minimum energy needed to displace atoms but also annealing effects, such as thermal recombination of Frenkel pairs. Literature (e.g. [7] and the

sources therein) states values of E_d between 21 and 25 eV after annealing. For low-energy electrons (< 2 MeV), the exact value of E_d is a significant source of uncertainty.

The maximum kinetic energy an electron can transfer to a silicon nucleus at rest can be calculated using the kinematics of a two-body scattering process as

$$T_{\max}^e \simeq 19.96 \tau(\tau + 2) [\text{eV}], \quad (3.13)$$

where $\tau = E_k^e/m_e$ is the ratio of the electron's kinetic energy and its mass [7]. Using equation 3.13 and $E_d = 21$ eV yields a minimum kinetic energy of 220 keV that allows electrons to cause displacement damage.

The threshold for cluster defects is significantly higher, and cluster defects become relevant at $E_k^e \gtrsim 5$ MeV [7].

3.4 QUANTIFICATION OF RADIATION DAMAGE

The *amount* of radiation is quantified with different observables depending on the type of damage that should be described. The Total Ionizing Dose (TID), i.e. the amount of ionizing energy absorbed per unit mass m , is used for surface damage in electronics and medicine. TID is measured in Gray: $1 \text{ Gy} = 1 \text{ J/kg}$.

For radiation with different particles, *particle flux* is used. The particle flux Φ is the rate of incident particles per unit area. Integrating the particle flux over time yields the *fluence* Φ . This unit area is to be understood perpendicular to the direction of the particle. The fluence in a radiation environment is independent of the orientation of the area. For surfaces that are exposed to particles from various directions, the angle of incidence θ must be considered. As shown in figure 3.6, that angle of incidence leads to a longer path length in a given volume. When counting particles crossing the surface of that volume, particles have to be weighted with a factor $\frac{1}{\cos(\theta)}$ to account for that greater path length.




Figure 3.6: Particles crossing a thin surface with thickness δ . The path length of a perpendicular particle (left) is δ , but under an angle of incidence θ the path length is $\delta / \cos(\theta)$

The dissipated non-ionizing energy per unit length $\frac{dE}{dx}$ is proportional to the displacement damage cross-section D , typically stated in units of MeV mb. The American

Society for Testing and Materials defines the displacement damage cross-section of 1 MeV neutrons as 95 MeV mb [67]. That way, radiation can be normalized to an equivalent fluence of 1 MeV neutrons Φ_{eq} . The displacement damage cross-section for electrons is shown in figure 3.7. Electrons create slightly more displacement damage than positrons. The ratio is close to unity in the energy range expected in Mu3e and for silicon [68].

$$1 - \frac{D_{e^+}}{D_{e^-}} \lesssim 10 \% \quad (3.14)$$

For the sake of simplicity, D_{e^+} and D_{e^-} are treated as equal in chapter 5.

3.4.1 NIEL HYPOTHESIS

Numerous observations have led to the hypothesis that radiation damage in electronics scales proportionally to the energy deposited by non-ionizing energy loss (NIEL) [7, 69–71]. This is called NIEL scaling hypothesis. Comparing the amount of NIEL deposited by the fluence of any particle to the NIEL from 1 MeV neutrons, one can normalize radiation to the standard. Instead of the actual fluence, an equivalent fluence of 1 MeV neutrons is stated.

LIMITS OF THE NIEL SCALING HYPOTHESIS

Irradiation with different particles and energies leads to a different distribution of microscopic defects. Therefore, the damage effects do not necessarily scale just with NIEL. The NIEL scaling hypothesis has been shown to break down when comparing radiation damage from different particle types and energies. The transferred energy affects the distance between displaced atoms and vacancies and, thus, the probability of recombination [72]. Microscopic simulations show a violation of NIEL scaling [73]. Particularly for electrons, the NIEL hypothesis does not hold up [74]. Better models, such as *effective NIEL*, have been developed to improve the description of radiation damage due to electrons [5, 6]. This model uses molecular dynamics simulations, improving the description for low energy transfers to the PKA.

3.4.2 OTHER FACTORS

Radiation damage depends on the total dose and other factors, such as dose rate, temperature, and the composition of the atmosphere. Special attention has to be given to seemingly unintuitive relations. For example, TID induces larger degradation of bipolar transistors at lower dose rates [7].

Defects can react with impurities and doping agents and form compound defects. Examples of compounds are vacancy-oxygen (*V-O*) and vacancy-phosphorus (*V-P*). Since impurity concentrations and doping depend on the manufacturing process of the

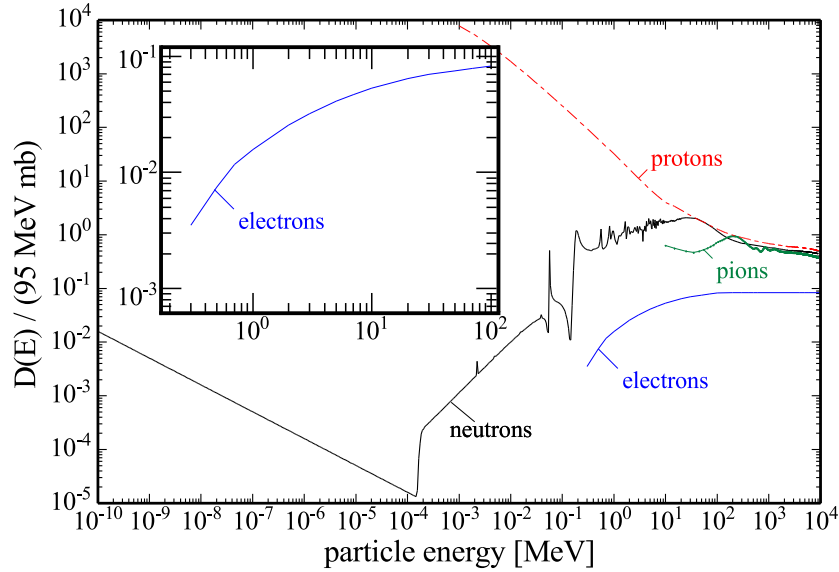


Figure 3.7: Displacement damage cross-section of neutrons, protons, pions and electrons in silicon as a function of the kinetic energy E . The function is normalized to 95 MeV mb, which corresponds to a neutron with $E = 1$ MeV [67]. An electron with $E = 15$ MeV, for example, causes a factor of 20 less damage than a 1 MeV neutron. The inset shows a zoomed-in part of the electron curve, covering the range of energies in Mu3e (up to 53 MeV). Adapted from [75], originally summarized in [76], based on data from [77] (also published in [78]), [79–82].

silicon, radiation effects can differ between seemingly identical detectors from different silicon foundries.

It is therefore strongly advised to conduct radiation hardness studies with the same devices of interest in an environment that is as similar as possible to the final conditions.

3.5 IMPACT ON MU3E

The Mu3e SciFi detector (see chapter 2.1.4) contains both plastic scintillating fibres and SiPMs. As shown in chapter 4, the SciFi detector with its length of 30 cm is affected strongly by attenuation that cannot be described with a single exponential decrease. With the goal of stopping $2.5 \cdot 10^{15}$ muons in Mu3e Phase I and the short distance between the stopping target and the detectors, a large positron fluence is expected. Chapter 5 also contains simulation results of the expected fluence in the SciFi SiPMs. Due to the limitations mentioned above of the NIEL scaling hypothesis, irradiation with neutrons is not considered suitable to predict the performance of positron-irradiated SiPMs. Therefore, SiPMs were irradiated with Michel positrons, and the effects of radiation damage are presented in the same chapter.

4

Attenuation studies of scintillating fibres

4.1 MOTIVATION

The amount of light detected in the Mu3e scintillating fibre detector limits the time resolution and efficiency. That amount depends on how much light is produced, transmitted to the fibre ends, and then detected in the SiPMs. The focus of this chapter lies on light attenuation inside the fibres. Light attenuation is well-described by an exponential decrease as a function of the distance for long fibres. Earlier measurements (e.g. [2]) have shown that this model is insufficient to describe the light yield at distances under one metre. In addition to the one attenuation length at long distances, the light attenuation follows a double-exponential decay with a short and a long attenuation length with possibly even more components due to different light loss mechanisms (see section 3.1.3).

Based on the thorough evaluation by the LHCb collaboration [57], Mu3e made an early tentative decision to use the same type of scintillating fibres: high-purity Kuraray SCSF-78 with a round cross-section and a diameter of $250\ \mu\text{m}$. In the LHCb SciFi Tracker, the fibres are cut to a length of 2.5 m. In contrast, the Mu3e SciFi detector has a length of just 30 cm, giving rise to different attenuation over such short distances.

The manufacturer determines the attenuation length stated in the datasheet [41] by exciting the fibres between one and three metres from a photodetector. Therefore, their value of $> 4.0\ \text{m}$ does not necessarily reflect the performance in Mu3e. An additional shorter component of the attenuation length would not be problematic per se, as it corresponds to additional light that can be detected. However, to evaluate fibres for Mu3e, the attenuation over short distances needs to be understood. This chapter presents the evaluation of the attenuation lengths of Kuraray SCSF-78 and six other types of scintillating and wavelength-shifting fibres.

4.2 SETUP

Shining UV light onto the side of a fibre stimulates photon emission. This technique is known as side-induced fluorescence. The fibres are placed in a setup as shown in figures 4.1 and 4.2. A movable block containing a UV-emitting Light Emitting Diode (LED) stimulates the fibre in different positions. One end of the fibre is either coupled to a calibrated PIN photodiode¹ attached to a power meter² or a spectrometer³. The attenuation behaviour of scintillating and wavelength-shifting fibres are studied with measurements of the radiant flux Φ_e and the spectral flux $\Phi_{e\lambda} = \frac{d\Phi_e}{d\lambda}$ of light emitted the fibre end.

The calibrated photodiode and the power meter allow high-precision radiant flux measurements. Studying the spectral flux with the spectrometer provides information on wavelength-dependent attenuation.

4.2.1 FIBRES UNDER STUDY

The following fibres are studied. All fibres under study have a diameter of 250 μm and two cladding layers.

- Kuraray SCSF-78, round fibre, *high purity**
- Kuraray SCSF-81, round fibre, *high purity**
- Kuraray NOL-11 round fibre*
- Saint-Gobain Crystals BCF-12 round* and square[†] scintillating fibres
- Saint-Gobain Crystals BCF-20 square scintillating fibre[†]
- Saint-Gobain Crystals BCF-92 square blue-to-green wavelength shifting fibre[†]

While the exact composition is typically a trade secret, Saint-Gobain Crystals states, “The core contains a combination of fluorescent dopants selected to produce the desired scintillation, optical and radiation-resistance characteristics” [85]. Similarly, Kuraray states, “The scintillating core contains a combination of fluorescent dopants” [41]. The NOL-11 fibres are the only fibre type under study that does not use a combination of fluorescent dopants but NOL molecules.

The square fibres by Saint-Gobain Crystals were cut from spools at PSI, where the measurements took place. The round fibres were cut in Geneva and transported to Villigen.

¹Newport 818-UV [83]

²Newport 843-R [83]

³FLAME-S-XR1-ES with a 1000 μm solarization-resistant fibre adapter (UV/VIS), both from Ocean Insight, USA [84].

*Provided by University of Geneva

[†]Provided by Paul Scherrer Institute




Figure 4.1: Fibre characterization setup. The fibre is excited with a UV LED in different positions. The end of the fibre is coupled to either a PIN photodiode (a) or a spectrometer (b)




Figure 4.2: Photograph of the characterization setup

4.3 RADIANT FLUX MEASUREMENTS

The first data set is recorded using the PIN photodiode and power meter mentioned above. The radiant flux Φ_e emitted from the fibre is measured via the current flowing through the diode and converted with built-in calibration.

4.3.1 MEASUREMENT UNCERTAINTY

Multiple sources contribute to the uncertainty of the radiant flux discussed in the following section (e.g. figure 4.3). No independent detector of comparable accuracy was available for intercalibration at the time of measurement. Therefore, the NIST-traceable calibration curves supplied with the device have to be assumed to be accurate within the stated uncertainty of $\pm 1\%$ [83]. Further uncertainties of the PIN photodiode and power meter are summarized in table 4.1. Assuming they are independent, the total uncertainty is 2.3%. Additional sources of uncertainty include the light output of the UV LED and the uniformity of the fibre.

Table 4.1: Sources of uncertainty for the PIN photodiode measurements due to the PIN photodiode and power meter, taken from the corresponding datasheets [83]. The uncertainties are combined quadratically. The total uncertainty containing unknown systematic effects is estimated.

*The absolute uncertainty of ± 20 pA corresponds to ± 14 pW

Source	Uncertainty
photodiode calibration	$\pm 1\%$
photodiode linearity	$\pm 0.5\%$
photodiode uniformity	$\pm 2\%$ over $\varnothing 2$ mm
power meter accuracy*	$\pm 0.25\% \pm 20$ pA
combined	$\pm 2.3\%$
total	$\pm 5\%$

The remaining uncertainties, such as the coupling between fibre and LED, are estimated to be included in a total uncertainty of $\sigma = 5\%$ per data point. The uncertainties are assumed to be uncorrelated. This assumption breaks down if attenuation does not occur gradually but in a singularity, for example, in kinks where light leaks out. For all data points where the singularity is between the PIN photodiode and the location of excitation, a part of light escapes and less light arrives compared to an intact fibre. The fraction of light loss in the singularity leads to a correlated systematic error in those data points. Such regions are excluded from the analysis if detected.

The independent variable, the distance between the PIN photodiode and the point of excitation, is known to a precision of 1 mm. It is therefore assumed to be negligible compared to the uncertainty of the dependent variable, the radiant flux Φ_e .

4.3.2 ANALYSIS

The double-exponential attenuation is described by a function of the form

$$f(z) = \Phi_e^{\text{short}} \exp^{-z/\Lambda_{\text{short}}} + \Phi_e^{\text{long}} \exp^{-z/\Lambda_{\text{long}}}, \quad (4.1)$$

where $\Lambda_{\text{short/long}}$ denote the two different attenuation lengths and $\Phi_e^{\text{short/long}}$ are the respective fluxes extrapolated to $z = 0$. The attenuation lengths of each individual fibre sample are extracted by fitting equation 4.1 to the radiant flux data obtained with the PIN photodiode. For a given set of parameters $\vec{\beta} = (\Phi_e^{\text{short}}, \Lambda_{\text{short}}, \Phi_e^{\text{long}}, \Lambda_{\text{long}})$, the residual $r_i = \Phi_e(z) - f(z, \vec{\beta})$ describes the discrepancy between the double-exponential model and a data point i . The sum of all residuals weighted with a factor $\frac{1}{\sigma_i^2}$

$$S = \sum_i \frac{r_i^2}{\sigma_i^2}, \quad (4.2)$$

is minimal for the best-fitting $\vec{\beta}$ (least squares principle [86]). As $f(z, \vec{\beta})$ does not depend linearly on the attenuation lengths, the solution is calculated using the iterative Levenberg–Marquardt algorithm provided by ODRPACK [87] through the SciPy interface [88]. The choice of algorithm assumes a negligible uncertainty in the independent variable, the distance between the PIN photodiode and the point of excitation. Multiple fibre samples of the same fibre type have been measured to account for variation between different samples.

Round BCF-12 fibres showed such strong attenuation that a double-exponential fit does not yield sensible results. The data from these three fibres are therefore described with a single-exponential model”

$$f(z) = \Phi_e^{\text{long}} \exp^{-z/\Lambda_{\text{long}}}. \quad (4.3)$$

FIT UNCERTAINTY

In addition to $\vec{\beta}$, ODRPACK provides the variance-covariance matrix \mathbf{V}_β of the fit parameters and so-called “standard errors” σ_β , calculated as

$$\sigma_{\beta, i} = \sqrt{\mathbf{V}_{\beta, ii} \cdot \frac{S}{(n-4)}}, \quad (4.4)$$

where $n - 4$ is the number of degrees of freedom in the fit to n data points. Scaling the standard deviation by $\frac{S}{(n-4)}$ is equivalent to artificially scaling the uncertainties such that S matches its expectation value if it follows a χ^2 distribution. This “trick” does not affect the best parameters $\vec{\beta}$.

4.3.3 RESULTS AND DISCUSSION OF THE PHOTODIODE MEASUREMENTS

Table 4.2 summarizes the fit results for all fibre samples. Figures showing all data and fit results are presented in appendix A.

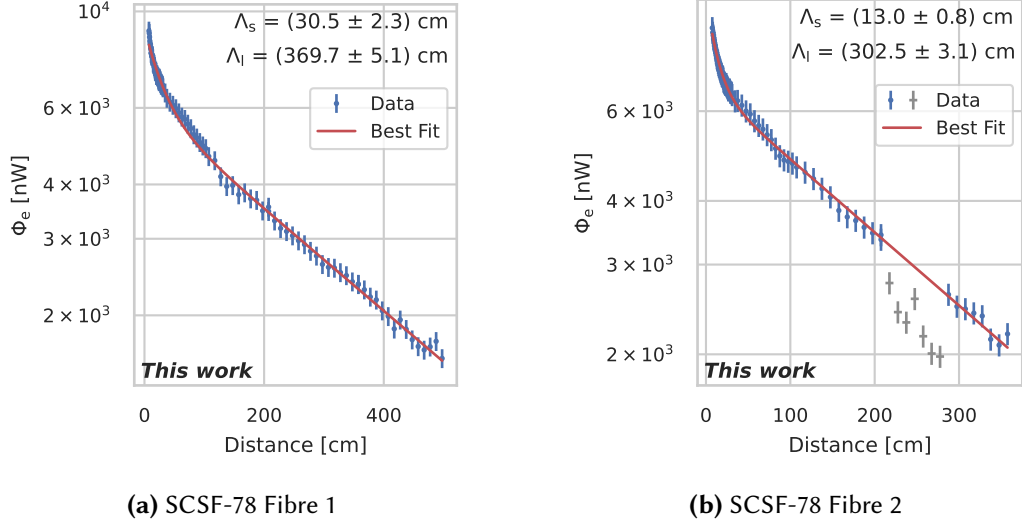


Figure 4.3: Radiant flux at one fibre end as a function of distance to the side-induced fluorescence. A double-exponential function is fit to the data in blue; grey data points are excluded from the fit. The two subfigures show data from two fibre samples of the same fibre type.

Fibres produced by Kuraray (SCSF-78, SCSF-81 and NOL-11) generally show longer attenuation lengths than the fibres produced by Saint-Gobain (all BCF types). Except for the round BCF-12 fibres, a second exponential component is observable in all samples.

The green-emitting scintillating fibre BCF-20 and the wavelength-shifting fibre BCF-92 have longer attenuation lengths than the blue-scintillating fibres BCF-12 by the same manufacturer. Although the round and square BCF-12 fibres are made of the same material and have the same scintillator agents, the attenuation length in the round fibres is significantly shorter. The round BCF-12 fibre was provided by the University of Geneva, where previous tests showed low light yield. It is unclear whether this only affects the particular spool in Geneva or all round 250 μm BCF-12 fibres. Additionally, the data for the round fibres show a dependence that is not well described by a single or double exponential function. The fit result for two of the three fibres is unphysical since the amplitude Φ_e^{short} of the short component is negative (see figures A.10e and A.11e in the appendix). Therefore, the short component of these two samples is not stated.

Upon closer inspection, the data obtained from the two Kuraray SCSF-78 show a dependence of z that is not sufficiently described by only two exponential functions, as shown in figure 4.4a. A third exponential term with an ultra-short attenuation

Table 4.2: Fit results for the short and long components of the attenuation length for each fibre. Λ_{long} and Λ_{short} are determined from the radiant flux measurements. “BCF-12 \circ ” refers to fibres with a round cross-section, while “BCF-12 \square ” refers to fibres with a square cross-section. The attenuation length Λ_{int} is calculated by numerically integrating the obtained spectra and performing the first step of the analysis (see section 4.4.2). Numbers 1 to 3 after the fibre type denote individual fibre samples of that type. The uncertainties represent the standard error as explained in section 4.3.2.

Fibre	$\Lambda_{\text{short}}[\text{cm}]$	$\Lambda_{\text{long}}[\text{cm}]$	$\Lambda_{\text{int}}[\text{cm}]$
SCSF-78 1	$(30.5 \pm 2.3) \text{ cm}$	$(369.7 \pm 5.1) \text{ cm}$	$(495 \pm 21) \text{ cm}$
SCSF-78 2	$(13.0 \pm 0.8) \text{ cm}$	$(302.5 \pm 3.1) \text{ cm}$	$(314 \pm 21) \text{ cm}$
SCSF-81 1	$(11.6 \pm 0.8) \text{ cm}$	$(216.3 \pm 8.7) \text{ cm}$	$(252 \pm 26) \text{ cm}$
SCSF-81 2	$(17.0 \pm 0.7) \text{ cm}$	$(364.2 \pm 13.0) \text{ cm}$	$(345 \pm 41) \text{ cm}$
SCSF-81 3	$(25.9 \pm 0.7) \text{ cm}$	$(269.0 \pm 4.0) \text{ cm}$	$(244 \pm 13) \text{ cm}$
NOL-11 1	$(20.8 \pm 1.4) \text{ cm}$	N/A	$(360 \pm 34) \text{ cm}$
NOL-11 2	$(29.5 \pm 1.9) \text{ cm}$	$(294.8 \pm 4.1) \text{ cm}$	$(256 \pm 7) \text{ cm}$
NOL-11 3	$(15.8 \pm 0.6) \text{ cm}$	$(199.9 \pm 3.2) \text{ cm}$	$(246 \pm 12) \text{ cm}$
BCF-12 \circ 1	N/A	$(25.1 \pm 0.3) \text{ cm}$	$(29 \pm 2) \text{ cm}$
BCF-12 \circ 2	N/A	$(26.5 \pm 0.4) \text{ cm}$	$(23 \pm 3) \text{ cm}$
BCF-12 \circ 3	N/A	$(27.7 \pm 0.4) \text{ cm}$	$(25 \pm 1) \text{ cm}$
BCF-12 \square 1	$(10.1 \pm 0.6) \text{ cm}$	$(62.2 \pm 0.2) \text{ cm}$	$(64 \pm 1) \text{ cm}$
BCF-12 \square 2	$(18.5 \pm 0.7) \text{ cm}$	$(75.2 \pm 0.8) \text{ cm}$	$(71 \pm 2) \text{ cm}$
BCF-12 \square 3	$(14.6 \pm 0.7) \text{ cm}$	$(72.9 \pm 0.7) \text{ cm}$	$(74 \pm 2) \text{ cm}$
BCF-20 1	$(11.0 \pm 1.1) \text{ cm}$	$(97.7 \pm 0.8) \text{ cm}$	$(110 \pm 11) \text{ cm}$
BCF-20 2	$(13.9 \pm 5.0) \text{ cm}$	$(105.5 \pm 5.0) \text{ cm}$	$(129 \pm 7) \text{ cm}$
BCF-20 3	$(11.7 \pm 1.3) \text{ cm}$	$(134.8 \pm 7.6) \text{ cm}$	$(120 \pm 8) \text{ cm}$
BCF-92 1	$(18.9 \pm 1.2) \text{ cm}$	$(90.6 \pm 1.7) \text{ cm}$	$(83 \pm 3) \text{ cm}$
BCF-92 2	$(12.2 \pm 0.5) \text{ cm}$	$(92.5 \pm 0.8) \text{ cm}$	$(89 \pm 4) \text{ cm}$
BCF-92 3	$(13.6 \pm 0.7) \text{ cm}$	$(92.7 \pm 1.1) \text{ cm}$	$(91 \pm 1) \text{ cm}$

length Λ_{us} improves the apparent agreement between the model and data significantly (figure 4.4b). The reduced χ^2 , i.e. $S/n = \chi_{\text{red}}^2$, where n is the number of degrees of freedom, is a measure of the goodness of fit. The two-component fit has a value of $\chi_{\text{red}}^2 \approx 0.27$, almost twice as large as the three-component fit with $\chi_{\text{red}}^2 \approx 0.13$. A larger value indicates greater disagreement between the data and the fitted model. Note that both values are less than the expected value of $\chi_{\text{red}}^2 = 1$ and a larger deviation from 1 typically indicates a *worse* fit. Visually, figure 4.4b shows better agreement than figure 4.4a. This indicates an overestimation of the total uncertainty of 5 % (as discussed in section 4.3.2). As most fibres did not show that feature, the analysis was not modified to incorporate a third component.

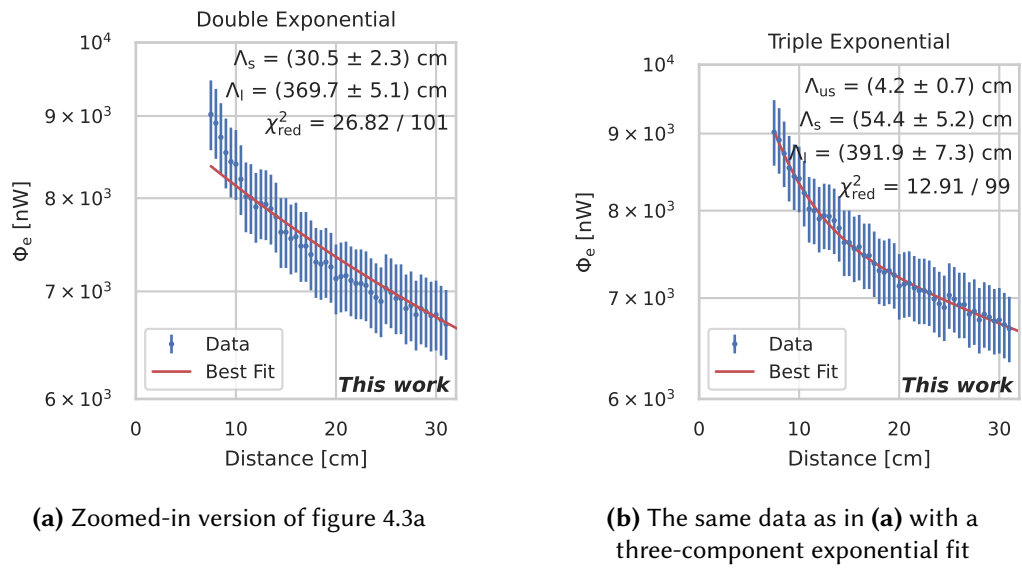


Figure 4.4: For individual fibres, the data are not well-described by just two components of the attenuation length. Fitting a model with an additional ultra-short component Λ_{us} improves the agreement between model and data significantly.

The second SCSF-78 sample shows data points between approximately 200 and 300 cm that appear to be outliers with too little light (see figure 4.3b). When running the LED beyond 300 cm, more light is recorded at the fibre end than in this “dark region”. This indicates that the light is neither absorbed nor scattered out in this region, but rather that less scintillation light is produced. Most likely, less scintillation light is generated by fluorescence, while the transmission through this region appears to be unaffected. For the data analysis, these points are excluded. This second fibre sample has significantly shorter attenuation lengths for both components.

The long components of the SCSF-81 and NOL-11 fibre samples vary even more drastically. One sample of each type shows very low attenuation (corresponding to a large attenuation length) before the light drops off after a certain distance is reached

(figure A.6e). Assuming a singular defect in the position after which the light drops off suddenly, the fit range is set up to this point. The cross-check using integrated spectra (column Λ_{int} and see below) shows a much shorter attenuation length. However, in each group, these two samples have a significantly larger attenuation length than the other samples. For the NOL-11 fibre, the drop-off point is within the range covered with the spectrometer (figure A.6d). This effect is visible in both data sets.

A possible cause of such an effect could be a defect that scatters light back into the fibre. That way, part of the light that normally is propagated in the opposite direction of the pin diode is also detected. On one side of the defect, more light would be measured than expected, while on the other side, less arrives at the pin diode. However, a visual inspection could neither confirm nor refute that hypothesis.

Different fibre samples of the same type often yield attenuation lengths incompatible within their standard errors. The differences are larger for the samples transported to PSI from Geneva. Either the Kuraray fibres have a larger variance in material properties, or the fibres were damaged during transport or storage.

4.4 WAVELENGTH-DEPENDENT ATTENUATION

An optical spectrometer allows measuring the radiant flux in a series of narrow wavelength ranges to obtain the spectral flux $\Phi_{e\lambda}$. A reflective grating diffracts the light onto a one-dimensional Charge-Coupled Device (CCD) array so that each channel of the CCD sensor is illuminated by the light of a different wavelength.

Some pixels are purposefully kept dark to determine thermal noise in the CCD sensor. The signal of these pixels is then subtracted from the rest, known as “electric dark correction”. Since the detector response is not perfectly linear with integrated radiant flux, a “non-linearity correction” is applied. These two features are provided by the manufacturer with a claimed linearity of $> 99.8\%$ [84]. The wavelength spectra are recorded using the OceanView spectroscopy software in version 1.6.7.

Dark spectra are recorded before a measurement using the same integration time and other settings but without side-induced fluorescence. This dark spectrum is then subtracted from the spectrum of interest to account for light leaks. Figure 4.5 shows a set of spectra from the same fibre. The measured radiant flux is reduced when the fibre is stimulated far from the spectrometer. This reduction is strongest at short wavelengths and around 530 nm and 610 nm.

Each spectrum is obtained by stimulating the fibre at a different distance from the fibre end coupled to the spectrometer. The spectrometer records spectral flux as “intensity” in *counts*. This intensity can be compared between two spectra in the same set but can not be compared to the absolute spectral flux obtained with the PIN diode due to a lack of calibration. Furthermore, the fibre coupling to the two devices is independent and not controlled. As the fibres were excited using UV light, no

comparison of the absolute light yield from ionizing particles can be made. The radiant flux of the same fibre sample ranges up to three orders of magnitude for excitation in different positions. Just like the exposure time of a camera is adjusted to lighting conditions, the CCD integration time is adjusted for every recorded spectrum.

4.4.1 DIFFERENTIAL ATTENUATION LENGTH

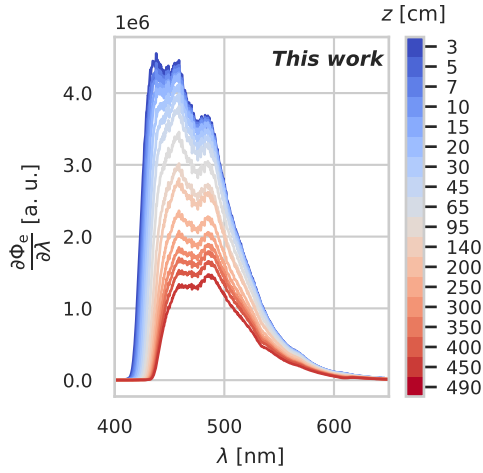
The decrease in spectral flux in the corresponding CCD pixel is studied to determine the attenuation length at a specific wavelength. This corresponds to a vertical slice in figure 4.5c at each point in λ . The light intensity as a function of the distance between the fibre end and stimulation is fitted with a single exponential function (see equation 4.3). The best fit value and its standard error are presented as a function of the wavelength (e.g. figure 4.6 and appendix A). The spectra are smoothed by a Gaussian filter with a width of 15 CCD pixels to mitigate noise.

Naturally, the differential attenuation length is shorter for fibres with short and longer for those with longer attenuation lengths, reproducing the previous results qualitatively. The minimum attenuation length lies at the lower edge of the spectra. The attenuation length appears to increase for the tail at short wavelengths, although the spectral flux at these wavelengths is close to zero. In detail, however, differences can be observed that are discussed in the following.


KURARAY FIBRES The differential attenuation length in SCSF-81 fibres seems to increase relatively steadily over most of the emission spectrum. The slope of the differential attenuation length of the other two Kuraray fibres, SCSF-78 and NOL-11, changes at around 450 nm. Local minima are visible at 530 nm and 610 nm for all three Kuraray fibres. These minima correspond to absorption maxima in polystyrene.

SAINT-GOBAIN CRYSTALS FIBRES The fibres produced by Saint-Gobain Crystals also show a change in the slope of the differential attenuation length, albeit in different locations. The blue-emitting fibres show this feature below 450 nm, while the green-emitting fibres show it above 500 nm. Like the Kuraray fibres, local absorption maxima are visible at 530 nm and 610 nm. The minima are less pronounced in the Saint-Gobain fibres than in the Kuraray fibres. Since the overall attenuation length is shorter, effects due to the polystyrene core have a smaller relative effect.


Comparing the blue-emitting BCF-12 and green-emitting BCF-20 scintillating fibres, one can see that above 500 nm, light attenuation is stronger in the blue-emitting fibres (figure 4.7). Combined with the shifted emission spectrum, this explains the longer overall attenuation length of 97.7 to 134.8 cm in BCF-20 fibres compared to 62.8 to 73.5 cm in square BCF-12 fibres.




(a) Radiant flux spectra of SCSF-78



(b) The spectra from (a), divided by the one with the shortest distance between the fibre end and excitation



(c) Radiant flux spectra of BCF-12 (square)



(d) The spectra from (c), divided by the one with the shortest distance between the fibre end and excitation

Figure 4.5: Emission spectra taken from two different samples stimulated at multiple distances to the fibre end coupled to the spectrometer (colour coded)

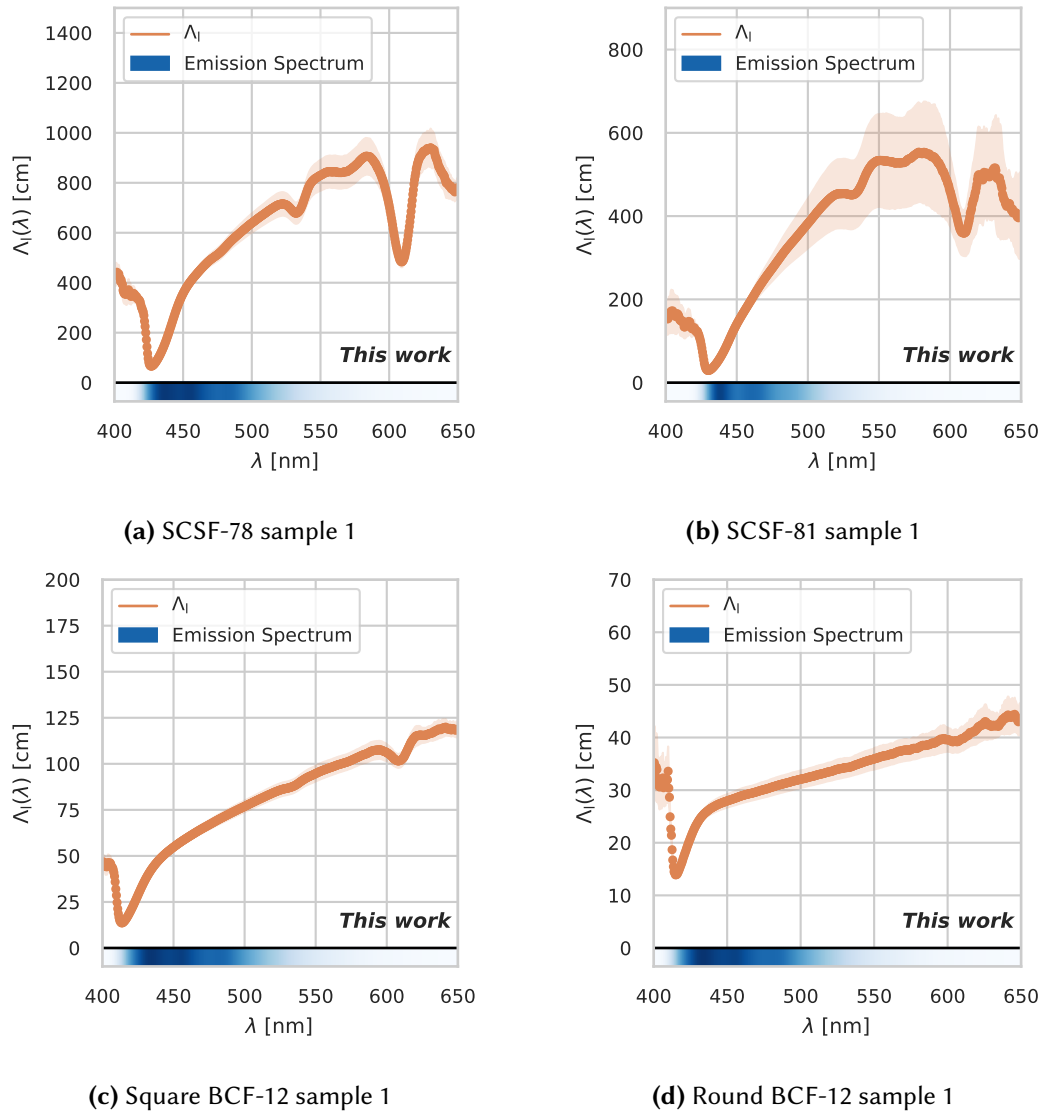
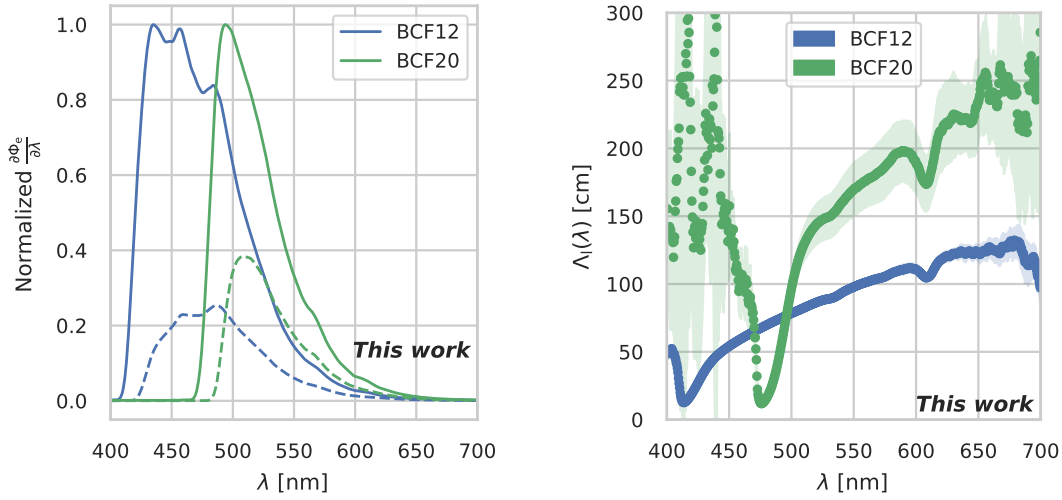


Figure 4.6: The *differential* attenuation length of different fibres. For each wavelength on the abscissa in e.g. figure 4.5, the intensities of all distances are fit to determine the attenuation length at this particular wavelength. The statistical uncertainty from the fit algorithm is shown as a light orange band. This does not include any systematic uncertainties (see also section 4.3).



(a) Normalized emission spectra of BCF-12 and BCF-20, the solid line is at 7 cm, the dashed line at 95 cm (b) Differential attenuation length of BCF-12 and BCF-20

Figure 4.7: Comparison of emission spectra and differential attenuation length between blue- and green-emitting scintillating fibres by Saint-Gobain

4.4.2 ATTENUATION OF INTEGRATED WAVELENGTH SPECTRA

As a cross-check to the values obtained using the power meter, the spectra are integrated numerically, and the wavelength-independent attenuation length is determined again. The wavelengths associated with each CCD pixel are not equidistant: The first two are 0.471 nm apart, while the last two are 0.352 nm apart. This is accounted for in the numerical integration by using the composite trapezoidal rule [89]. The number of wavelength spectra recorded is lower than the number of radiant flux measurements. Apart from the number of data points, the fit method is the same as for the long component (see section 4.3.2).

Especially at short distances, the spectra of certain fibres show seemingly non-physical results. Examples can be seen for all three NOL-11 fibres (see figures A.4.1 through A.4.3 in the appendix). The spectra in which the fibres are excited 3 cm from the fibre end appear to have a lower radiant flux than the ones at 5 and 7 cm. A possible reason lies in the coupling to the spectrometer. The shorter the distance between the spectrometer and the LED, the stronger the bending due to slight misalignments, which leads to light loss. A relative uncertainty of 10 % is assumed for each integrated spectrum. The values Λ_{int} and their standard errors (see section 4.3.2) are included in table 4.2 in column Λ_{int} . One example of such a fit is shown in figure 4.8. For all other fibre samples, see appendix A.

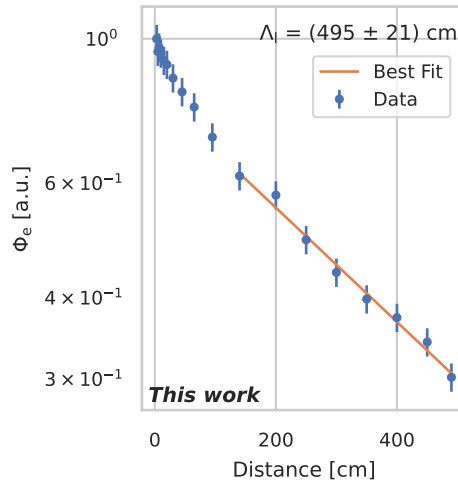


Figure 4.8: Determination of the long component of the attenuation length by integrating the spectra (see figure 4.5) and fitting a single exponential function in the range covered by the solid line. Each point in this figure corresponds to one integrated wavelength spectrum. The result is used to cross-check the attenuation length determined by measuring the radiant flux.

4.5 DISCUSSION OF THE ATTENUATION STUDIES

The attenuation behaviour of scintillating fibres at distances in the order of the size of the Mu3e SciFi detector is studied. In the following, a comparison to the existing literature is made, and the findings are discussed.

4.5.1 COMPARISON TO LITERATURE

There are several publications presenting attenuation lengths of scintillating fibres. Selective measurements are summarized in table 4.3 and discussed in the following.

Antonelli et al. published studies of blue-emitting scintillating fibres with a diameter of 1 mm [90]. They report the following attenuation lengths: BCF-12 (produced in 1992): (226 ± 3) cm, BCF-12 (produced in 1993): (286 ± 8) cm, SCSF-81 (produced in 1992): (321 ± 5) cm. The attenuation length of the BCF-12 fibres is several times greater than the measurements presented in this chapter (25.1 cm to 27.7 cm for round fibres and 62.2 cm to 75.2 cm for square fibres). The manufacturer states that “In small fibers (≤ 0.5 mm), the fluor concentration is increased, usually at the expense of light attenuation length” [59]. The increased fluor concentration is a possible explanation for the observed differences.

The datasheet by Kuraray [41] states attenuation lengths of greater than 4.0 m and greater than 3.5 m for 1 mm SCSF-78 and SCSF-81 fibres, respectively. With 303 cm to

370 cm the SCSF-78 fibres tested herein have a shorter attenuation length than stated by the manufacturer. All samples of SCSF-81 but one lie below the stated value with 216 cm to 269 cm. The exception is the sample where the result is potentially inflated by a measurement artefact, as discussed in section 4.3.3.

Furthermore, the datasheet includes differential attenuation in units of dB/km. Using equation 3.11, their data are converted to attenuation lengths to allow easier comparison, as shown in figure 4.9. Qualitatively, their differential attenuation length agrees with the results presented here. The local minima at 530 nm and 610 nm can be explained by attenuation in the PS of the core material, as shown in figure 3.2 and are therefore expected. However, overall their attenuation lengths are far larger than the ones obtained here. The only fibre that shows an attenuation length of more than 10 m for any given wavelength is again the second SCSF-81 sample. Their fibres have a diameter of 1 mm, potentially increasing the attenuation length significantly. Unlike Saint-Gobain, Kuraray does not dwell on fluor concentrations, but thin fibres may have a higher fluor concentration as well.

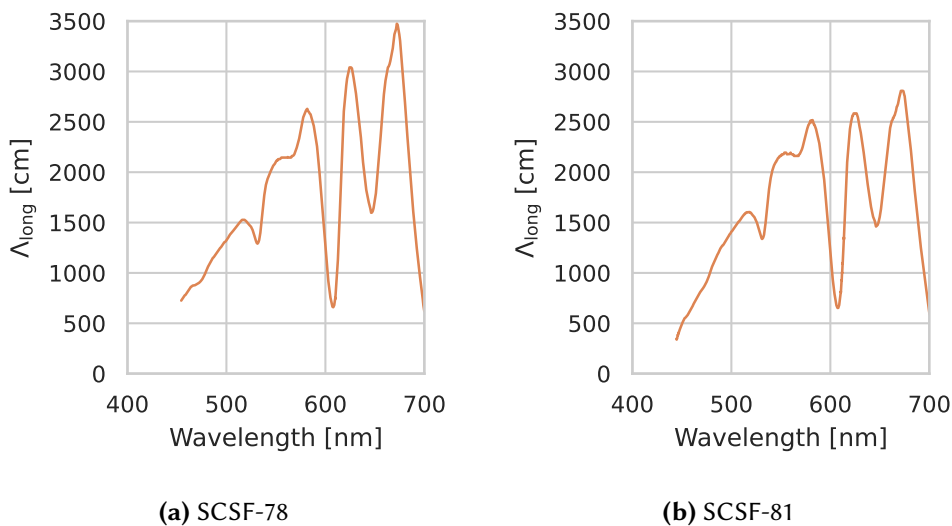


Figure 4.9: Attenuation lengths for $\varnothing 1$ mm scintillating fibres published by the manufacturer, derived from the figures titled “Transmission Loss” in [41]

Figure 3 in [58] from Borshchev et al. shows the attenuation per wavelength of SCSF-78 and NOL-11 fibres in a similar manner as section 4.4.1 of this chapter. As here, their fibres have a diameter of $250 \mu\text{m}$. To compare the work of this thesis to theirs, figure 4.10 shows an overlay of their data on the results herein. The measurements agree significantly better than the comparison to thicker fibres before. The “best” performing fibres I measured perform similarly to what Borshchev et al. present. The consistency between these measurements of $250 \mu\text{m}$ fibres and the ones herein leads to the hypothesis that $250 \mu\text{m}$ fibres from Kuraray have an intrinsically shorter

Table 4.3: Summary of literature comparison between the measurements in this work (on 250 μm fibres) and selective other measurements

Fibre	Result	This work	Comment	Reference
BCF-12 (1992)	$\Lambda_{\text{long}} = (226 \pm 3) \text{ cm}$	62.2 cm to 75.2 cm (square)	1 mm diameter, higher fluor concentration than in thin fibres [59]	[90]
BCF-12 (1993)	$\Lambda_{\text{long}} = (286 \pm 8) \text{ cm}$	62.2 cm to 75.2 cm (square)	1 mm diameter, higher fluor concentration than in thin fibres [59]	[90]
SCSF-81 (1992)	$\Lambda_{\text{long}} = (321 \pm 5) \text{ cm}$	216 cm to 269 cm, ex. sample 2	1 mm diameter	[90]
SCSF-78	$\Lambda_{\text{long}} > 4.0 \text{ m}$	303 cm to 370 cm	1 mm diameter	[41]
SCSF-81	$\Lambda_{\text{long}} > 3.5 \text{ m}$	216 cm to 269 cm, ex. sample 2	1 mm diameter	[41]
SCSF-78	Figure 4.9a	Figure 4.6a	1 mm diameter, shape matches, but they have less attenuation	[41]
SCSF-81	Figure 4.9b	Figure 4.6b	1 mm diameter, shape matches, but they have less attenuation	[41]
NOL-11	Figures 4.10a and 4.10b		Shape matches. Sample 1 agrees; sample 2 has 40 % smaller attenuation lengths	[58]
SCSF-78	Figures 4.10c and 4.10d		Shape matches. They have 36 and 11 % longer attenuation lengths	[58]

attenuation length than their 1 mm counterparts, even though it is not explicitly stated. What remains to be explained is why some fibres have a short attenuation length without visible damage. Potential deterioration could have occurred either before or during transport. Different storage conditions and consequent damage to the fibres are additional sources of systematic uncertainties.

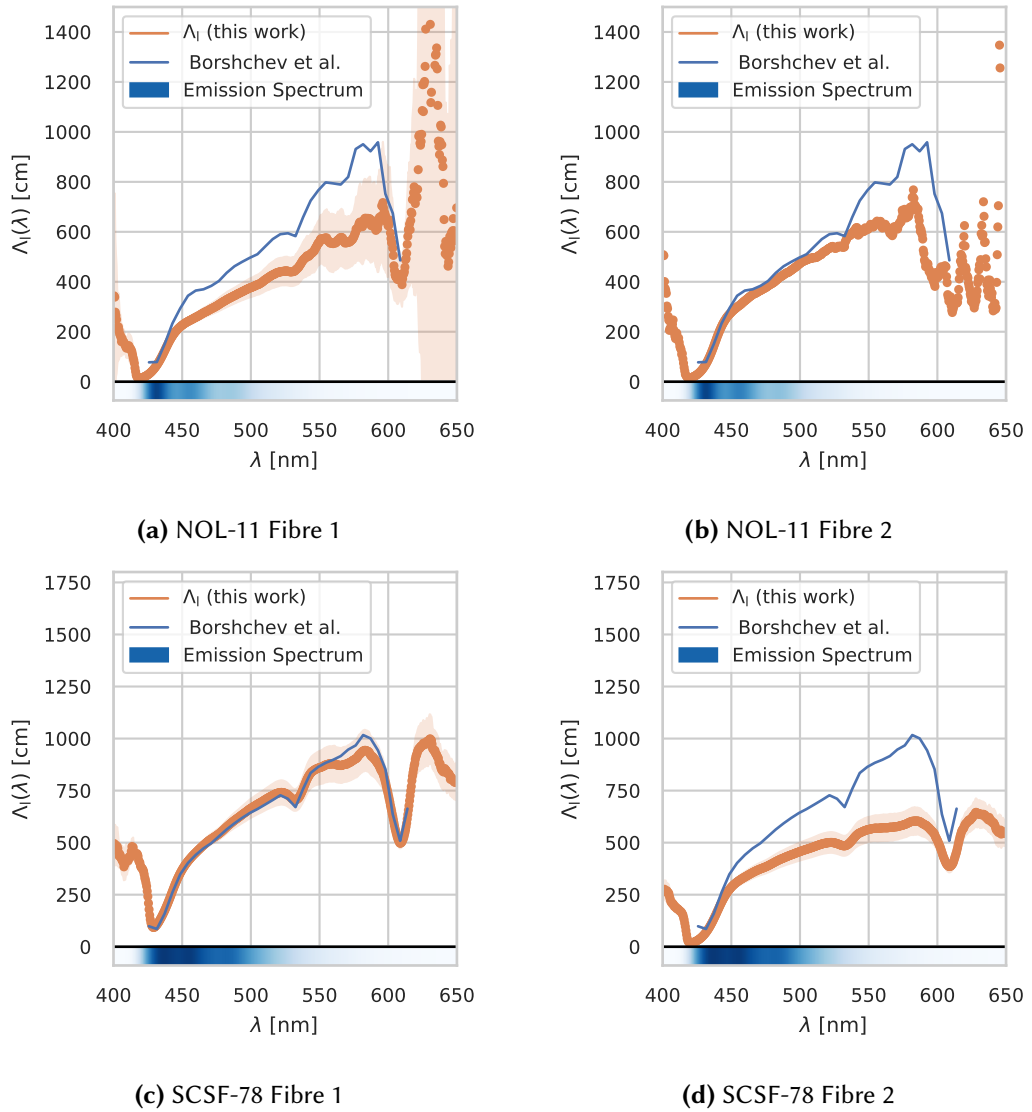
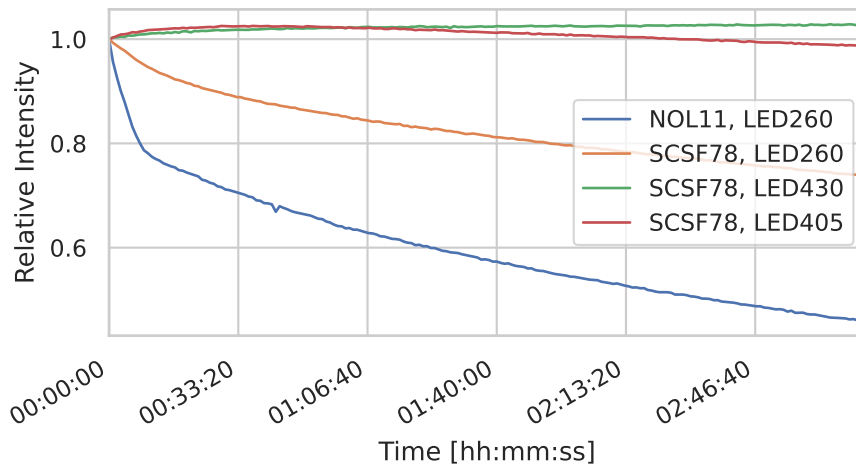


Figure 4.10: The differential attenuation length of different fibre samples from this work. Superimposed are the data from Borshchev et al. [58]. Both NOL-11 fibres studied here show slightly more attenuation than their measurement (36 % and 11 %, respectively). The differential attenuation of the first SCSF-78 sample in (c) is in good agreement to their measurement, while the second fibre shows 40 % stronger attenuation

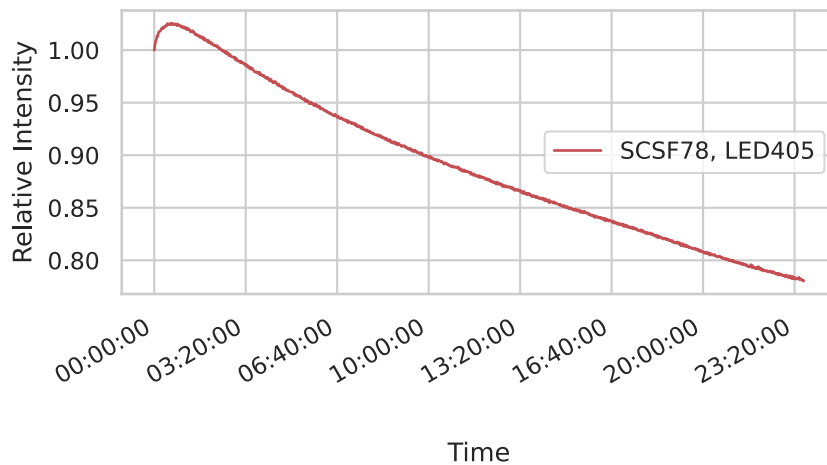
4.6 UV DAMAGE IN SCINTILLATING FIBRES

A possible source of error with side-induced fluorescence is the destruction of the scintillating fibre with UV light. Fibre samples were exposed to LEDs with different wavelengths at their maximum rated current to investigate the potential damage. In regular time intervals, a spectrum is recorded, which is then integrated to determine if the light output changes. The result is shown in figure 4.11.

Especially when illuminated with light of very short wavelengths, one can observe a decrease in light output. Wavelengths above 400 nm seem to not damage SCSF-78 fibres significantly over hours. However, a longer exposure with 405 nm light does reduce the light output by more than 20 %. Moving the LED further away from the fibre end, thus placing the illuminated region between the LED and the fibre end, the light yield increased. Therefore, one can assume that the decrease in light yield can primarily be attributed to less scintillation light being produced rather than stronger attenuation in this region. The illumination was kept to a minimum to keep the potential damage induced during measurements negligible. The implications of the damage caused by UV light on detector construction are discussed in section 4.7.1. Reversibility or annealing of this kind of damage was not studied.



(a) Comparison of UV damage in NOL-11 and SCSF-78 fibres and different LEDs over three hours



(b) Long-term (24 hours) exposure of an SCSF-78 fibre to light from an LED with a wavelength of 405 nm

Figure 4.11: The relative light intensity (radiant flux) detected from fibres during long-term excitation with UV light from different LEDs. The number indicates the wavelength in nanometres.

4.7 CONCLUSIONS OF THE ATTENUATION MEASUREMENTS

When a scintillating or wavelength-shifting fibre is excited far from a fibre end, the light attenuation is well-described by an exponential decay of the radiant flux Φ_e as a function of distance. The attenuation length of different fibre types excited at large distances varies from well under 30 cm to several metres. The fibres under study by Kuraray have consistently larger attenuation lengths than those by Saint-Gobain Crystals. Six of the seven studied fibre types show an attenuation pattern better described by a double-exponential function with an additional light component that extinguishes at much shorter lengths between 10 and 35 cm. A two-component fit of the light yield from various fibre types is presented. For individual fibres, indications for additional components of the attenuation length beyond two were found. Additionally, the attenuation is studied as a function of the wavelength of the scintillation light.

Small-scale detectors such as the Mu3e scintillating fibre detector with a length of approximately 30 cm are strongly affected by this effect. On the one hand, an extrapolation from the vendor-stated long attenuation length at large scales *underestimates* the amount of light detected in a short detector. On the other hand, the attenuation length is shorter than just the long attenuation length suggests. If the light of fibres is detected with photodetectors whose Photon Detection Efficiency (PDE) has a strong wavelength dependence, the change in spectrum needs to be considered.

In detail, the measurements showed difficulties when handling fibres. The first set of fibres from Geneva was wrapped in a paper envelope, and the fibres were so damaged that no meaningful measurements were possible. A second batch was wrapped from one spool to another and transported with utmost care. Nevertheless, the samples under study here were still of heterogeneous quality. The fibres cut at PSI did not suffer from this effect. Before advancing, the cause for the heterogeneity should be evaluated. It is unclear if the fibres were already damaged on the spool in Geneva or afterwards during the transfer to a second spool and transport by train. Visual inspection did not yield any clues.

The performance of the round BCF-12 fibres remains an unresolved issue. Technically, one would expect little to no difference between the round and square fibres. A difference in *trapping efficiency* (see section 3.1.2) does not translate to vastly different path lengths inside the fibres. Earlier studies (e.g. [40, 91]) have shown that round BCF-12 fibres have a significantly lower light yield than other fibre types. The measurements of this chapter were obtained with samples from the same spool as in [40] and [91]. The results here show that the light output is low, and the attenuation is stronger than in the other fibres. While measuring different quantities, the results herein agree with the conclusion that the performance of round BCF-12 fibres is significantly worse than all other fibres under test.

Figure 4.12 shows the emission spectrum of an SCSF-78 fibre at 3 and 30 cm (the length of the SciFi detector), as measured herein. Additionally, the PDE of a Hamamatsu

S13552 SiPM⁴ is overlaid. Despite the slight shift towards long wavelengths, the main part of the spectrum remains in the range in which the PDE exceeds 40 %. Therefore, the change in shape is not expected to significantly impact the Mu3e SciFi detector’s performance. The tentative decision to use this fibre type (and SiPM) could be backed up by these measurements.




Figure 4.12: Relative emission spectra of Kuraray SCSF-78 fibre 1 at 3 and 30 cm, the size of the Mu3e detector. Overlaid is the Photon Detection Efficiency (PDE) of the Hamamatsu S13552 SiPM (data from [42]).

4.7.1 EXPERIMENTAL ASPECTS AND DETECTOR CONSTRUCTION

Several observations during these measurements have implications for constructing the Mu3e SciFi detector. Figure 4.11b shows that light with a wavelength of 405 nm reduces the light output of SCSF-78 fibres by more than 20 %.

Fluorescent lamps in many laboratories work with mercury vapour that emits ultraviolet light through excitation. Although most of the light is converted into visible light in the fluorescent coating of the tube, some light at 405 nm is emitted [92]. Therefore, Kuraray SCSF-78 scintillating fibres (and potentially others) should not be stored without protection in rooms with fluorescent lighting or sunlight. Short exposure appears not to cause significant damage, as it took almost 24 hours and an LED in close proximity to reduce the light output by 20 %. In direct comparison, NOL-11 fibres showed a faster decrease in light output with UV illumination. Exposure of these fibres should be even more thoroughly minimized. For handling, artificial lighting without UV light is recommended. During transport, fibres and fibre assemblies should be carried in a light-tight enclosure.

⁴the other tentative choice for Mu3e, see also the next chapter

4.7. CONCLUSIONS OF THE ATTENUATION MEASUREMENTS

Although the fibres were transported and handled with a high level of care, the fibres from Geneva had more defects than the ones obtained at PSI where the measurements took place. This suggests that they could have been damaged during the transfer to a second spool and transportation with public transport. The scintillation light escaping in damaged regions was visible to the naked eye. Kinks in the fibre are also immediately visible, but light-escaping regions could also be found in straight, seemingly undamaged regions. It is, therefore, crucial to handle the bare fibres as little as possible during the production of the SciFi ribbons for Mu3e. Single fibres should only be transported on the spool by the manufacturer and not individually.

Damage to fibres is often very localized. Passing a localized defect while scanning over the fibre with an LED, the light output decreases abruptly. This can be seen in sections A.4.1 and A.3.3. A significant amount of light escapes the fibre, which can be seen as a blue glowing at that position. Such defects could be identified visually by exciting the fibre at regular intervals. The affected portion of the fibre could be cut out and excluded during ribbon production. The production workflow must allow cutting the fibre and continuing with a new fibre end. The short length of 30 cm offers many positions for potential cuts, thus minimizing the length of the discarded fibre. In the same manner, local variations of fibre diameter can be cut out if they are detected visually.

5

Irradiation studies of silicon photomultipliers

5.1 INTRODUCTION

Radiation damage in silicon photomultipliers is a phenomenon studied in different contexts. Most commonly, the sensors are irradiated with neutrons to extrapolate the resistance against damage caused by hadrons, be it in high-energy experiments or in space. The radiation environment in Mu3e differs from those in typical high energy physics experiments. While the detectors at the Large Hadron Collider are exposed to hadrons with energies up to the TeV scale, the radiation inside Mu3e is dominated by positrons from muon decays at rest with an energy of up to 53 MeV. Furthermore, the Mu3e detector is exposed to virtually no neutrons in contrast to high energy experiments. Comparing the displacement damage cross-section (see figure 3.7) of electrons and hadrons shows that electrons are expected to cause several times less radiation damage in silicon than hadrons.

The SiPM model tentatively chosen for the Mu3e SciFi detector is identical to the one used in LHCb: the Hamamatsu S13552, formerly denoted as “high R_Q ”. The LHCb collaboration expects an equivalent fluence of 1 MeV neutrons of up to $\Phi_{\text{eq}} = 13 \cdot 10^{11} \text{ cm}^{-2}$. Part of the evaluation by the LHCb collaboration was the irradiation with neutrons that showed an increase in dark current proportional to the neutron-equivalent fluence [57]. That dark current increase is explained by an increase in DCR. The LHCb detector is operated at -40°C and equipped with a neutron shield to keep the DCR low enough throughout the lifetime of the experiment.

Despite the seemingly “harmless” positron radiation inside Mu3e, radiation hardness must also be considered. As part of the evaluation of the SiPM for Mu3e, the expected fluence must be determined. This chapter presents a detailed simulation of the positron

fluence from the $2.5 \cdot 10^{15}$ muon decays needed to achieve the Phase I goal of Mu3e [1]. Due to the large number of positrons and the close vicinity to the decay target, the SciFi SiPMs are exposed to up to $2.8 \cdot 10^{11} \text{ cm}^{-2}$ of neutron-equivalent fluence—almost a third of the fluence in LHCb.

A dedicated setup to irradiate SciFi SiPM arrays was designed and built at ETH Zürich as part of this thesis. Using this setup, six SiPMs were irradiated with positrons from muon decays, emulating the radiation inside Mu3e as closely as possible.

The irradiated SiPMs allow the study of the impact on the dark current, DCR and a possible change in breakdown voltage. A critical limitation is the maximum data rate of the readout ASIC MuTRiG. With an expected rate of 650 kHz of real signals and an upper limit of 1.1 MHz the DCR must not exceed 450 kHz [40].

Section 5.5 presents these studies of the irradiated SiPMs. While dark current and any change in breakdown voltage appear manageable, the DCR at room temperature lies in the range of 100 MHz when extrapolated to a full Phase I. In reaction to the findings of this thesis, the design of the SciFi detector was revised to allow operation below -10°C . Furthermore, a new feature was implemented in MuTRiG 2 that allows DCR suppression through hit validation by neighbouring channels. In addition to these countermeasures, this chapter discusses further ways of mitigation, such as a modified operation, thermal annealing, and as a last resort, a replacement of the SiPMs, which is possible due to the modular design of the detector.


The focus on radiation damage in the SiPMs is owed to literature research suggesting that the SiPMs are more prone to radiation damage than the scintillating fibres [3, 93]. In addition to the Hamamatsu S13552 SiPMs, the irradiation setup allowed the irradiation of SiPMs and scintillators of the Mu3e tile detector [94] and polyimide used in the Mu3e pixel tracker, where it is exposed to muons and positrons.

5.2 SETUP


A high rate of muons is necessary to emulate the radiation environment of Mu3e and conduct irradiation in a reasonable timeframe. The PiE5 beamline at PSI offers rates in the order of 10^8 muons per second and was therefore chosen (see section 2.2.2). During irradiation with Michel positrons, the SiPMs were placed in a dry helium atmosphere at a controlled temperature and biased above their breakdown voltage.

5.2.1 IRRADIATION CHAMBER

An aluminium chamber with an entrance window was designed and constructed to keep the required conditions and allow direct mounting to the PiE5 beamline. The chamber consists of two parts: a “lid” mounted to the beam pipe and a detachable “base”. A drawing of the base is shown in figure 5.1. When closed, the box dimensions are $50 \times 35 \times 8 \text{ cm}^3$ with a usable depth of 4.6 cm inside.



(a) Technical drawing missing tubing, meteorology and SciTile feedthrough



(b) Photograph with all components

Figure 5.1: Open base of the irradiation chamber: The green boards to the left and right serve as electrical feedthroughs for the SiPMs. “Aare” and “Berg” denote two landmarks at PSI which are used as references. The “Aare” side lies at positive x values, the “Berg” side at negative x . Seen in $+z$ direction.

Inside the chamber, a copper cooling block holds up to two custom-made Printed Circuit Boards (PCBs) for silicon-photomultiplier arrays and one tile matrix. Cold water can be pumped through the block to keep the SiPMs at a regulated low temperature. Four openings on the base and two on the lid allow for gas exchange inside the chamber and evacuation.

The lid has a window in which a Mylar (biaxially-oriented polyethylene terephthalate) foil is the barrier between the chamber and the beam pipe. This *beam window* must be thin enough to allow surface muons to pass through while sturdy enough to withstand a pressure difference of more than 1 bar since the beam pipe is evacuated. Right below the beam window is space for a smaller chamber used to irradiate polyimide foil and a 4 mm thick polyethene plate that is thick enough to stop all incoming surface muons and large enough to cover the entire entrance window.

Various sensors monitor the conditions inside the chamber: An electrochemical oxygen sensor by AlphaSense [95] and readout electronics designed for this setup, a USB weather station by Yoctopuce [96] and one temperature sensor on each SiPM carrier board. The electrical feedthroughs for all sensors are designed specifically for this chamber.

A desiccant agent is placed inside the chamber to bind water vapour outgassed by materials inside the chamber.

5.2.2 SiPM CARRIER BOARD




Figure 5.2: Rendering of the SiPM carrier board used for irradiation. Apart from the two connectors, the backside is flat.

A custom carrier board for the SiPM arrays was designed for irradiation (see figure 5.2). Each board contains one Hamamatsu S13552 SiPM array with 128 channels and measures $51 \times 32 \text{ mm}^2$. Those channels can be biased and read out via four flexible flat cables compatible with the prototype boards used by the Mu3e fibre group. While the existing prototypes route neighbouring channels on the same flexible printed circuit,




Figure 5.3: Schematic of the lowpass filter and bias of the SiPM on the carrier board. The two halves of the SiPM can be biased individually or combined if a zero-ohm link (R3) is used.

two neighbouring channels are routed on opposite sides of the carrier board. Therefore, the carrier boards are not drop-in replacements, but the channel mapping needs to be considered. For mechanical compatibility with existing and future setups, two sets of mounting holes match the different configurations. A digital temperature sensor is placed less than 5 mm from the SiPM to monitor the temperature. Both halves of the SiPM array can be biased independently or with a single high-voltage line configurable via a removable zero-ohm link (see figure 5.3). Each high-voltage line is equipped with a lowpass filter close to the SiPM array. The board fulfils standard design rules and presents a cost-effective alternative to rigid-flex modules with printed cables. In the final design of Mu3e, the connection between SiPMs and readout ASICs will be fixed. Therefore the SiPMs can only be connected to them via soldering or a custom adapter (see section 6.4.2).

5.2.3 SAMPLES

For this irradiation campaign, eight carrier boards were assembled, of which seven were equipped with a SiPM array. The board without a SiPM served as a blank to determine the effects of irradiation on the board or the components used in the lowpass filter. One board with a SiPM array was kept as a control sample and not irradiated. The remaining six samples were irradiated with different doses, as presented in section 5.4. The irradiation chamber holds up to two carrier boards which can be irradiated simultaneously. The two locations are denoted as “Aare” and “Berg”.

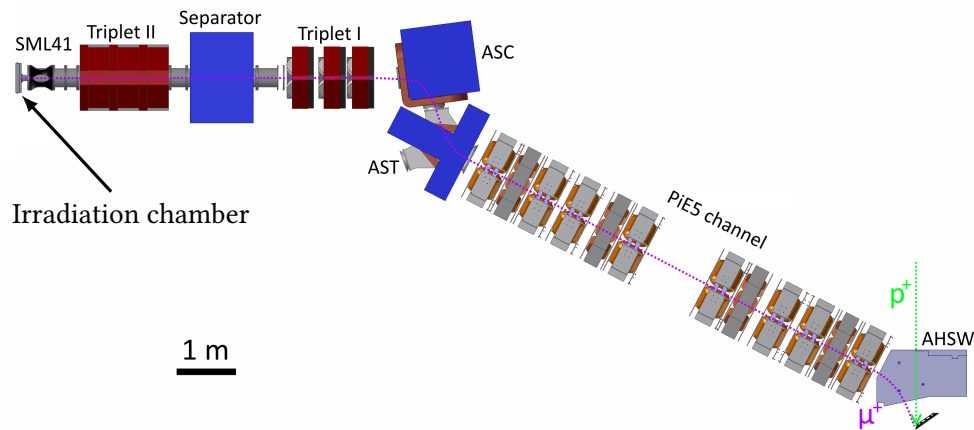


Figure 5.4: PiE5 beam line and elements with irradiation chamber. Adapted from [50].

5.2.4 BEAMLINER

The irradiation took place at the PiE5 beamline at PSI, chosen because of its high rate of surface muons. The lid of the irradiation chamber was mounted to the collimator and steering system SML41, serving as the interface between atmospheric pressure and the vacuum inside the beam pipe (see figure 5.4). As the irradiation campaign was the first use of the beamline in 2020 after the scheduled shutdown in winter, the beamline parameters had to be re-optimized for a maximum muon rate. A cylindrical scintillator with a cross-section of 3.113 mm^2 was used with a photomultiplier tube to determine the beam profile and rate. Muons that hit the pill scanner can be distinguished from positrons by the higher amplitude of the produced signal that is discriminated at two thresholds. One is optimized for muons only, and the other is set to identify muons and positrons. The discriminated signals are fed to a scaler to measure the hit rate. Additionally, the current of the proton beam at PSI is saved, and the hit rate is then normalized to the nominal beam current of 2.2 mA. At the optimal settings, the beam muons have a momentum of about 28 MeV [50].

Except for the separator SEP41 (a Wien filter in the beamline), the settings of each beamline element are read out and logged using EPICS [97].

5.3 DOSE SIMULATION

First estimations of the expected dose in the scintillating fibre detector were published in [40]. This section provides an update incorporating the final geometry of the experiment as well as the goal number of muon stops on target of $2.5 \cdot 10^{15}$ for Mu3e Phase I as described in [1].

The setup for accelerated irradiation is simulated to provide normalization to the dose relative to what is expected in the Mu3e experiment. A common quantity has to be

determined to map irradiation in a test setup to the final experiment. In both settings, the predominant particle type is positrons from Michel decays, i.e. with a kinetic energy of up to 53 MeV. The most naive scaling would be the number of positrons passing through each SiPM. However, the average path length inside the silicon differs due to the geometry and helical trajectories in the final experiment. Comparing the fluences would take directional irradiation into account. However, the fluence disregards the differences in energies and the subsequent difference in potential for damage.

Weighing the *spectral* fluence with the displacement damage cross-section (see figure 5.5) allows taking those differences into account. In section 3.4.1, the shortcomings of the NIEL hypothesis and the conversion to neutron-equivalent fluence are discussed. While conversions across particle types require a more elaborate treatment, here, the conversion to neutron-equivalents serves the main purpose of comparing fluences of < 55 MeV positrons with each other. Errors that affect both positron-to-neutron fluence conversions cancel out. For simplicity, the displacement damage cross-section based on the Robinson partition function provided by [98] is used. The same displacement damage cross-section is used for electrons and positrons.




Figure 5.5: Displacement damage cross-section of electrons in silicon as a function of the electron kinetic energy E . This function is used to weigh the fluences in Mu3e and the irradiation setup. Based on [98].

The simulation of the Mu3e detector and the irradiation setup both use the same coordinate system explained in chapter 2. For the position inside Mu3e, the SiPMs are numbered as depicted in figure 2.5. The positions are enumerated clockwise around the detector's central axis, as seen from the upstream side. The highest point of the scintillating fibre detector is the boundary of the SiPMs with indices 0 and 11. Since the fibre ribbons are staggered along z , two neighbouring SiPMs are at different positions in z . On the downstream side, SiPM 0 is at $z = 145.375$ mm, SiPM 1 at $z = 157.875$ mm,

while on the upstream side, SiPM 0 is at $z = -157.875$ mm and so on.

5.3.1 RADIATION IN MU3E FROM e^\pm

Due to the polarization of the muon beam in the $+z$ -direction and the magnetic field of 1 T inside the experiment, the upstream sensors will receive a higher dose than the downstream sensors. The staggering accounts for a difference of 1.5 % to 6 % between sensors closer to the target (staggered inwards) and those farther away from the target. In addition, the shape and position of the muon beam spot and the unequal stopping distribution¹ on the target lead to an inhomogeneous dose distribution. Figure 5.6 shows the stopping distribution inside the Mu3e detector, simulated for this study using the Mu3e framework. Furthermore, the upstream SiPMs are closer to the beampipe, where muons stop and decay (see *collimator* and *narrowing* in figure 5.6).




Figure 5.6: Vertex position for Michel positrons reaching the SciFi SiPMs in simulation. The muon beam is generated upstream (left) and a lead collimator at $z = -400$ mm reduces the beam diameter. Due to the helical tracks of the muons, some muons pass the collimator and are stopped in a narrowing of the aluminium beampipe at $z = -135$ mm where the beam pipe changes diameter from 30 mm to 20 mm. Most muons reaching the SiPMs decay on the stopping target around $z = 0$ or the first layer of pixel sensors. The SciFi SiPMs are located at approximately $r \approx 60$ mm and $z \approx \pm 152$ mm. Muons stopped in the narrowing are responsible for a significant part of the radiation in the upstream SiPMs.

¹See Fig. 6.1 in the Mu3e technical design report [1].

NEUTRON-EQUIVALENT FLUENCE

The neutron-equivalent fluence Φ_{eq} is derived using the spectral electron and positron fluence and the damage function shown in figure 5.5. Figure 5.7 shows the spectral fluence Φ_e in the upstream and downstream SiPMs as simulated herein.

With the geometry of the scintillating fibre detector, particles originating from the target region can impinge on the silicon photomultipliers either via a direct path or by recurling in the magnetic field. At the point of impact \vec{x} , incoming and outgoing particles are distinguished using the 2D scalar product of the transversal projection of the location and momentum vectors

$$\vec{x}_{\text{trans}} \cdot \vec{p}_{\text{trans}} = (x, y) \cdot (p_x, p_y) = xp_x + yp_y. \quad (5.1)$$

A value greater than zero corresponds to a particle that propagates away from the central axis, while a value less than zero indicates a particle moving inwards, i.e. a recurler. In figure 5.8, particles originating from the stopping target and moving inwards and outwards are labelled as “Target, in” and “Target, out”, respectively.




Figure 5.7: Simulated spectral fluence through the scintillating fibre SiPMs upstream and downstream for $2.5 \cdot 10^{15}$ muon stops on target. Shown is the fluence of all electrons and positrons regardless of the decay type.

The simulated spectral fluence of Michel positrons is shown in figure 5.8. The spectrum of incoming Michel positrons peaks at approximately 18 MeV, This can be explained by the kinematics of particles reaching the SiPMs. A high-energetic positron from the target region needs a particular direction to hit the relatively small SiPM while recurling in the magnetic field. The angular range in which positrons of approximately 18 MeV can reach the SiPM is much larger than for other energies, as

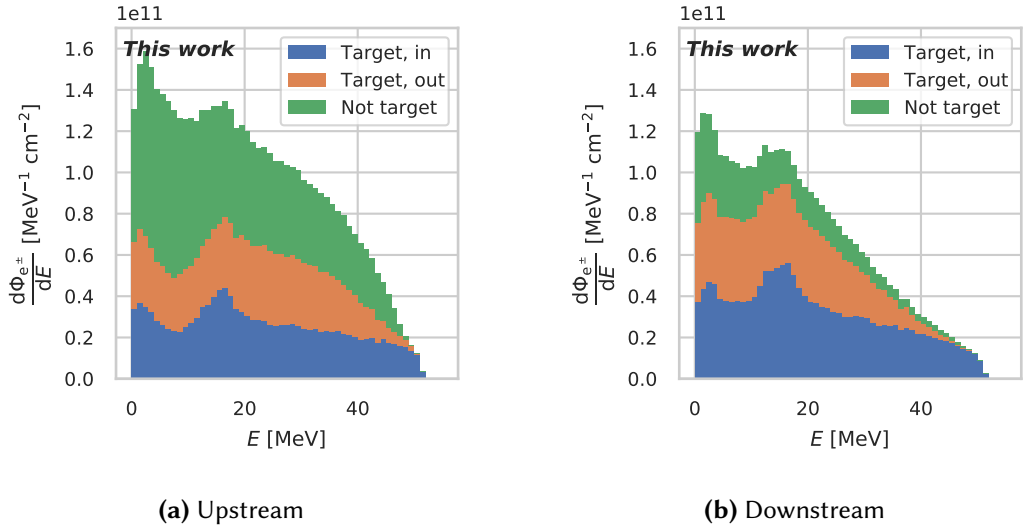


Figure 5.8: Simulated spectral fluence of Michel positrons through the scintillating fibre SiPMs for $2.5 \cdot 10^{15}$ muon stops on target, the Mu3e Phase I goal. The fluence is subdivided in positrons originating from the target region (target and first pixel layer) that are either approaching the z axis (Target, in) or moving away from it (Target, out). The third category of positrons stems from muons decaying in places away from the target (Not target). Michel positrons make up the largest share of the total positron fluence shown in figure 5.7

seen in figure 5.9. The differential decay width $\frac{d^2\Gamma}{dx d\cos\theta_e} \propto x^2 [(3 - 2x) + \cos\theta_e(2x - 1)]$ predicts a favoured decay in the upstream direction (see also equation 1.1). Here, the angle θ_e is defined by the projection of the particle momentum on the z -axis as $\cos\theta_e = p_z/p$. For $x < 0.5$, meaning Michel decays in which the positron carries energy less than a fourth of the muon rest mass, the decay in the downstream direction is preferred. This can be seen in figure 5.8, where the differential flux of positrons from the target region through the downstream SiPMs is larger than through the upstream sensors.

Figure 5.10 shows the conversion of the Michel positron fluence into neutron-equivalent fluence. Almost half of the radiation damage in the upstream SiPMs is caused by positrons from muons not decaying on target. Most of these muons do so in the narrowing of the aluminium beam pipe (see figure 5.6). Detailed values are presented in table B.2 in the appendix. The high number of positrons from the narrowing implies that better beam collimation further upstream reduces the radiation damage in the SciFi SiPMs significantly.

Figure 5.11 shows the neutron-equivalent fluence due to electrons and positrons from all relevant production mechanisms. Michel positrons account for the largest share of displacement damage in the SiPMs, followed by positrons from radiative muon decays. The remaining processes account for less than 10 % of the total neutron-

equivalent fluence: Electrons from Bhabha and Compton scattering in detector material and electrons from photon conversion. Electrons and positrons from the radiative muon decay with internal conversion $\mu^+ \rightarrow e^+e^+e^- \bar{\nu}_{P\mu} \nu_e$ account for about 10^{-4} of the total damage and are therefore omitted in figure 5.11. The detailed numbers can be found in table B.3 in the appendix.




Figure 5.9: Kinetic energy of incoming (per equation 5.1) Michel positrons hitting a SiPM vs their azimuth angle after the muon decay. The range of angles under which a SiPM is hit is the largest for positrons with $E \approx 18$ MeV, thus explaining the peak in figure 5.8




Figure 5.10: Cumulative distribution of the neutron-equivalent fluence due to Michel positrons by decay location. Positrons from the target region are split into incoming and outgoing particles as defined in equation 5.1. Most muons decaying outside the target region decay in the beampipe (see also figure 5.6).

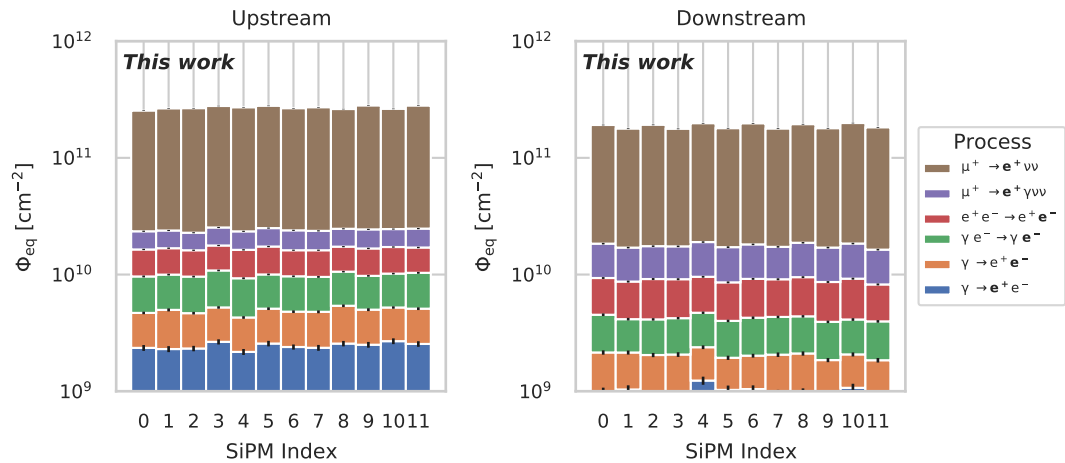


Figure 5.11: Cumulative distribution of the neutron-equivalent fluence per SiPM broken down by particle type and origin. Positrons from Michel decays, and radiative muon decays are responsible for most of the deposited non-ionizing energy, which is converted to neutron-equivalents using [98].

Finally, adding all decay types, decay locations and trajectories, the total neutron-equivalent fluence adds up to approximately $2.7 \cdot 10^{11} \text{ cm}^{-2}$ in the upstream SiPMs and $1.9 \cdot 10^{11} \text{ cm}^{-2}$ in the downstream SiPMs. This corresponds to approximately 20 % of the expected neutron-equivalent fluence through the identical SiPMs in the LHCb experiment [57]. The per-SiPM fluence is shown in figure 5.12, and the individual values are listed in table B.1 in the appendix. Note the staggering of odd and even SiPM numbers, which is explained by the geometrical staggering in z as explained in section 2.1.4.

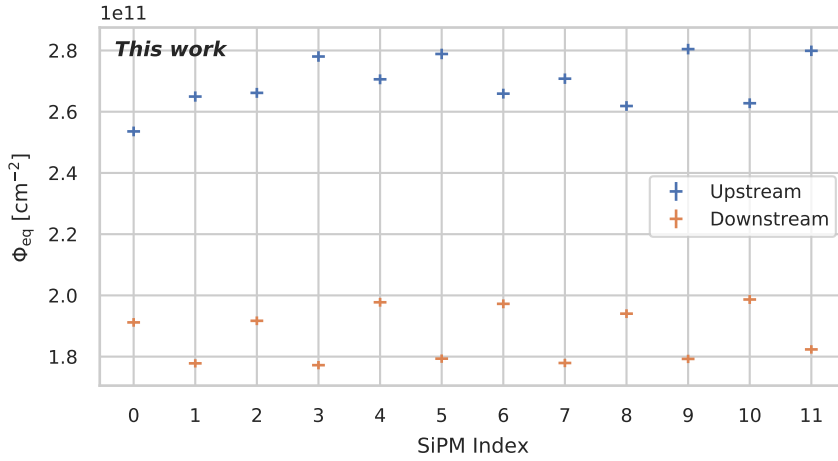



Figure 5.12: Simulated neutron-equivalent fluence per SiPM array inside the Mu3e scintillating fibre detector. The fluence corresponds to $2.5 \cdot 10^{15}$ muon stops on target and is calculated from positron and electron flux under the NIEL hypothesis using [98]. The SiPM index corresponds to the number in figure 2.5.

TOTAL IONIZING DOSE


The Total Ionizing Dose (TID) due to ionizing energy deposition is another potential source of radiation damage. The SiPMs are simulated as solid silicon with a 200 nm thick layer of silicon oxide at the surface. As the electron stopping power of silicon and silicon dioxide are equal within 5 % [99], the bulk silicon and the silicon oxide doses are expected to be approximately equal. In both parts of the SiPMs, the expected TID is in the order of 1 kGy, as shown in figure 5.13. The TID follows the same staggering pattern as the neutron-equivalent fluence caused by the detector geometry.

The values are calculated from the mean energy deposition and the number of simulated interactions. The large error bars in figure 5.13b result from the limited Monte Carlo statistics in the simulated thin silicon oxide layer. As the uncertainties are obtained from the simulated values themselves, they depend on the number of

interactions and the distribution of the energy deposition. In theory, the deposited dose per crossing positron follows a Landau distribution with an undefined mean and variance. Since the tail of the Landau distribution is unphysical (a particle cannot deposit more energy than its own), the average would eventually converge [100].



(a) TID in the bulk silicon



(b) TID in the silicon dioxide. The silicon dioxide is simulated with a thickness of 200 nm, leading to a high variance of the deposited dose per crossing positron. This variance is reflected in the error bars calculated per SiPM array, leading to large error bars in some arrays.

Figure 5.13: Simulated total ionizing dose (TID) in each SiPM array for $2.5 \cdot 10^{15}$ muon stops on target. A silicon dioxide layer of 200 nm is simulated on top of the 300 μm thick bulk silicon. The SiPM index corresponds to the number in figure 2.5.

5.3.2 RADIATION IN PIE5




Figure 5.14: Visualization of the radiation setup and simulated particle trajectories during 64 ns. Muons pass through layers of polyimide foil before they are stopped in a polyethylene target and decay. The spectral fluence of the decay positrons as well as secondary particles through the SiPMs and the tile matrix is counted and used to estimate the equivalent fluence of 1 MeV neutrons using [98]. This figure shows the simulated decay of 315 muons generated 1 cm away from the target (green lines in the centre of the figure).

Based on the Mu3e simulation framework, the irradiation setup described in section 5.2 was simulated. A visualization of this simulation, including the setup and decaying muons, is shown in figure 5.14. Implemented in the simulation is the geometry of the stopping target, cooling block, SiPMs, scintillating tile matrix and polyimide sheets (see figure 5.14). As the tile matrix and polyimide sheets were irradiated parasitically, this section focuses on the dose and fluence in the two scintillating fibre SiPMs.

The dose calculation is performed in multiple steps, shown in figure 5.15. The following section provides a detailed explanation of the individual steps and quantities.

5.3.3 NORMALIZATION TO THE PROTON CURRENT

Secondary particle production at PSI secondary beamlines is proportional to the *proton beam current* I_p . As the current is not constant, any beam characterization must be




Figure 5.15: With a pill scan, the beam shape and total muon rate $R_{2.2\text{mA}}$, normalized to the HIPA proton beam current are determined. These results are used in a Geant4 simulation to determine the spectral flux in the irradiated samples, normalized to the proton beam current.

normalized. A measurement of this current is available via Experimental Physics and Industrial Control System (EPICS) [97]. The current monitor MHC4 at the HIPA facility gives the best available measurement of the proton beam on Target E. The same normalization is necessary when comparing real-world irradiation to simulations. By convention, quantities are normalized to $I_p = 2.2\text{ mA}$.

For simulation, the beam is characterized by the normalized rate of muons $R_{2.2}$ and the shape of the beam spot in the stopping target. This characterization is performed using a beam scanner consisting of a cylindrical scintillator with a cross-section of 3.113 mm and a photomultiplier tube. In the first step, the beamline settings—mainly currents in electromagnets—are optimized iteratively to maximize the muon flux in the centre of the setup. The highest flux corresponds to surface muons with a momentum of approximately 28 MeV. Afterwards, the pill scanner is moved across the transverse plane of the beam to determine the beam spot shape.

The data from the pill scan with the final settings are shown in figure 5.16, reflecting the muon beam distribution. A bivariate Gaussian model of the form

$$\dot{\Phi}_{P\mu}(x, y) = \frac{R_{2.2}}{\sqrt{(2\pi)^2 \sigma_x \sigma_y}} \cdot \exp \left[-\frac{1}{2} \left(\frac{(x - \mu_x)^2}{\sigma_x^2} + \frac{(y - \mu_y)^2}{\sigma_y^2} \right) \right] \quad (5.2)$$

is fit to the measured muon flux $\dot{\Phi}_{P\mu}$ in the pill.

The rates in the scanner are normalized to the proton current I_p . With R being the




Figure 5.16: The muon beam spot in PiE5. The measured muon rate and 1σ , 2σ , 3σ contour lines of a bivariate Gaussian distribution fitted to data are shown. Note that the coordinates refer to the internal coordinates of the pill scanner. The origin of the irradiation setup lies at $x = 7.4$ mm, $y = 0$.

total muon rate, the normalized rate is

$$R_{2.2} = R \cdot \frac{2.2 \text{ mA}}{I_p}. \quad (5.3)$$

The best fit yields the parameters for equation 5.2 presented in table 5.1.

Table 5.1: Fit results for the beam spot measurements

$R_{2.2}$	$(2.346 \pm 0.037) \cdot 10^8 \text{ s}^{-1}$
σ_x	$(16.45 \pm 0.40) \text{ mm}$
σ_y	$(21.61 \pm 0.47) \text{ mm}$
μ_x	$(2.78 \pm 0.43) \text{ mm}$
μ_y	$(-1.19 \pm 0.51) \text{ mm}$

The beam characteristics are needed as input to the Geant4 simulation to determine the flux through the SiPMs in the irradiation chamber. The total neutron-equivalent fluence can be obtained by the time integral of the flux over all times while the shutter is open:

$$\Phi_{\text{eq}} = \int \dot{\Phi}_{\text{eq}} dt \propto \int_{\text{shutter}} I_p dt. \quad (5.4)$$

The constant of proportionality depends on the measured beam spot and the setup geometry. The irradiation chamber is simulated in Geant4 to obtain this constant that connects the neutron-equivalent fluence to the integrated beam current.

As the fluence is proportional to the rate of muons, equation 5.4 can be rewritten as

$$\Phi_{\text{eq}} = \int_{\text{shutter}} R_{2.2} \frac{I_p(t)}{2.2 \text{ mA}} g dt. \quad (5.5)$$

so that all unknowns are embedded in a single factor g of the same dimension as Φ_{eq} . The simulation yields the spectral positron and electron fluence $d\Phi_{e^\pm}/dE$ in both SiPM, the Aare and Berg side (refer to figure 5.1a). Figure 5.17 shows the flux, i.e. the fluence per unit time of both sides as a function of energy E .

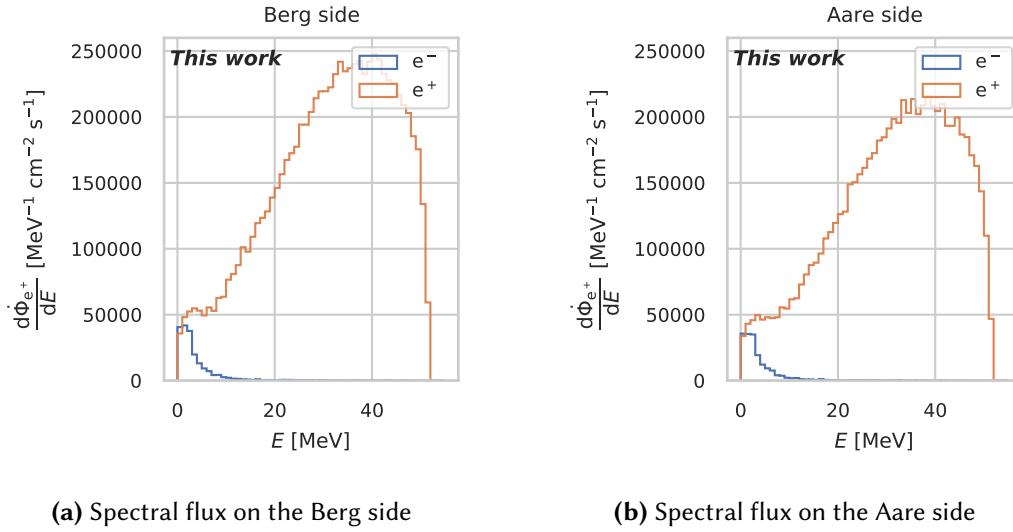


Figure 5.17: Simulated spectral flux in both positions of the irradiation chamber. This flux is normalized to $I_p = 2.2 \text{ mA}$.

The neutron-equivalent fluence is obtained using the hardness factor $\kappa(E)$, which is the ratio of the damage function for electrons $D_{e^\pm}(E)$ and 1 MeV neutrons $D_n(1 \text{ MeV}) = 95 \text{ MeV mb}$:

$$\Phi_{\text{eq}} = \int \kappa(E) \frac{d\Phi_{e^\pm}}{dE} dE = \int \frac{D_{e^\pm}(E)}{D_n(1 \text{ MeV})} \frac{d\Phi_{e^\pm}}{dE} dE \quad (5.6)$$

These damage functions are discussed in section 3.4 and reflect the amount of non-ionizing energy deposited in matter.

Using the measured beam spot (see table 5.1) and the simulated setup, the neutron-equivalent flux in the two slots per integrated proton charge is obtained:

$$\frac{\Phi_{\text{eq}}}{\int I_p dt} = \begin{cases} (6.550 \pm 0.017) \cdot 10^8 \text{ cm}^{-2} \text{ mAh}^{-1}, & \text{on the Aare side,} \\ (7.633 \pm 0.018) \cdot 10^8 \text{ cm}^{-2} \text{ mAh}^{-1}, & \text{on the Berg side.} \end{cases} \quad (5.7)$$

The quoted uncertainties are only the statistical uncertainties from the simulation.

5.3.4 SYSTEMATIC UNCERTAINTIES

The explained technique is affected by several sources of systematic uncertainties. In the following, the most critical uncertainties are discussed.

POLARIZATION

Due to the weak interaction's chiral structure, a surface muon's spin is antiparallel to its momentum (see section 1.4). Therefore, higher-energetic positrons are more likely to be emitted in the upstream direction than downstream (see also equation 1.1). The muon polarization at PiE5 is about 86 %, as measured in the MEG experiment at the same beam line [101]. The main reasons for the partial depolarization are contamination with cloud muons and beam divergence. The depolarization in the irradiation setup is not measured and could differ from the one measured in MEG due to the target material and difference in magnetic fields.

The nominal polarization in the simulations of this chapter is 86 %. Simulations with an initial muon polarization between 0 and 100 % are compared to estimate the effect on the neutron-equivalent flux (see figure 5.18). For a polarization of 70 %, the neutron-equivalent flux is about 2.6 % lower than for a completely polarized beam. A systematic uncertainty of 1.3 % is assumed due to the unknown muon (de-)polarization.




Figure 5.18: Simulated effect of the muon polarization on the neutron-equivalent flux at the nominal beam current of $I_p = 2.2$ mA, as a function of the muon polarization.

INTEGRATION LIMITS AND NUMERICAL INTEGRATION

As explained above, the total fluence depends on the measured beam current and beam shutter status. Continuous logging of the beam shutter status allows for precise

integration limits. In detail, EPICS reports the beam shutter status S as 1 for an open shutter and 0 for a closed shutter. The only unknown is the exact movement of the beam shutter. During the time of movement, the beam shutter variable is reported as 0.5, posing an error source. Compared to the times of irradiation in the order of days, the movement of the beam happens within seconds. The potential error is therefore neglected.

Another potential error source is the numerical integration of the discrete measurements of I_p . The effective sampling rate of 6 min^{-1} is sufficient to detect short beam interruptions with subsequent resumption (see figure 5.19a). Therefore, the systematic uncertainty of numerical integration is also neglected.

STABILITY OF BEAMLINE PARAMETERS

Almost all beamline settings (see section 2.2.2) are controlled via a program called “Set Point”. One important exception is the high voltage for the separator SEP41, controlled separately. After initial problems with the stability of the high voltage and large leakage currents were resolved, the voltage was stable.

Since the potential difference between the anode and cathode of the separator is about 360 kV, leakage current produces hard x-rays through bremsstrahlung. Even after improving the stability, the radiation levels inside the beam area were too high for human safety due to these x-rays. Therefore, to access the area, the high voltage needed to be manually reduced and reset after the access. During the irradiation of SiPMs 1 and 2, a high voltage lower than the nominal voltage was found, possibly by human error. The lower voltage caused improper beam alignment and thus a reduction of the dose rate for SiPMs 1 and 2. After the incident, the high voltage was monitored closely, and even though it could not be logged automatically, no irregularities were observed. It can be assumed that the large systematic uncertainty only affects the first two samples.

The uncertainty due to beamline parameter drift in the irradiation of SiPMs 3 to 6 is negligible, while the effect on SiPMs 1 and 2 is discussed in detail in section 5.4.1.

BEAM MEASUREMENT AND MUON-STOPPING DISTRIBUTION

The pill scanner is aligned to fixed points at the beamline using alignment lasers. This alignment has a precision of 1 mm, which directly translates to a systematic uncertainty of the central beam position (μ_x, μ_y) .

In the lateral (z) direction, the stopping target extends from 16 to 20 mm, measured from the inner surface of the lid. Due to spatial and safety constraints, the muon beam was measured 38 mm from the target. The simple beam scanner used to measure the beam spot is only sensitive to the spread in the beam position, not the divergence, i.e. the angular spread.

The measurements from table 2.2 in [50] allow an estimation of the uncertainty due to the z offset and beam divergence. At the same position in the beamline, downstream of the collimator system, the values in table 5.2 were obtained.

Table 5.2: “Summary of beam profile parameters at a standard pill measurement position DS [downstream] the SML/collimator system [...] for the G4BL simulation starting from TgE.” [50] The quoted values are used to estimate the uncertainty of the beam spot size due to the distance between the stopping target and where the beam spot was measured

	x	y
σ (mm)	34.0	26.5
ρ	0.89	0.71
ϵ_{rms} (mm · mrad)	935	421

Table 5.2 contains the beam size σ in the horizontal and vertical dimensions, x and y , respectively.

The emittance $\epsilon_{x/y}$ and the correlation coefficient between the spatial and angular coordinates $\rho_{x,\theta/y,\phi}$ are used to obtain the angular spread θ/ϕ as explained in [50]²:

$$\epsilon_x = \sigma_x \cdot \theta \cdot \cos(\arcsin(\rho_{x,\theta})) \quad (5.8)$$

From this, the angular spreads are $\theta = 56$ mrad and $\phi = 22$ mrad.

During the irradiation, the beam was more focussed, which (if everything else were to remain the same) would increase the angular spread by the ratio of the widths σ [102]. In the x component, the increased angular spread is 116 mrad and in y 27 mrad. The distance of about 20 mm between the beam measurement and the stopping target translates to a lateral displacement of 2.3 mm and 0.54 mm, respectively, which are treated as a systematic uncertainty.

The fit uncertainties in table 5.1 are also considered systematic uncertainties. The total uncertainties of those parameters are summarized in table 5.3.

To determine the effect of these uncertainties on the neutron-equivalent flux $\dot{\Phi}_e$, the setup was not only simulated with the nominal parameters in table 5.1 but also with modified parameters. Each parameter was varied in steps of 1 mm, and the flux of this varied configuration is compared to the “nominal” fluence $\dot{\Phi}_{\text{eq}}^{\text{nom}}$ as shown in figure 5.20. The central position (μ_x, μ_y) and the shape (σ_x, σ_y) are varied independently.

Variations of μ_y have little influence on the flux due to the shape and orientation of the SiPM shown in figure 5.1. As the box has its centre at $x = 7.4$ mm, the “Aare” and “Berg” SiPMs are at 15.2 mm and -0.4 mm, respectively. The nominal centre of the beam is, therefore, better aligned with the SiPM on the “Berg” side. The change in flux scales with the gradient of the beam shape: Where the beam has its maximum, the flux does

²On page 47, x_m is used to denote the beam size; in table 22, the author uses σ_x

Table 5.3: Systematic uncertainties of the beam parameters and their effect on the neutron-equivalent flux

Parameter	syst. uncertainty [mm]	effect on $\Phi_{\text{eq}}^{\text{Berg}}$	effect on $\Phi_{\text{eq}}^{\text{Aare}}$
μ_x	1.1	< 1 %	0.3 %
μ_y	1.1	< 1 %	< 1 %
σ_x	2.3	9.5 %	4.6 %
σ_y	0.7	2 %	2 %
combined	—	9.7 %	5 %

not change significantly (figure 5.20a). Under the falling edge of the distribution, the flux in the “Aare” SiPM has a stronger dependence on the beam position (figure 5.20b) of about 3 % per millimetre of beam displacement.

The more focussed the beam is, i.e. the smaller σ_x and σ_y are, the higher the flux in both SiPMs, as shown in figures 5.20c and 5.20d. If the beam is broadened in one dimension and narrowed in the other, the change in flux is tiny. This explains the diagonal shape in the figures. An error of 2 mm in either parameter would induce an error in flux between 5 and 10 %.

BEAM RATE

The pill scan is based on counting the number of events in a scintillator and is thus a statistical process. Close to the beam centre, the rate inside the pill was in the order of 10^5 s^{-1} . As the used scaler counts in time windows of 100 ms, the number of events per data point is in the order of 10^4 . That rate has to be scaled using the measured proton beam current and the size of the scintillator to convert this rate into what is shown in figure 5.16. The number of hits in a given time window follows a Poissonian distribution, so a statistical uncertainty of $1/\sqrt{10^4} = 1\%$ has to be assumed per data point. As these measurements’ uncertainties are unknown, a total uncertainty of 1 % is assumed.

The total rate of the beam spot is determined via a fit that yields a fit uncertainty of 1.6 %. That fit relies on the assumption of a multivariate Gaussian distribution with no correlation between the x and y components. As shown in chapter 4.5 of [50], the two components are only very weakly coupled in PiE5. When one compares the total beam rate determined with a simple multivariate Gaussian fit to a model-independent measurement, a “most often $\lesssim 3\%$ higher beam rate” is obtained, possibly “due to non-Gaussian tails in the distribution” [50]. An additional systematic uncertainty of 3 % accounts for a model mismatch of the chosen distribution. The combined uncertainty of the beam rate is taken as 3.5 %.

GEOMETRY AND POSITION

The irradiation setup was produced with great precision and coupled to the beamline using a flange leaving virtually no space for misalignment. This geometry is translated into the simulation accordingly, and tolerances are deemed to have a negligible effect on the total dose compared to other uncertainties, such as the beam spot position.

SIMULATION EFFECTS AND SCALING

To quantify the damage over time, a Geant4-based simulation was used. The limitations of quantifying radiation using neutron-equivalent fluence discussed in section 3.4.1 mainly concern the comparison to other particles. As shown in figure 5.17b, the spectral flux during the irradiation differs from the one expected in Mu3e. The main reasons for the difference are the different geometry and the absence of a magnetic field. Nevertheless, converting positron fluence into neutron-equivalent fluence accounts for the difference in NIEL deposited by positrons of different energies.

TOTAL SYSTEMATIC UNCERTAINTY

The total systematic uncertainty of the neutron-equivalent flux is presented in table 5.4. The same uncertainties carry over to the fluence as the total fluence is proportional to the flux.

Table 5.4: Systematic uncertainty of neutron-equivalent flux. The uncertainties are treated as uncorrelated, so the total uncertainties for each side are obtained by summation in quadrature.

Source	$\Delta\dot{\Phi}_{eq}^{Aare}$	$\Delta\dot{\Phi}_{eq}^{Berg}$
Polarization	1.3 %	1.3 %
Numerical integration	-	-
Beamline parameters	Only SiPM 1	Only SiPM 2
Beam position/shape	9.7 %	5 %
Beam rate	3.5 %	3.5 %
Geometry	-	-
Total	10.4 %	6.2 %

5.4 FLUENCE PER SAMPLE

As explained above, the neutron-equivalent fluence is calculated in several steps. The log file containing the status of the beam shutter KSF41 and data from the proton

current monitor MHC4 is used to integrate the proton current I_p over the corresponding time. Applying the previously discussed calculations to each SiPM sample yields the neutron-equivalent fluence for each sample. The results are summarized in table 5.5.

The positron flux for both positions is slightly different. Therefore the total fluence depends on whether the SiPM was irradiated in the “Aare” or “Berg” slot.

Table 5.5: Integrated proton current for all SiPMs. The *side* refers to the two positions inside the setup, *Aare* and *Berg* are landmark references at PSI. Together with the results in equation 5.7, the integrated proton current is converted to the neutron-equivalent dose. During the irradiation of SiPMs 1 and 2, the stated range consists of the upper limit and an estimate of the lower limit. SiPM 7 is an unirradiated control sample. In addition to the six samples, a populated PCB missing the SiPM was irradiated to observe possible effects of the irradiation of the board itself.

SiPM	Side	$\int I_p dt$ [mAh]	Φ_{eq} [10^{10} cm^{-2}]
1	Aare	263.2	11.6 ± 1.2 to 17.2 ± 1.8
2	Berg	263.2	13.50 ± 0.84 to 20.0 ± 1.2
3	Aare	88.9	5.82 ± 0.61
4	Berg	27.6	2.11 ± 0.13
5	Berg	61.3	4.68 ± 0.29
6	Berg	98.2	7.49 ± 0.46
7	control	0.0	0.0
blank	Aare	98.2	-

5.4.1 SEPARATOR INCIDENT

During the beamtime, the separator SEP41 suffered from a relatively large leakage current while operating at the nominal voltage of ± 180 kV. As the electric field inside a Wien filter is a plate capacitor, a leakage current means that electrons are accelerated by a potential difference of 360 kV before hitting the anode. With a kinetic energy of 360 keV, these electrons produce hard X-rays when they are stopped, scaling with leakage current. The high voltage had to be reduced to ± 150 kV to keep the radiation levels near the separator at safe levels every time someone accessed the area. As explained in section 5.3.4, this voltage was the only beamline setting not controlled by EPICS and, therefore, not logged.

After one access during the irradiation of the first two SiPMs, 1 and 2, the high voltage was not reset before the irradiation was resumed. This mistake was only noticed on the morning of the second day after the last access.

A reduced voltage in the separator leads to a stronger deflection of the positrons that are separated, but the muons are also deflected in the same direction. Therefore,

no additional positrons were directed on the SiPMs, but the muon beam partially intersected with the beamline, reducing the positron flux in the irradiation setup.

Under the assumption that the voltage was only too low between the last access and the time of discovery, a lower limit can be determined by treating this time as if *no irradiation* had taken place. With the normal calculation as an upper limit, the effective integrated current for these two SiPMs lies in the following range:

$$177 \text{ mAh} < \int I_p dt < 263 \text{ mAh.} \quad (5.9)$$

Potentially that error did not happen during the previous access, which would reduce the lower limit even further.

This error does not affect the other SiPMs, 3 through 6.

5.5 RADIATION EFFECTS

During the irradiation, the SiPMs were biased and cooled to temperatures between 9 and 13 °C. The current flowing through the SiPMs and the temperature were monitored in intervals of 10 s.

After irradiation, the SiPMs were immediately cooled below -20 °C, and an unbroken cold chain was maintained until the first post-irradiation studies.


5.5.1 CURRENT DURING IRRADIATION

The bias currents of the SiPMs during irradiation are summarized in figure 5.21 and presented separately with the bias voltage for each individual SiPM in appendix B.6. Since the lowpass filter on the carrier board contains a 1 kΩ resistor in series with the SiPM (see figure 5.3), a current of 1 mA leads to a voltage drop of 1 V. The externally applied voltage was adjusted twice to compensate for this voltage drop during the irradiation of SiPMs 1 and 2, explaining the steps in figure 5.21. During the irradiation of the last four SiPMs, no such compensation was performed. The reduced effective bias voltage at higher currents (and thus at higher neutron-equivalent fluence) affects the current through the SiPM. As a result, the increase in the current flattens out.


The Villigen midsummer weather increased the dew point, which caused condensation on the outside of the irradiation chamber. The temperature was increased from 9 °C to approximately 13 °C to protect the electronics from water damage (see figure 5.19c). The higher temperature leads to larger currents during the irradiation of SiPMs 4 through 6.

The currents of all samples exceeded the limit of the ammeter monitoring the dark current during irradiation. Therefore all curves are capped at approximately 2.25 mA.

A control sample consisting of a carrier board populated with all components, excluding the SiPM, was also irradiated. No current through this board was measurable, thus proving that the increase in the current is exclusively due to effects in the SiPM.

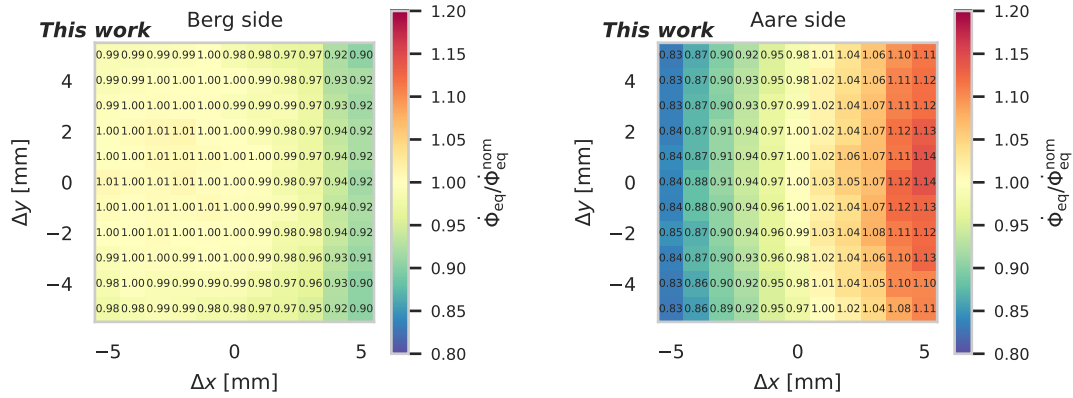


(a) Data points from EPICS. On the left, the proton current of HIPA I_p is shown. A brief interruption and subsequent resumption of the beam operation can be observed. The beam shutter KSF41 was open the entire time, as can be seen on the right (a value of 1 corresponds to an open shutter, whereas 0 would indicate a closed shutter). The relatively high sampling rate minimizes uncertainties from numerical integration.




(b) Proton current (MHC4) and beam shutter status (KFS41) during irradiation (c) The temperature of the SiPMs during irradiation

Figure 5.19: Monitoring variables during irradiation. In the middle of the campaign there was a beam interrupt during which no irradiation took place



(a) Simulated effect of variation of the beam spot centre on the SiPM being positioned at $x < 0$ (b) Simulated effect of variation of the beam spot centre on the SiPM being positioned at $x > 0$



(c) Simulated effect of variation of the beam spot width on the SiPM at $x < 0$ (d) Simulated effect of variation of the beam spot width on the SiPM at $x > 0$

Figure 5.20: Two-dimensional distribution of the change in neutron-equivalent flux $\dot{\Phi}_{eq}$ due to variations of the beam position and width (colour coded). The parameters μ_x , μ_y , σ_x and σ_y are varied by ± 5 mm in steps of 1 mm. The results are used as systematic uncertainties for the dose measurements. (All *This work*)




Figure 5.21: Currents of each SiPM during irradiation. SiPMs 1 and 2 were irradiated at a lower temperature, and the bias voltage was adjusted twice. All other SiPMs were irradiated with a constant bias voltage.

5.5.2 IV CURVES

Before soldering the SiPMs to the carrier boards, and therefore before irradiation, IV curves of every channel of the SiPMs were measured³ at room temperature. After irradiation, the SiPMs were warmed to room temperature, and a second IV curve per channel was measured. During the second set of measurements (after irradiation), the laboratory temperature was up to approximately 3 °C warmer. The temperature differences are shown in figure 5.23. In figure 5.22, the curves for one 128-channel SiPM array are shown. IV curves for all SiPMs can be found in appendix B.3.




Figure 5.22: IV curves for SiPM 6 before (blue) and after irradiation (red). Different curves represent individual channels. The approximate region of the breakdown voltages V_{BD} is marked. The current increases after irradiation above and below breakdown.

Figure 5.22 shows that the current after irradiation is significantly higher, both above and below the breakdown voltage. The change in current at the manufacturer-recommended voltage 3.5 V above breakdown (determined using the tangent method, see section 5.5.3) is determined per channel. The per-channel current ratios of one SiPM are shown in figure 5.24a (see also appendix B). Below, in figure 5.24b, the relative increase in the current below breakdown is shown.

³Courtesy of Y. Demets




Figure 5.23: Temperature differences during IV measurements before and after irradiation for each channel. The measurements before irradiation were conducted at warmer temperatures for almost every individual channel.



(a) Current ratio above breakdown at $V = V_{BD} + 3.5\text{ V}$ (b) Current ratio below breakdown at $V = 50\text{ V} < V_{BD}$

Figure 5.24: Per-channel ratio of the current in SiPM 6 after and before irradiation at voltages above and below breakdown. Each data point resembles the current ratio $I_{\text{after irradiation}}/I_{\text{before irradiation}}$ in an individual channel. (a) shows the operational voltage recommended by the vendor, while (b) shows a fixed voltage below breakdown.

The change in leakage current is not uniform across an SiPM, as can be seen as outliers in figure 5.24b. Some outliers can be explained with instrumentation errors, e.g. the channel that flattens out in figure B.1. As the formation of defects is a statistical process, individual channels can have a particular defect that leads to a high leakage current that bypasses the amplification region. In such a case, the current below breakdown would increase significantly, but that effect would be overshadowed as soon as the breakdown voltage is reached. No further studies of these outliers have been performed.

The distribution of the per-channel current ratios of one SiPM is not well-described by a Gaussian approximation (see figures 5.24b and 5.24a, as well as section B.4 in the appendix). Box plots have been chosen to summarize the information from all SiPMs as a function of the neutron-equivalent dose in figure 5.25. Each box represents the spread of 128 channels of one SiPM. The orange line in the centre marks the median current ratio, while the lower and upper edges mark the first and third quartiles. Each box contains whiskers spanning from the 5th to the 95th percentile to show the spread beyond the quartiles. Studying percentiles rather than the mean and standard deviation mitigates the impact of outliers and allows for better a description of skewed distributions. The width of the box shows the error of the neutron-equivalent dose.


The variability of a skewed distribution can be described by the interquartile range (IQR), the difference between the 25th and 75th percentile. A normal distribution with a standard deviation σ has an IQR of approximately 1.35σ . A linear model is fit to the median current ratios for SiPMs 3 to 7 with weights of $(\text{IQR}/1.35)^{-2}$. SiPMs 1 and 2 are excluded from the fit due to the uncertainty of the fluence (see section 5.4.1). Fitting a linear function assumes a linear increase of damage with NIEL, i.e. NIEL scaling.

At a bias voltage of 3.5 V above V_{BD} , the results shown in figure 5.25a show a linear dependency. The linear fit yields an increase of current by a factor of 2000 ± 250 for a neutron-equivalent fluence of $1 \cdot 10^{11} \text{ cm}^{-2}$.


For currents below the breakdown voltage, the distribution of the current ratios per SiPM differs vastly, as shown in figure 5.25b. The increase is determined by the fit is 21.8 ± 8.9 per neutron-equivalent fluence of $1 \cdot 10^{11} \text{ cm}^{-2}$, significantly smaller than the increase above breakdown.

5.5.3 BREAKDOWN VOLTAGE

SiPMs are operated with a reverse bias larger than their breakdown voltage V_{BD} . The difference between the operating voltage and V_{BD} has a large influence on the SiPM performance. A change in V_{BD} due to radiation would therefore require an adjustment of the operating voltage. The power supply providing the bias voltage can supply voltages more than twice the typical operation voltage before irradiation. In the following, a possible change is studied, so that proper action can be taken if the voltage needs adjustment.



(a) Current ratios above breakdown, $V = V_{BD} + 3.5 \text{ V}$



(b) Current ratios below breakdown, $V = 50 \text{ V} < V_{BD}$

Figure 5.25: Change in current in relation to the neutron-equivalent fluence above (a) and below (b) breakdown. Each box describes the distribution of the per-channel current ratios after and before irradiation of one SiPM. The central line marks the median ratio, and the box edges mark the first and third quartile. The width represents the total uncertainty of the neutron-equivalent fluence in the corresponding SiPM. A linear model is fit to the median values of SiPMs 3 to 7 (an unirradiated control sample).

IV curves (e.g. figure 5.22) allow the extraction of the breakdown voltage of a SiPM. Many textbooks only provide a qualitative definition of the reverse breakdown voltage of a diode V_{BD} as the voltage above which it conducts a significant amount of current (e.g. [103]). Two different methods of quantitative determination of the breakdown voltage described in the literature [104] are presented in the following. The selection is based on the resolution of the data available. Based on each method, the breakdown voltages before and after irradiation are compared.

RELATIVE DERIVATIVE METHOD

The first method is based on the $\log(I(V))$ curve, which has an inflexion point at $V = V_{peak}$ in the notation of Hamamatsu [65]. At this inflexion point, the derivative of the logarithm

$$\frac{d}{dV} \log(I(V)) = \frac{1}{I} \frac{dI}{dV} \quad (5.10)$$

has a local maximum. The difference $V_{peak} - V_{BD}$ is independent of V_{BD} and only depends on the type of SiPM [65]. Under this hypothesis, a change in V_{peak} would indicate a change in V_{BD} .

As the current is measured in equidistant voltage steps ΔV , the derivative is calculated numerically for every step n using the symmetric difference quotient

$$\left[\frac{dI}{dV} \right]_n = \frac{I_{n+1} - I_{n-1}}{2 \Delta V}. \quad (5.11)$$


As a heuristic, a Gaussian function is fitted to the data around the maximum to find V_{peak} . An example is shown in figure 5.26.

TANGENT METHOD


The second method is also based on $\log(I)$. In this case, a linear function is fit to the steep rise in $\log(I)(V)$. The breakdown voltage is then defined as the interception point of the extrapolation with the baseline current below breakdown (see figure 5.27). For the unirradiated SiPMs, this baseline constant can be approximated by another linear function. After irradiation, the baseline has a more complex shape. A second-order polynomial is fitted to the data to describe the baseline after irradiation (see figure 5.27b).

TEMPERATURE CORRECTION

The breakdown voltage depends linearly on the temperature with a coefficient of approximately 54 mV/K [42]. During the recording of the waveforms, the temperature was logged (see figure 5.23). For the sake of comparability, the values presented here




(a) Before irradiation




(b) After irradiation

Figure 5.26: Determination of the breakdown voltage of one SiPM channel using the relative derivative $\frac{d}{dV} \log(I(v)) = I^{-1} \frac{dI}{dV}$. The voltage at the maximum V_{peak} has an offset with respect to V_{BD} that only depends on the type of SiPM. Therefore, a change in V_{peak} is equivalent to a change in V_{BD} .



(a) Before irradiation



(b) After irradiation

Figure 5.27: Example of how the tangent method is used to determine the breakdown voltage. The current I , a second-order polynomial fit of the baseline and the tangent are shown. The vertical line V_{BD} marks the intersection of the baseline and the tangent.

are corrected to what $V_{\text{BD}}^{\text{measured}}$, obtained at a temperature T , would correspond to at $T_0 = 25^\circ\text{C}$:

$$V_{\text{BD}} = V_{\text{BD}}^{\text{measured}} - 54 \text{ mV/K} \cdot (T - T_0). \quad (5.12)$$

RESULTS

For each SiPM array, V_{BD} is determined for all 128 channels before and after irradiation. Figure 5.28 shows the breakdown voltages of all channels for a single SiPM array determined with the tangent method. The difference

$$\Delta V_{\text{BD}} = V_{\text{BD}}^{\text{after irradiation}} - V_{\text{BD}}^{\text{before irradiation}} \quad (5.13)$$

is calculated for every single channel (see the bottom part of figure 5.28).

The per-channel V_{BD} and ΔV_{BD} can be found in appendix B.5 for all SiPMs and using both methods. In the following, the different SiPMs are compared. As the distribution of the 128 per-channel ΔV_{BD} is not necessarily Gaussian, the analysis is based on percentiles. That way, no assumption about the distribution is made. Figure 5.29 shows a boxplot of the V_{BD} for each SiPM as a function of the neutron-equivalent fluence Φ_{eq} of the respective SiPM. The orange line in the middle signifies the median ΔV_{BD} , and the vertical boundaries are the first and third quartile, thus covering the central 50 % of the values. The whiskers mark the interval that contains the central 90 % of the values. In the horizontal, the position of each box marks the neutron-equivalent fluence and the width of its uncertainty.




Figure 5.28: Temperature-corrected breakdown voltage for each channel of SiPM 6 before and after irradiation. The bottom shows the change after irradiation.

The median ΔV_{BD} is negative for all SiPMs. The magnitude is below 500 mV for all SiPMs and, therefore, in the same order of magnitude as the typical variation of V_{BD}

within one 128-channel SiPM array [42, 91] (also visible in appendix B.5). Unexpectedly, the breakdown voltage of the unirradiated control sample SiPM 7 is also reduced. Possible causes are discussed in section 5.7.

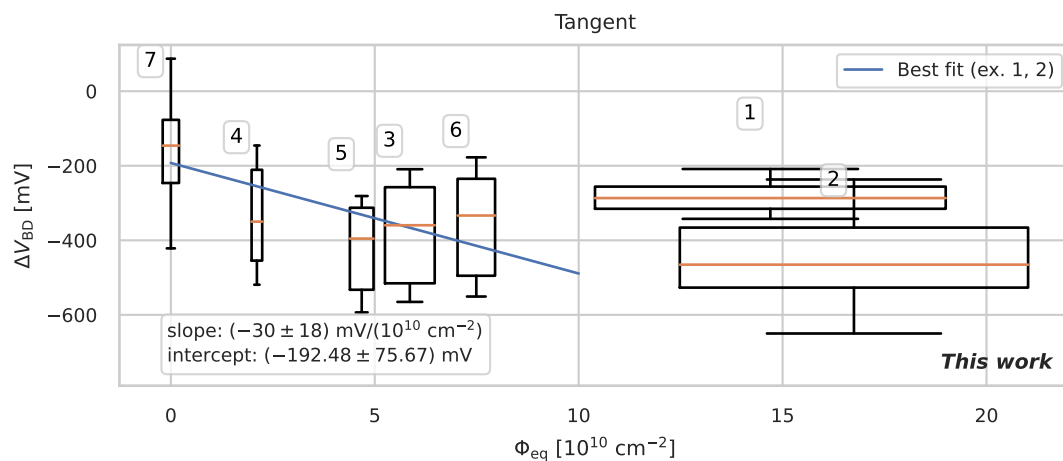
For the following steps, two assumptions are made. First, it is assumed that the reason for $\Delta V_{\text{BD}} \neq 0$ in SiPM 7 affects all SiPMs equally. The second assumption follows from the NIEL hypothesis that changes are proportional to the deposited NIEL, i.e. scale linearly with Φ_{eq} .

Under these assumptions, a linear function is fit to the median values of SiPMs 3 through 7, excluding the two SiPMs with the previously discussed large fluence uncertainty. The weights for each point are $(\text{IQR}/1.35)^{-2}$ (see above). For ΔV_{BD} obtained with the tangent method, the fit yields a dependence of


$$dV_{\text{BD}}/d\Phi_{\text{eq}} = (-30 \pm 18) \text{ mV}/(10^{10} \text{ cm}^{-2}). \quad (5.14)$$

The derivative method yields $dV_{\text{BD}}/d\Phi_{\text{eq}} = (-22 \pm 17) \text{ mV}/(10^{10} \text{ cm}^{-2})$.

With those assumptions and uncertainties of 60 % to 77 %, these values can hardly be regarded as significant. Even though the measurement is *inconclusive* regarding whether ΔV_{BD} decreases with irradiation, the result shows that if there is a decrease, it is not larger than the typical variation between channels of one Hamamatsu S13552 SiPM column array.



(a) Using the tangent method



(b) Using the relative derivative method

Figure 5.29: Box plots of the change in breakdown voltage, as determined by two different methods. The median of the difference for each SiPM is indicated by an orange line, surrounded by a box marking the first and third quartile. The whiskers cover the central 90 % of the data points. For all SiPMs except number 7, the width of the box shows the uncertainty of the neutron-equivalent fluence (SiPM 7 is not irradiated and has therefore zero fluence). A linear function is fitted to SiPMs 3–7

5.5.4 WAVEFORMS

Waveforms, i.e. the signal output as a function of time, allow the evaluation of SiPMs performance. The irradiated SiPMs 1, 6, and 7 were covered to shield from light, biased with 55 V, and read out with a DPNC 342 readout board. That way, a large number of waveforms was recorded for each channel, each with a length of 200 ns. No trigger was used, so the time when a waveform was recorded was independent of the signal itself. The current of a signal is converted into a voltage and amplified for easier processing. Therefore, a single-photoelectron signal has an amplitude of approximately 100 mV.

Figure 5.30 shows an overlay of 100 such waveforms of a representative channel for each of the three SiPMs. Darker areas in the figure represent combinations of time and voltage that occurred in more waveforms. Most waveforms do not show any signal for the unirradiated SiPM 7 (figure 5.30c). They consist only of a baseline at approximately 0 V. Only one of the 100 waveforms contains a single signal, visible as a faint peak. Although occurring in the dark, this signal is indistinguishable from a signal caused by a single photon interacting with the SiPM.

In contrast to SiPM 7, the waveforms from SiPMs 6 and 7 show substantially more dark counts. The DCR is so large that the baseline shifts by up to 75 mV, as shown in figure 5.30b. Due to capacitive coupling on the DPNC 342 board, the baseline is measured between -50 and -100 mV, even though the signals are positive. The current signal does not change polarity, and a directly coupled readout system would see a positive baseline shift due to overlapping signals. When a signal with an amplitude of 100 mV per photoelectron is shifted by a random value up to 75 mV, advanced analysis techniques are necessary to achieve single-photon resolution.


5.5.5 DARK COUNT RATE (DCR)

The DCR ν of a SiPM channel is an essential quantity for the operation of a SiPM. Waveforms taken from single channels can be used to determine the DCR by identifying any signals. Here, two methods to determine the DCR from waveforms are presented, based on leading-edge discrimination, the same method used in MuTRiG.


COUNTING METHOD

For uncorrelated events occurring at a fixed mean rate ν , the probability of observing n events in an interval Δt is described by a Poisson distribution. “Memory effects” in SiPMs, such as afterpulses and crosstalk, have probabilities below 5% [42, 105], allowing the approximation of dark counts as a Poisson process.


The first method of determining the dark count rate is straightforward: counting the number of dark counts n in a waveform of length Δt . The probability density function



(a) SiPM 1, irradiated



(b) SiPM 6, irradiated



(c) SiPM 7, not irradiated

Figure 5.30: Shown are heat maps (overlays) of 100 waveforms with of 200 ns length each. The irradiated SiPMs 1 and 6 show a significant amount of dark counts and baseline shift. SiPM 7 (figure 5.30c) is not irradiated, and most waveforms do not contain any signal.

of the Poisson distribution is

$$f(n) = \frac{(v\Delta t)^n}{n!} e^{-v\Delta t}. \quad (5.15)$$

The expectation value and variance of this distribution are the mean rate multiplied by the time interval.

$$E[n] = V[n] = v\Delta t. \quad (5.16)$$

Dark counts are determined with leading-edge discrimination in waveforms of $\Delta t = 200$ ns. The distribution of the number of dark counts for one representative channel of SiPM 6 is shown in figure 5.31a. A Poisson distribution is fit to the data and shows good agreement, confirming the assumption of independent dark counts. The DCR $\nu = 24.85$ kHz is extracted from the best value $\nu\Delta t = 4.97$ using $\Delta t = 200$ ns.

INTERARRIVAL TIME METHOD

Using the Poisson approximation again, the probability of observing zero events in a time interval δt follows from equation 5.15. Since n only takes on discrete values, the probability of observing no events $P(n = 0)$ is equal to $f(0) = e^{-v\delta t}$.

The complement of that expression then gives the cumulative distribution function for the random variable δt :

$$F(\delta t) = 1 - e^{-v\delta t} \quad (5.17)$$

This variable is not to be confused with the fixed interval Δt , for example, defined by an oscilloscope.


The derivative of $F(\delta t)$ yields the probability density function

$$f(\delta t) = \nu e^{-v\delta t}. \quad (5.18)$$


The expectation value and variance of this exponential distribution are

$$E[\delta t] = \frac{1}{\nu} \quad \text{and} \quad V[\delta t] = \frac{1}{\nu^2}. \quad (5.19)$$

The time during which zero events are observed is exactly the interarrival time, i.e. the observation of δt allows measuring the average rate ν . However, direct observation of δt requires infinitely long waveforms. The limited waveform length imposes that at least two events must be observed within 200 ns, which significantly alters the distribution. This condition is circumnavigated by “stitching together” different 200 ns waveforms, as they are statistically independent. One example distribution and the corresponding fit are shown in figure 5.31b, showing good agreement between the data and an exponential fit. The best value for the DCR is $\nu = 24.74$ kHz here.






(a) A histogram of the number of dark counts in a time interval of 200 ns and a Poisson distribution fitted to the data. The fit result corresponds to a DCR of $4.97/200 \text{ ns} = 24.95 \text{ kHz}$.



(b) Interarrival times of dark counts. An exponential function is fit to the distribution with $\delta t > 30 \text{ ns}$.

Figure 5.31: Two methods of determining the dark count rate of the same SiPM channel using waveforms

(a) SiPM 7, not irradiated. $T = 25$ to 27°C , $V = 55$ V(b) SiPM 6, irradiated uniformly. $T = 27$ to 28°C , $V = 55$ V(c) SiPM 1, irradiated, higher particle fluence in channels with low numbers. $T = 25$ to 27.5°C , $V = 55$ V**Figure 5.32:** Per-channel dark count rate for three different SiPM arrays, determined by counting events in windows of 200 ns and fitting a Poisson distribution to the result

RESULTS

Applying the two methods to waveforms recorded of SiPMs 1, 6, and 7 allows the determination of the per-channel DCR. SiPM 7 was not irradiated, and most channels have a DCR of approximately 20 kHz. These rates are lower than the typical rates of 60 to 300 kHz stated by the manufacturer. A possible reason for that is the lower overvoltage than what the datasheet is based on (2.5 V instead of the nominal voltage of 3.5 V) [42].

As expected from figures 5.30a and 5.30b, the DCR is significantly higher than for SiPM 7. The per-channel DCR of SiPM 6 (figure 5.32b) is more than a factor of 1000 larger than the unirradiated SiPM. This factor is comparable to the mean increase in the current above breakdown for SiPM 6, which is 1244 with a standard deviation of 496 (see B.13a in the appendix). Most channels have a DCR between 20 and 25 MHz and a larger spread than the channels of the unirradiated SiPM. The channels in the centre (around channel 64) have a slightly higher DCR, possibly caused by the finite size of the beam spot that was centred with respect to the SiPMs.

The DCR of the channels of SiPM 1 depends on the channel number, as shown in figure 5.32c. During the irradiation of SiPM 1, the beam was misaligned for a certain time (see section 5.4.1). Considering the setup geometry and the misalignment of the beam in the $+y$ direction, which corresponds to lower channel numbers, the observed gradient from about 10 to 25 MHz is explained qualitatively.

The two presented methods are based on the same assumption of uncorrelated, random events. Applying both methods to the same data does not yield two independent results but serves as a consistency check.

5.6 ANNEALING

After the irradiation and measurements presented so far, the SiPMs were cooled again. The literature suggests partial recovery of the radiation damage at room temperature [7]. Furthermore, it suggests that the biggest change would happen early, and the SiPMs approach a more stable state.

To avoid this during beamtime, I built an ad-hoc “annealing setup” based on an alienated Peltier device whose sole purpose is heating the irradiated SiPMs. With the temperature sensor on the carrier board (figure 5.2) and a software-based PID controller, the SiPMs were heated to 80 °C. The SiPMs were biased above breakdown and kept in the dark. Over days, the current through the SiPMs decreased, showing clear evidence of annealing (see figure 5.33).

The curves show a fast drop in current in the first minutes after being heated to 80 °C, followed by a much slighter decrease under the assumption of two annealing components. Fitting a function of the shape

$$I(t) = I_0 + I_1 e^{-t/\tau_1} + I_2 e^{-t/\tau_2} \quad (5.20)$$




Figure 5.33: Current during annealing as a function of time. The left half shows the first ten hours, while the right shows the current over the course of five days. The dark solid lines show the result of a double-exponential fit with a constant baseline shown as a dotted line.

reveals the time scales of the two processes in terms of the decay times τ_1 and τ_2 . The quantity $1 - I_0/I(0)$ conveys the relative reduction of current expected at $t \rightarrow \text{inf}$.

Table 5.6: Time scale and relative recovery of high-temperature annealing

SiPM	τ_1	τ_2	$1 - I_0/I(0)$
SiPM 1	(0.54 ± 0.01) h	(1.84 ± 0.01) d	6.1 %
SiPM 3	(1.85 ± 0.05) h	(2.49 ± 0.06) d	8.3 %
SiPM 4	(1.03 ± 0.01) h	(2.02 ± 0.01) d	14.9 %
SiPM 5	(1.99 ± 0.01) h	(2.65 ± 0.01) d	7.7 %
SiPM 6	(1.37 ± 0.01) h	(1.61 ± 0.01) d	6.5 %

The fast annealing happens on a time scale of half an hour to two hours (see τ_1 in table 5.6), while the slow part has a time constant between 1.61 and 2.65 days (see τ_2 in table 5.6). The constant baseline I_0 and the current at $t \approx 0$ are compared to estimate the amount of recovery. The baseline I_0 is also shown as dashed lines in figure 5.33. The recorded data show a maximum recovery of 14.9 %.

5.6.1 DISCUSSION AND LITERATURE

Before the start of this controlled annealing, the SiPMs were at room temperature while IV curves and waveforms were taken. Installation in the annealing setup took different

amounts of time, during which the SiPMs could warm up while the current was not yet monitored. Additionally, irradiation took place at a higher temperature than expected for Mu3e. These facts imply that figure 5.33 does not show the full annealing process but rather the process after partial annealing already occurred. Therefore, the potential for recovery in figure 5.6 must be understood as a lower limit to the total recovery. Thermal annealing as a means to reduce the impact of radiation damage in Mu3e is discussed in the next section.

For a prototype version of the Hamamatsu S13552 SiPM used in Mu3e, a relative reduction of the dark current of approximately 50 % was observed after annealing at 35 °C over two weeks [106]. This gives rise to the hope that a similar recovery can be achieved with the final version of the SiPM used in Mu3e and irradiation with protons.

Table 5.7: Overview of partial annealing studies in SiPMs

Relative reduction	Radiation	Annealing Conditions	SiPM model	Reference
15 %	Positrons	5 d at 80 °C after partial annealing	Hamamatsu S13552	this work
50 %	Neutrons	2 w at 35 °C	Hamamatsu S13552 prototype	[106]
93 %	Neutrons	1 mo at up to 175 °C	Hamamatsu S13360	[107]
>90 %	Neutrons	“days to weeks” at 250 °C + forward bias	Hamamatsu S13360	[108]

Other studies have shown that annealing with applied *forward* bias at 250 °C leads to significantly better recovery than annealing with no or reverse bias [108]. Even at 250 °C, the authors emphasize that it takes “days to weeks” to recover SiPMs exposed to a neutron fluence of approximately 10^{12} cm^{-2} . Similarly, by annealing at temperatures up to 175 °C, dark current can be reduced by a factor of 15 [107]. These two studies were conducted with Hamamatsu S13360 SiPMs [46] with different properties than the Hamamatsu S13552 SiPMs arrays used in Mu3e [42]. In the context of the Mu3e SciFi detector, this could open up the possibility of recycling irradiated SiPMs exposed to a fluence so large that in-situ annealing does not suffice.

5.7 CONCLUSION OF THE IRRADIATION AND OUTLOOK

In Mu3e, a neutron-equivalent fluence of up to $2.8 \cdot 10^{11} \text{ cm}^{-2}$ is expected in the SciFi SiPMs based on the simulations presented in this chapter. That is approximately one-fifth of the fluence expected in the identical SiPMs in LHCb. To study the effect of radiation damage by Michel positrons, six Hamamatsu S13552 SiPMs were irradiated

with different fluences up to about 27 % of what is expected in Mu3e. In the following, the findings, their impact on Mu3e and possible mitigation strategies are discussed.

5.7.1 GENERAL

Irradiation of SiPMs with Michel positrons increased the leakage current and DCR of these SiPMs. The current below the breakdown voltage increased by a factor of 2.19 ± 0.89 per neutron-equivalent fluence of 10^{10} cm^{-2} . Above breakdown, the current through the SiPM increased by a factor of 200 ± 25 per 10^{10} cm^{-2} neutron-equivalent fluence.

Dark current is closely connected to dark counts: signals mimicking those of a single photon. Per single SiPM channel, such signals are observed at rates up to 30 MHz at $\Phi_{\text{eq}} = 7.5 \cdot 10^{10} \text{ cm}^{-2}$.

Literature suggests that irradiation leads to an increase in the breakdown voltage of SiPMs [3, 109]. In contrast to these findings, a slight *reduction* of V_{BD} is observed here. The authors of [3] found no significant increase up to $\Phi_{\text{eq}} = 5 \cdot 10^{13} \text{ cm}^{-2}$ and only 300 mV after $5 \cdot 10^{14} \text{ cm}^2$. Similarly, in [109] a study of SiPMs irradiated with an equivalent fluence between $1.2 \cdot 10^{13}$ and $2.2 \cdot 10^{14} \text{ cm}^{-2}$ is presented. Neglecting the shortcomings of the NIEL hypothesis in detail, their doses are several orders of magnitude higher than what is presented in this thesis. A linear increase of 4 V per $2.2 \cdot 10^{14} \text{ cm}^{-2}$, as described in [109] for Hamamatsu SiPMs, would be immeasurable at Mu3e's $2.8 \cdot 10^{11} \text{ cm}^{-2}$ and the doses administered herein.

The current at the breakdown voltage lies in the order of 1 nA per channel. At a voltage in the order of 50 V, the heat dissipation is in the order of 50 nW. Self-heating effects with such a low power seem unlikely, but the temperature inside the individual pixels was not measured.

A possible explanation would be a type inversion in the silicon. Silicon strip detectors irradiated in an 800 MeV proton beam show a decrease in depletion voltage up to a fluence of $1.5 \cdot 10^{13} \text{ cm}^{-2}$ ($\Phi_{\text{eq}} \approx 10^{13} \text{ cm}^{-2}$), followed by an increase with further irradiation [110]. The bulk damage changes the effective doping concentration and causes an inversion from n-type silicon to t-type silicon. The question of whether a type inversion really takes place in Hamamatsu SiPMs exceeds the scope of this thesis. Higher fluences and means to detect said inversion would be necessary to investigate this possibility.

However, not only the breakdown voltage of the irradiated SiPMs was reduced, but also the control sample. The IV curves before irradiation were taken before the SiPMs were bonded to the carrier board. The SiPMs were not removed from the carrier board for the measurements after irradiation. That carrier board contains a lowpass filter that causes a voltage drop over the 1 k Ω resistor (see figure 5.3). The current flowing through that resistor during the measurements to determine the breakdown voltage is in the order of 10^3 nA. Therefore, the voltage drop is below the order of 1 mV and

cannot explain the observations. Both measurements were taken at slightly different temperatures, which is corrected for using equation 5.12.

The two techniques to determine V_{BD} might not be equally suitable for irradiated and unirradiated SiPMs and introduce a bias that was not taken into account, though such a bias would not affect the control sample. A systematic user error at the probe station used to record the IV curves appears unlikely, as the measurement is conducted automatically after manual alignment. A possibility that cannot be excluded is an alteration to the setup that I was unaware of.

The slope of a linear fit is unaffected by a constant offset of V_{BD} . Based on the assumption that the unexplained difference between V_{BD} before and after irradiation is the same for all SiPMs, a reduction of $20\text{--}30\text{ mV}/(10^{10}\text{ cm}^{-2})$ as a function of neutron-equivalent fluence is obtained. The fit uncertainty is larger than 50%. Together with the unexplained difference with the control sample, the quantitative results are inconclusive. If the breakdown voltage decreases due to irradiation with positrons, the effect is smaller than the per-channel variation of one SiPM array.

5.7.2 IMPACT ON MU3E AND POSSIBLE MITIGATION

The measurements presented in this chapter were taken under conditions that do not fully reflect the ones in Mu3e. The following differences between the conditions in Mu3e and the measurements herein must be considered to understand the impact on the final detector.

A large share of the positron fluence in the upstream SiPM is caused by muons decaying in the beampipe. The simulation herein assumes a baseline collimation concept. However, better collimation far away from the detector would significantly reduce the amount of radiation exposure in the upstream detectors, thus extending the lifetime of the SiPM.

Linearly extrapolating the increase in the current above breakdown to the expected dose of Mu3e means an increase of up to 5300 ± 780 times. At room temperature and 3.5 V above breakdown, the dark current lies at 1.8 mA for the sample irradiated with 27% of the maximum fluence of Mu3e. Scaled to Mu3e, the current would be almost 7 mA, which exceeds the current high voltage supply specification. An upgraded version of the high voltage supply expected in 2022 will be able to supply 20 mA per channel, which is sufficient to handle such high currents.

BREAKDOWN VOLTAGE When the breakdown voltage changes, the bias voltage must be retuned. As external factors, such as the temperature, also affect the breakdown voltage, regular retuning is generally advised. The inconclusive result of $dV_{BD}/d\Phi_{eq}$ makes it challenging to predict the exact behaviour during Mu3e. However, all observed ΔV_{BD} were significantly smaller than 1 V. Therefore, no significant drift in the breakdown voltage of the SiPM is expected. What must be considered, however, is the fact that

not only the carrier board for irradiation contains a lowpass filter, but so does the SMB (section 6.4.2). The current version 2.5 is equipped with one 2.55 k Ω per 64 channels of the SiPM. A current of 1 mA through the entire SiPM array would cause a voltage drop of more than 1 V, which outweighs any intrinsic change in the SiPM itself.

DARK COUNT RATE Plain extrapolation to the total Mu3e dose would lead to a per-channel DCR in the order of 100 MHz, exceeding the acceptable limit of 450 kHz. A more realistic expectation has to include the details of the operation in the final experiment, which potentially reduce the true DCR. Personal experience and literature ([65, 105]) suggest a DCR reduction due to cooling by a factor of two for every 10 K in unirradiated SiPM. The DCR was measured at temperatures between 20 and 30 °C. As a reaction to the findings herein, the Mu3e SciFi group adapted its specification to make use of the full potential of the cooling system and operate the SiPMs at temperatures below -10 °C [111]. Based on the abovementioned factor, that cooling would mean a reduction by up to a factor of 16 compared to the values found herein. A factor 16 would already reduce the DCR to \approx 6 MHz.

Another option to reduce the DCR is by lowering the bias voltage. The measurements here were performed with a bias voltage of 55 V, approximately 2.5 V above breakdown, leaving only little room for reduction, though. For comparison, the manufacturer recommends operating the SiPM 3.5 V above breakdown. Previous measurements have shown that a larger bias voltage is associated with a better time resolution [40] and PDE [65]. Therefore the bias voltage has to compromise between DCR, PDE, and time resolution, but no significant reduction is expected in comparison to the measurements from this chapter.

Since dark counts occur spontaneously and independently, the signal corresponds to the one of a single photoelectron. The probability of two cells of one SiPM experiencing simultaneously is low compared to the probability of prompt crosstalk. With a probability of less than 5 %, a signal in one cell triggers a neighbouring cell [42]. Applying this to DCR, the trigger frequency of signals with two or more photoelectrons is more than a factor of 20 lower than the one-photoelectron rate. This factor, combined with cooling, could reduce a rate of 100 MHz to acceptable levels. Operating the SciFi detector at thresholds above one photoelectron is, therefore, a potent way to suppress DCR with a minor impact on the time resolution and efficiency [40, 91, 111].

As shown in figure 5.30, a high DCR affects the baseline, making single-photon resolution impossible with leading-edge discrimination. Fine-tuning the thresholds during operation will be increasingly more challenging with irradiation.

Another way of reducing the rate of such false signals is using clustering (see also section 6.4.1). Signals from the scintillating fibres seldom only reach one SiPM channel. Requiring a coincident signal in neighbouring channels would reduce the rate of dark counts that are recognized significantly. As a reaction to the large values of DCR

observed in the works leading to the results of this chapter, such a feature was added to the second version of the MuTRiG ASIC. The probability of observing zero dark pulses from a SiPM channel with a DCR of ν within a time window Δt is $p = e^{-\nu\Delta t}$. Together with the expected rate of real signals of 650 kHz, the probability that a dark pulse is coincidentally validated by one of n possible channels is

$$e^{-n\Delta t(\nu+650\text{ kHz})}. \quad (5.21)$$




Figure 5.34: Calculated hit rates due to dark pulses that are validated by a coincident signal, assuming 650 kHz signal rate plus DCR. The rates are shown for two different time window settings and three different configurations of neighbouring channels that can validate hits. These settings correspond to the hit validation function of MuTRiG 2. The maximum tolerable rate of 450 kHz follows from the MuTRiG limitation of 1.1 MHz.

Figure 5.34 shows the random coincidence rates for the possible coincidence settings. When choosing the maximum number of channels that can validate hits and the longer time window of 19 ns, the maximum tolerable hit rate of 450 kHz is only reached once the DCR exceeds 1.8 MHz. The strictest setting in which only the two adjacent SiPM channels can validate a hit within a time window of 6 ns suppresses random coincidences up to a DCR of 5.9 MHz. Thus, MuTRiG 2 could prolong the lifetime of the SciFi SiPMs by a factor of 13.

Together, these options allow for the reduction of the DCR to tolerable levels. When the final readout hardware becomes available, a course of action can be developed, balancing the positive and negative impacts of the individual options.

Matching signals from both fibre ends and requiring a coincidence of upstream and downstream does not reduce the data rate of MuTRiG. The switching board is the earliest point in which the upstream and downstream data streams are combined

(see section 6.4.6). Like in the case of the single-photon resolution, the extent of this reduction is yet to be studied once the final detector hardware is available.

THERMAL ANNEALING Apart from adapting the operating conditions to the changes due to radiation, the SiPMs can be treated themselves. The irradiated SiPMs were heated to 80 °C for five days to accelerate thermal annealing. A fast and a slow component were observed with respective time scales in the order of hours and days, leading to a reduction of the dark current of up to 15 %. This number must be considered a lower limit on the true potential of a controlled annealing recovery since annealing of the SiPMs studied herein started before these measurements.

In the Mu3e experiment, similar accelerated annealing is possible by heating the SiPMs in situ during beam interrupts and detector maintenance. Such heating could be performed by using the cooling infrastructure to conduct heat *to* rather than *from* the SiPMs, considering the surrounding detector components.

The annealing studies, summarized in table 5.7, show high potential for mitigating radiation damage by thermal annealing, especially at high temperatures. The potential for in-situ annealing in Mu3e depends on the surrounding detector components and has to be evaluated. Thermal annealing outside the detector at high temperatures is a promising alternative to discarding irradiated SiPMs. That way, the drawbacks of the other reactions to radiation damage could be minimized.

6

Data acquisition and SciFi integration

Working on the development of the Data Acquisition System (DAQ) was a substantial part of my doctorate. The main contributions concerned the integration of the SciFi detector into the overall DAQ. These contributions include hardware design (section 6.4.4) and software development (section 6.6). Due to the nature of the work going into this chapter, it is more descriptive than the previous ones. It presents the status of the Mu3e DAQ [112] with a strong focus on the readout of the SciFi detector and my contributions.

6.1 GENERAL OVERVIEW

Mu3e is a triggerless experiment that utilizes custom DAQ hardware and software. The readout allows the three subdetectors to share as much readout hardware design as possible to simplify development. The SciFi and SciTile detectors digitize the analog SiPM pulses with the same ASIC called MuTRiG described in section 6.4.1. Readout Field Programmable Gate Array (FPGA) boards share a common design for all different subdetectors. The detailed data flow from data generation to final storage is portrayed in section 6.3. DAQ hardware and software are described in sections 6.4 and 6.5. The main part of the work in this thesis in DAQ concerns the integration of the SciFi detector into the common DAQ system, culminating in the 2021 integration run described at the end of this chapter.

6.2 DATA PRODUCERS AND CONSUMERS

There are *data producers* and *data consumers* in Mu3e, as in virtually all DAQ systems. The data are sent via a network of FPGAs from the *producers* to the *consumers*.

There are two types of data producers in Mu3e: the MuPix pixel sensors and MuTRiG ASICs for the two scintillating timing detectors. Hits in the MuPix pixel sensors are digitized on-chip. The corresponding data are sent out via up to three 1.25 Gbit/s Low Voltage Differential Signalling (LVDS) data links per sensor connected to the rest of the DAQ. The chips of the inner two layers are connected via all three LVDS links. A single link is sufficient for the chips in the outer double layers, as the hit multiplicity is lower at larger radii. In total, 2844 MuPix chips are expected to produce 56 Gbit/s of data for a muon stopping rate of 10^8 per second.

The two scintillating timing subdetectors are both read out by the MuTRiG ASICs. Although the SciTile detector has almost twice the number of readout channels, it is expected to produce less data than the SciFi detector, as listed in table 6.1. The reason for the lower data rates lies in the geometry of the detectors (see chapter 2). The SciFi detector is in the central station, where a higher particle density leads to higher hit multiplicity and more radiation damage (see chapter 5). The data are transmitted via LVDS with a data rate of 1.25 Gbit/s.

Table 6.1: Comparison of numbers of channels and data rates of the two MuTRiG-based subdetectors. Each ASIC has 32 channels. Although the SciFi detector has fewer readout channels, it produces more data due to hit multiplicity and noise [1]

Quantity	Fibre Detector	Tile Detector
Number of ASICs per module board	4	13
Number of module boards	24	14
Total number of ASICs	96	182
Total number of readout channels	3072	5824
Expected data rate	28 Gbit/s	17 Gbit/s

The *end consumer* on-site is a data collection server connected to PSI computing services. Ultimately, the data are stored in the *PetaByte Archive* at Swiss National Supercomputing Centre (CSCS) in Lugano [113].

6.3 DATA FLOW AND FILTERING

In the following, the Mu3e DAQ system is explained by following the flow of detector data from its production to eventual storage depicted in figure 6.1. Figure 6.2a illustrates the data path from the SciFi detector to storage via the hardware described in detail in section 6.4.

The primary data producers connect to subdetector-specific Detector Adapter Boards (DABs) that pass the data to generic Front-End Boards (FEBs) inside the Mu3e solenoid. On each FEB is an FPGA serving as the interface to the readout system outside of

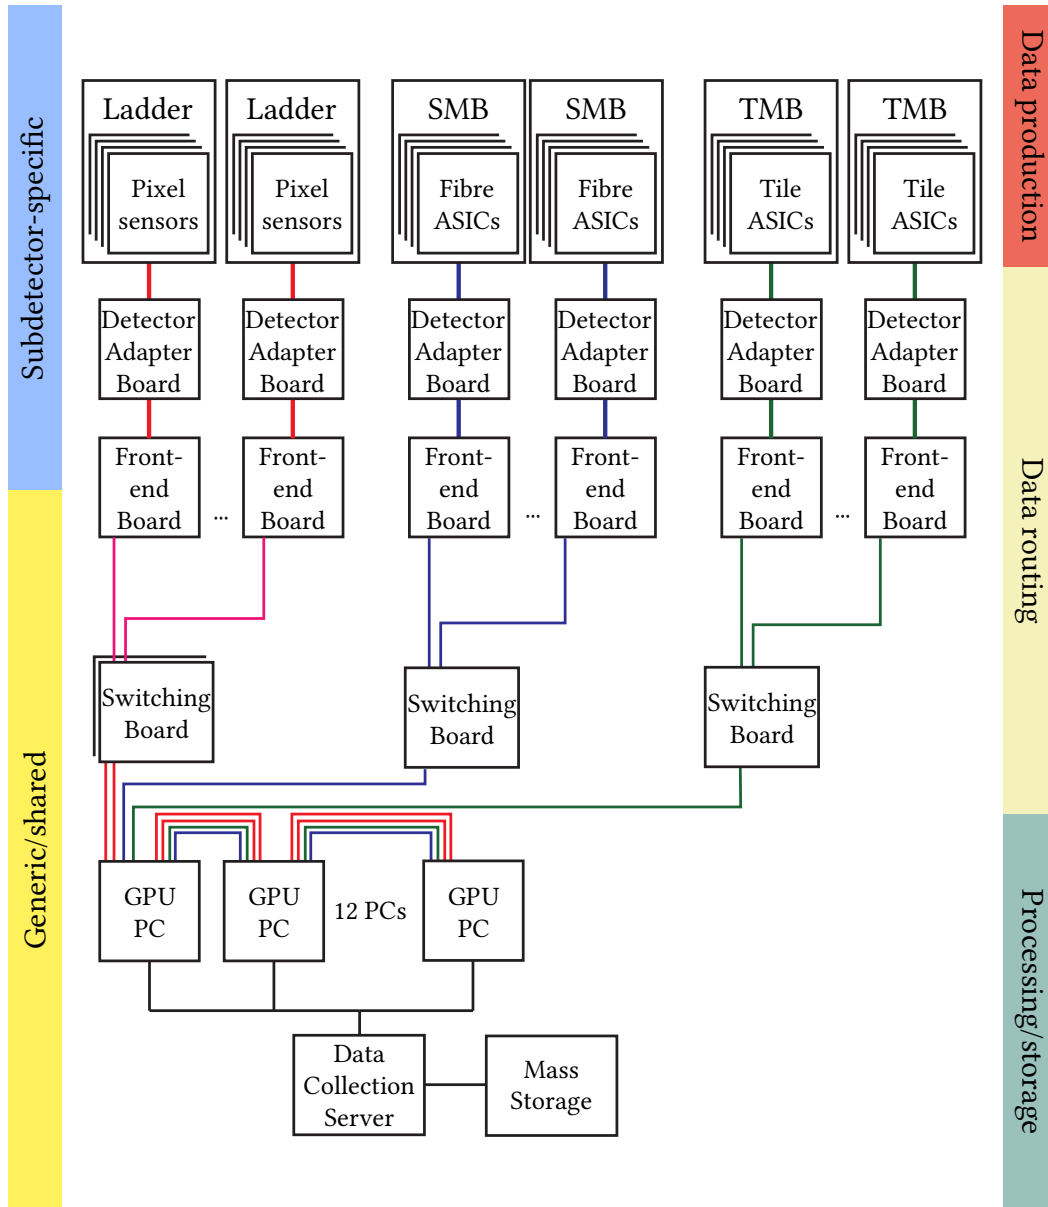




Figure 6.1: Overview of the Mu3e DAQ topology. The hardware of the front-end boards is generic, but the firmware is subdetector-specific. Adapted from [1].

the magnet via optical data links. Interleaved in the hit data are monitoring data from the detectors and FEB itself. Up to 34 of these links connect to one of four Switching Boards (SWBs) in the counting house, where the data streams are merged and time-aligned, and monitoring data is filtered out. Here, data from the upstream and downstream halves of the SciFi are merged, and coincidences between the two parts can be determined and used as a criterion for further data processing. All SWBs send the time-aligned data streams to the same PC of the filter farm. There, the information from the whole detector during overlapping time slices (frames) of 64 ns is used to select frames of interest. First, tracks are reconstructed from pixel hits on a Graphics Processing Unit (GPU). If a combination of three tracks is identified as a candidate for the decay $\mu^+ \rightarrow e^+e^+e^-$, the full detector data are used for a slower Central Processing Unit (CPU)-based reconstruction. Frames of interest passing the selection criteria are then stored on disk.

In the opposite direction, control and configuration commands are sent from the SWBs via an optical uplink to the FEBs and the detectors (see section 6.6). A synchronized clock and reset signals are injected into the front-end FPGAs via the same uplink. Each MuTRiG and MuPix ASIC in the experiment can be addressed and configured individually. Its position is defined by the path shown in figure 6.2b.



(a) Flow of SciFi data from the SciFi Module Board (SMB) via Detector Adapter Board (DAB), Front-End Board (FEB), Switching Board (SWB) and an event filter, to storage.



(b) The flow of configuration data.

Figure 6.2: Flow of SciFi data (a) and MuTRiG configuration data (b) in the opposite direction. One SWB serves all twelve FEBs with two links to an SMB each. Each MuTRiG is configured independently using a unique address.

6.4 DAQ HARDWARE

This section describes the hardware necessary to realize the data flow in figure 6.2 for the SciFi detector. The analog SiPM signals are digitized by MuTRiG, the primary SciFi

data producer.

6.4.1 MuTRiG

MuTRiG (**M**uon **T**iming **R**esolver including **G**igabit-link) is the ASIC that digitizes silicon-photomultiplier signals from the SciTile and the SciFi detector. It is a mixed-mode readout ASIC designed for the rate capabilities and timing requirements of Mu3e.

MuTRiG is a successor to the STiC family [114, 115]. STiC was designed to provide high time resolution for the EndoTOFPET-US (endoscope for time-of-flight positron emission tomography) and other time-of-flight applications. While meeting the Mu3e criteria regarding time resolution, the single 160 Mbit/s serial data link significantly limits the maximum event rate. The design goal of MuTRiG was to tune the existing STiC design to the needs of Mu3e, especially regarding the data rate. The analog front-end and the Time-to-Digital Converter (TDC) design of STiCv3 were kept, but the digital logic circuits required adaptations.

GENERAL OVERVIEW

The detector design leads to significant differences in the produced signals and background rates. In particular, the SciFi detector must detect and discriminate single photons at a rate of up to 620 kHz per channel [1]. At the same time, a typical SciTile hit produces in the order of 1000 photons per SiPM channel. The expected hit rates range from $\mathcal{O}(10\text{ kHz})$ in the most-downstream tiles up to 53 kHz [44]¹. Because the SciFi detector must be operated at very low thresholds, effects such as dark counts will produce additional signals and potentially take up bandwidth. This led to a design requirement of 10^6 events per channel, which is met by introducing a 1 Gbit/s serial data link and thanks to a modification in the hit protocol for SciFi hits (see section 6.4.1).

THE ANALOG FRONT-END

Each MuTRiG has 32 analog front-ends containing an SiPM input channel and signal discrimination. Each of the 32 SiPM input channels (see figure 6.3) can be biased with a constant voltage of up to 900 mV compared to ground. This allows fine-tuning the bias voltage of each SiPM connected to a single MuTRiG. The range covers the spread of breakdown voltages on Hamamatsu S13552 SiPM arrays [42, 91], the sensor used in the SciFi detector.

¹The cited source studies detector concepts with four segments of the tile detector, corresponding to the Phase II design and various tile sizes. Here, only the Phase I design with two tile segments is considered.

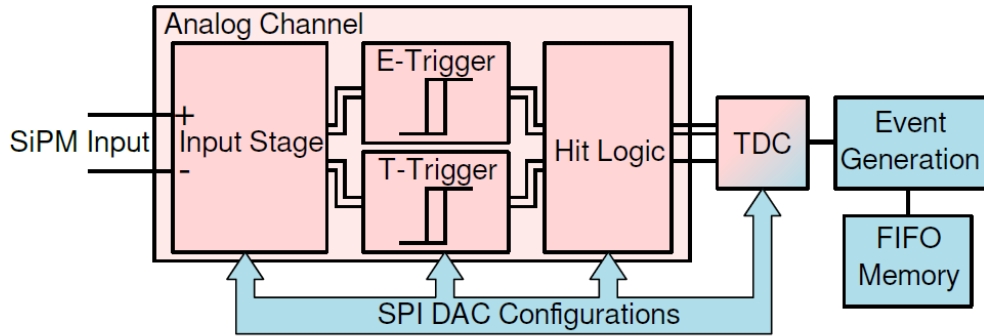


Figure 6.3: Diagram of a single MuTRiG Channel. A silicon photomultiplier can be connected differentially or single-ended. The analog signal is split into two branches for *timing* (T) and *energy* (E) discrimination. The hit logic unit combines the two trigger signals so that a single TDC can convert information from both branches to digital timestamps. The timestamps are then further processed and stored in on-chip memory. All stages can be configured via Serial Peripheral Interface (SPI) [1, 43, 116]

The input stage of the analog front-end consists of a symmetrical structure that allows for single-ended and differential inputs. After the input stage, the signal is split into one *timing* branch and one *energy* branch. The timing branch discriminates the signal against a low threshold to extract precise information on the arrival time while minimizing time walk effects [54]. The energy branch offers a linearised time-over-threshold measurement to determine the energy deposited in the scintillator material, via the number of photons detected. In the SciFi detector, only the timing information is used.

The two discriminated signals are combined in the hit logic unit, as shown in figure 6.4. At the top, the analog signal and the two thresholds are shown. The middle section shows the digital discriminator output for the timing and energy triggers and the exclusive disjunction (XOR) of the two. The bottom shows the coarse and fine counters whose state defines the corresponding timestamps (see section 6.4.1).

THE TIME-TO-DIGITAL CONVERTER

One Time-to-Digital Converter (TDC) per analog input channel converts the timing and energy information from the hit logic unit to digital data. The data are stored in timestamps, encoding the time interval between a fixed reference and the time of arrival of a signal in the analog front-end. The energy information is determined by the time difference between the first and second rising edge of the exclusive disjunction of the two discriminated signals (see middle and bottom of figure 6.4).

A single time base unit provides two counters for all 32 TDCs, a 625 MHz *coarse counter* and a *fine counter* with 32 bins per $\frac{1}{625 \text{ MHz}} = 1.6 \text{ ns}$. The time information is




Figure 6.4: Working principle of the two branches inside MuTRiG. The top shows the analog SiPM signal and the two thresholds. The outputs of the two thresholds and the exclusive disjunction are shown in the middle. The state of the two counters (bottom) determines the timing and energy timestamps. Adapted from [1, 43].

obtained from a Phase Locked Loop (PLL) and a 16-stage Voltage Controlled Oscillator (VCO) [43]. The coarse counter runs through $2^{15} - 1$ binary states mapped onto sequential timestamps later. Although each bin of the fine counter has an average width of 50 ps, the time bins are non-uniform, described by differential non-linearity, which needs to be corrected [40, 114]. The TDC needs about 30 ns to recover and cannot digitize another hit during that time, leading to a dead time of 3 % at a signal rate of 1 MHz.

THE DIGITAL PART

The digital logic circuitry of MuTRiG provides several improvements and adaptations to the STiCv3 design. The event generator takes timestamps from the TDC and converts them into *events*. There are two types of *events*: short and long events. Long events with 48 bits contain information from both the timing and energy timestamps. Short events with 27 bits only contain timing information and one bit indicating whether the input signal exceeded the energy threshold. The short event mode was introduced for the SciFi detector to allow for higher event rates. The shorter an event (in bits), the more events can be transmitted per unit of time at a given bit rate. Each timestamp has a possible *bad* flag raised if an illegal VCO state is observed. Detailed information about the data encoded in the two event types can be found in the appendix in table C.7.

These events are buffered and grouped into *frames*. Each frame contains up to 255 events plus additional header and trailer data. The trailer contains the result of a Cyclic Redundancy Check (CRC) (see [43] for details), which indicates potential transmission errors. The full event layout is in the appendix in table C.8.

Data from MuTRiG chips are 8b/10b encoded [117, 118], and hit information is zero-suppressed. Zero-suppression is a simple way of data compression that reduces data rates in detector regions with lower hit rates. In 8b/10b encoding, words consisting of 8 bits are mapped to 10-bit symbols. The average number of logical ones and zeroes is equal. The benefits of this encoding are DC balance, detection of possible errors, and clock recovery from the data stream. In addition, so-called *control symbols* are integrated into the data stream for synchronization and to mark data frame boundaries. These are valid 10-bit symbols that do not encode 8-bit words.

COINCIDENCE CLUSTERING LOGIC

In reaction to the increase of DCR from radiation damage (chapter 5), MuTRiG version 2 introduced a *coincidence clustering* feature. This feature allows detecting the coincidence of hits in nearby channels. A hit is only sent out if a second hit occurs in one of the closest six channels² within an adjustable time window. The coincidence window is 6 ns to 19 ns long, and for each channel, the condition, which of the six

²Three channels to either side

neighbouring channels are considered, can be adjusted. Connections to neighbouring MuTRiGs allow coincidences across chip boundaries, so channel 1 of one chip can be set to work in coincidence with channel 32 of the contiguous chip.


CHIP CONFIGURATION

The chip configuration is written via an SPI interface. A list of settings can be found in appendix C.1. Chip configuration and the rest of the “slow control” are integrated into the Mu3e data path between the SWBs and the FEBs (see figure 6.2a). The front-end FPGAs are connected to the SPI interface on the MuTRiG via the DAB and SMB. Each MuTRiG chip can be addressed and configured individually.


6.4.2 SCIFi MODULE BOARD



(a) First SMB prototype (SMB1) with four MuTRiG 1 chips and connectors at the right of the picture to connect to a SiPM array



(b) Second SMB prototype (SMB2) with an integrated SiPM array (hidden, on the right side) and two out of four possible bonded MuTRiG 2 chips. The black connector carries data signals; the top wires supply the SiPM bias voltage, and the bottom wires supply power.



(c) Rendering of the third SMB version (SMB2.5). While functionally almost identical to the SMB2, its connections and dimensions fit the Mu3e detector.

Figure 6.5: SciFi Module Board prototypes. The first two versions were used in test beam campaigns and proved the design. Only the third iteration will fulfil the spatial constraints inside the Mu3e detector.

Each SMB, designed at the University of Geneva, hosts four MuTRiGs. The first version of the board, denoted as *SMB1*, is equipped with four MuTRiG 1 chips and has four Hirose FH41 connectors to connect to silicon photomultipliers.

The second version, *SMB2*, is built as a single rigid-flex-rigid circuit board containing four MuTRiG 2 chips and a 128-channel silicon photomultiplier array, as in the final design. One *SMB2* can read out one side of a SciFi ribbon by coupling with screws.

This design allows for better signal transmission between the SiPMs and MuTRiG. The more compact design is a step toward the geometry required to fit the final experiment. Data connections are realized via a Samtec QSH-030 connector. Power and SiPM bias voltage are supplied via cables directly soldered to the board.

At the moment (Q3/2022), a revision of the second SMB version is under development at the University of Geneva, denoted as version 2.5. The changes necessary to use this SMB2.5 in the Mu3e experiment include a reduction of the physical dimensions of the board down to a width of 26 mm and the replacement of the Samtec connector by *micro-twisted pair* cables soldered to an interposer (see section 6.4.3) [1].

A fourth version (SMB3) will be needed for compatibility with a future revision of the MuTRiG chip. All current SMB prototypes can be seen in figure 6.5.

6.4.3 MICRO-TWISTED PAIR CABLES


As the Mu3e design is very compact, the spatial constraints for signal cables are very tight. To fulfil these constraints, data are transmitted from the SMBs to the DABs over AWG36³ twisted pair cables with polyimide coating and an impedance of 90 Ω [1, 119].

6.4.4 DETECTOR ADAPTER BOARD

The DAB is the adapter (as the name suggests) between the front-end electronics and the detector module boards (e.g. SMBs). The DAB is placed inside the service support wheels at the two detector extremes (see chapter 13.3 in [1]). These are located at a distance of more than one metre from the SciFi detector. To mitigate potential signal losses, the DAB is equipped with an equalizer circuit for each 1.25 Gbit/s data link.

By collaboration specifications, all signals between DAB and detectors must be differential links. However, the SMB2 prototype had both differential and single-ended signals. As the logic voltage level of MuTRiG is higher than that of the front-end FPGA, a level-shifter translates the signals between the two components, thus preventing damage and signal loss.

As shown in figure 6.7, a FEB has two *banks*, embodied by two separate Samtec QTH-040 connectors (CON2 and CON3). Each of these banks has enough bandwidth and electrical connections to allow the operation of one SMB. Therefore, the first DAB prototype was designed to connect a single SMB to one of the two FEB banks. After successful tests with the prototype, *DAB2* was designed to connect two SMBs to one FEB.



(a) First SciFi DAB prototype. This prototype connects one SMB to one FEB bank.
 (b) DAB2. The highlighted part corresponds to the circuitry in (a) and exists in two copies, one for each FEB bank

Figure 6.6: SciFi Detector Adapter Board (DAB) prototypes. One Front-End Board (FEB) can read out up to two SciFi Module Boards (SMBs), one per *bank*

6.4.5 FRONT-END BOARD

The core of the FEBs is an Arria V FPGA [120]. On one end, detector data are received via multiple 1.25 Gbit/s links from the DAB via CON2 and CON3. On the other end, a Samtec Firefly ECUO-B04 optical transceiver offers four inputs and four outputs, each running at 6.25 Gbit/s to connect a FEB to a SWB via optical fibre cables.

For the SciFi detector, the front-end FPGA can sort incoming hits by their timestamp, cluster those hits spatially and discard them based on cluster size. Sorting is necessary as the hit data do not necessarily arrive in the order in which hits occur. Real hits typically come in clusters of several SiPM readout channels, while *dark counts* appear as individual channels without correlation between channels. Discarding such *fake hits* helps to significantly reduce the data produced by the SciFi detector, especially after irradiation.

Embedded into the front-end FPGA is a NIOS II *soft microprocessor* [121]. This general-purpose microprocessor can execute instructions written in the C programming language. The FPGA is programmed either via a debug interface or a MAX 10 FPGA reading from onboard flash memory. That flash memory can be written remotely so that firmware modifications are possible even when the FEB is installed and not

³This corresponds to a diameter of 0.127 mm




Figure 6.7: A Front-End Board (FEB). The connection to the DAB is realized via the connectors CON2 and CON3, serving one SMB each. Data is sent to a Switching Board (SWB) via two optical transceivers (not installed here). The clock signal can either be supplied via the optical transceivers or a differential electrical signal (top left). A JTAG port allows debugging and control of the Arria V front-end FPGA. Picture adapted from [1].

accessible anymore.

6.4.6 SWITCHING BOARD




Figure 6.8: Switching board, developed under the name PCIe40. The heatsink obscures the Arria 10 FPGA. Figure from [1].

These *PCIe40* SWBs were developed for the LHCb and ALICE upgrade and offer 48 optical inputs and outputs (see figure 6.8) [122–124]. Each SWB connects between 12 and 34 FEBs to a computer via Peripheral Component Interconnect Express (PCIe) and the filter farm. The PCIe interface⁴ is used to transfer configuration data to the subdetectors. The data from the different FEBs are merged and forwarded to the filter farm.

The SWB is the first point at which data from the upstream and downstream halves of the SciFi detector are combined, and hits can be matched to suppress false signals. In the case of the SciFi detector, a single SWB is connected to all twelve FEBs via two 6.25 Gbit/s optical links, one uplink and one downlink.

⁴Technically, two eight-lane PCIe 3.0 interfaces connected to one 16-lane edge-connector

6.4.7 RECEIVER BOARD

Each PC of the filter farm is equipped with commercial Terrasic DE5a NET [125] *receiver boards* that carry an Intel Arria 10 FPGA, four QSFP optical transceivers and two banks of DDR4 memory. This board is used extensively in small-scale DAQ setups where it additionally takes the role of the SWB.

6.5 DAQ SOFTWARE

A custom software framework based on Maximum Integrated Data Acquisition System (MIDAS) [126] manages run transitions, chip configurations, monitoring, and channels data streams. The following section gives an overview of the software aspect of the Mu3e DAQ.

6.5.1 MIDAS

MIDAS is a data acquisition system developed for high-energy physics experiments at PSI and TRIUMF. The system allows controlling experiments in so-called *run states* and *transitions*, e.g. switching from the *idle* state to the *running* state in which data are recorded. MIDAS can run on a distributed network of machines, so it is possible to separate data taking and data storage servers, computers managing parts of the experiment, and online analysis. A so-called Online Database (ODB) is used to store settings, configurations, and the general state of the experiment in a way that eliminates ambiguity. Programs running on remote machines can register so-called *callback functions* that are called when an ODB entry changes. Remote clients can send data to the ODB via so-called *front-ends*: programs running locally with access to the ODB.

The *history* system periodically logs variables describing the experiment conditions (see section 6.6). Graphs of the logged data are generated in real-time and allow easy monitoring.

MIDAS contains a web server that allows access to the ODB and history plots and provides a web front-end to most MIDAS features. Custom pages can be integrated and hosted to control the experiment via the manipulation of ODB values. MIDAS provides a JavaScript library for the simplified development of these custom pages.

6.6 SLOW CONTROL

The slow control system was a particular focus of this work's contributions to the Mu3e DAQ system. *Slow control* describes handling all data of slowly changing quantities that can be read out infrequently. Examples are monitoring of voltages and temperatures, but also detector configuration. These data are also often *slow* when comparing data

rates. MuTRiG, for example, has a 1.25 Gbit/s interface for data, but configuration information is sent via a 128 kbit/s SPI interface.

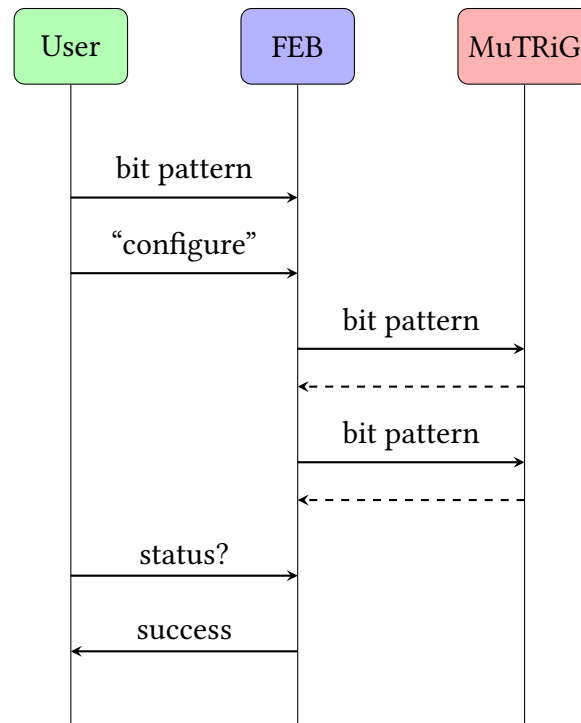


Figure 6.9: Message sequence chart of MuTRiG configuration. A bit pattern containing all settings is generated in software and written to the front-end FPGA on the Front-End Board (FEB). Via a *configure* command, transmission to the MuTRiG via SPI is initiated. The MuTRiG automatically responds with the previous bit pattern, so sending the same pattern twice allows the verification. The result of this verification is stored in a status register on the FEB that can be read out by the user

6.6.1 FUNCTIONALITY

The slow control system is embedded in the MIDAS-based online framework of Mu3e.

Configuration settings must be sent from a user to the ASIC. Such settings are stored in MIDAS' ODB or separately on disk and then sent through the DAQ network. The FEBs handle the direct communication with the detector ASICs. Therefore, commands are sent to the FEBs. The front-end FPGA allows remote *read* and *write* commands to its memory and registers, thus allowing bidirectional communication.

6.6.2 SciFi INTERFACE WITH MIDAS

A MIDAS custom page provides a graphical user interface for the MuTRiG settings stored in the ODB. The Mu3e framework generates a bit pattern in the format required by MuTRiG (see appendix C.1 for details). That information is then written to the front-end FPGA. Transmission to the ASIC is initiated via an Remote Procedure Call (RPC) to the NIOS II soft microprocessor on the FPGA⁵.

The same interface is used for communication in the opposite direction, reading back information from the FPGA. An RPC triggers the FPGA to write the requested data to a predefined memory address where it can be read out afterwards.

The communication between the FEB and MuTRiG is implemented via SPI (details in appendix C.2). The process is initialized via a MIDAS custom page that serves as a user interface to the ODB. This interface allows reading back the written pattern with the caveat that the read-back pattern is not the current pattern but the previous one. Writing the same pattern twice allows the detection of significant communication issues.

6.6.3 SciFi COUNTERS AND RESETS

At the time of writing, the DAQ development is a work in progress. The feature set is being expanded, and improvements are made to the framework. The following list of commands for the SciFi detector that can be issued via the slow control interface is not comprehensive but describes the state by the end of 2021.

- *All off*: A predefined configuration pattern stored in the NIOS memory can be sent to all MuTRiG chips connected to a FEB. This pattern is meant to bring the chips into a well-defined state that consumes little power by turning off most on-chip features. A chip configured like this sends an *idle signal* via the LVDS link and empty frames.
- *PseudoRandom BitStream (PRBS)*: All MuTRiGs connected to a FEB are configured in a state to send frames of hits containing a pseudorandom bit sequence (described in H. Chen's dissertation [43]). This mode allows for debugging the datapath.
- *PLL Test*: In this mode, the FEB injects test pulses with a frequency of 100 kHz into the MuTRiG TDC. The MuTRiGs are configured such that these test pulses bypass the analog front-end and generate hits. As the generation of those pulses is periodic, the difference between timestamps is expected to be uniform. An ill-configured PLL that is not properly locked can lead to seemingly irregular

⁵NIOS II is only used to configure the scintillating fibre and tile detectors

time intervals between test pulses. Therefore, this mode allows the tuning of the PLL to bring it to a functional state.

- *Configure MuTRiG*: Configure one of the MuTRiGs connected to the DAB with a pattern written to the FPGA. Typically that pattern is written just before the configuration command. The ASICs can be addressed and configured individually.
- *Read Counters*: Read back monitors for the datapath on the front-end FPGA. Each of the following counters is read per ASIC.
 1. *Number of Hits*: The number of observed *MuTRiG hits*, useful to see if an ASIC sends data.
 2. *Time*: A counter that simply counts up, useful to see that the system is working.
 3. *Number of Frames*: The number of MuTRiG frames received.
 4. *Number of Bad Frames*: The number of MuTRiG frames with a failed CRC.
 5. *Number of PRBS Words*: The number of 48-bit words received in the PRBS debug mode (see above).
 6. *Number of PRBS Errors*: The number of PRBS words that did not match the expected word. From one word, one can calculate the next word arithmetically.
 7. *Number of LVDS Words*: The total number of 8b10b symbols received by the FEB.
 8. *Number of LVDS Errors*: The number of 8b10b symbols received by the FEB with disparity errors.
 9. *Number of Data Synchronization Loss*: The number of losses of LVDS synchronization.
- *Reset Counters*: Reset all counters listed above to zero.
- *Configure Reset Skew Phases*: Configure reset skew phases on the front-end FPGA. Timing skew describes the phenomenon of a signal arriving at different components at different times. As skew cannot be prevented, the skew phases allow mitigation when it occurs.

6.7 INTEGRATION RUN 2021

In 2021, Mu3e members from the Universities of Heidelberg, Mainz, and Geneva, PSI and ETH Zürich set up the most complete Mu3e prototype until this point, denoted

as *integration run* [9]. Two layers of pixel detector chips and two SciFi ribbons were brought to PSI. After years of parallel development, the different subsystems were brought together and operated in a vertical slice containing the whole DAQ chain.

The word integration is deliberately ambiguous, as it involves several aspects. Mechanical integration means that the subdetector prototypes were mechanically compatible with one another and were mounted inside the Mu3e solenoid at the PiE5 beamline (see [1]). The systems were powered from a common source and cooled independently but simultaneously. The setup was integrated into the infrastructure at PSI, both the existing one and the infrastructure we set up.


The most important aspect for this chapter is DAQ integration. A large portion of the work in this thesis was devoted to the integration of the SciFi detector. In addition to the slow control described in this chapter, several MIDAS front-ends were created, controlling different parts of the detector, such as meteorology, power supplies, and cooling. DAQ integration also means compatible hardware. Figures 6.10b and 6.10a show a SciFi ribbon mounted above the pixel detector. Note that the SciFi ribbon is covered by UV protective foil (see section 4.7.1). Figure 6.10c shows the DAB2 mounted to the common minicrate and cooled with the same cooling block as the pixel DABs.

After the installation of the detector and the setup of all services, the MuTRiGs were to be commissioned. Injecting periodic test pulses is expected to produce equidistant timestamps. The distribution of the interarrival times between timestamps would show a very narrow peak at the time period of the injected signal. However, the received data did not match our expectations, as can be seen in figure 6.11. Rather than a narrow peak, a broad, falling distribution was observed. The explanation that was found later is that a bug caused a misinterpretation of the timestamps. If the bits constituting the timestamps are interpreted in the wrong order, a periodic signal can generate signals that resemble noise.


Even though the data was unusable, the integration run demonstrated that many parts did work. The readout infrastructure of the pixel detectors was shown to work. That demonstration can be projected to the SciFi part, at least for the components that are common or identical. SciFi data made it through the whole DAQ chain, albeit disordered in one step. Additionally, the configuration with the ASICs could be confirmed by simple tests such as changes in power consumption. As that configuration requires the software to initiate it and the DAB (among other components) to function, their performance could be verified.

6.8 CONCLUSIONS AND OUTLOOK


The Mu3e DAQ framework is still under constant development and constantly growing. As anticlimactic as the 2021 integration run ended for the SciFi group, the efforts put into preparation, commissioning and debugging brought Mu3e much closer to a final,



(a) Full SciFi ribbon on the beampipe



(b) SMB2 mounted on the beampipe



(c) DAB2 mounted in a minicrate

Figure 6.10: Photographs of the SciFi hardware used during the 2021 integration run




Figure 6.11: A distribution of the interarrival times between subsequent MuTriG timestamps during the 2021 integration run. The MuTriG receives a periodic signal expected to generate a very narrow distribution. A bug mangled up the data, so the data are random.

running state than ever before. The general functionality of the readout hardware and software was successfully tested.

Mu3e plans to commission its detector in 2023. By then, the hardware must be final, and the DAQ framework must work. For the SciFi detector, this means the development and production of the final SMB, which requires only minor modifications to the existing design. During the integration run, we demonstrated that the DAQ hardware works—only our data did not make sense. In 2022, the different groups continued the joint efforts, and with the readout bug fixed, it was possible to read sensible SciFi data. This missing puzzle piece and further developments since the 2021 integration run shape a positive outlook.

7

Conclusions and outlook

The Mu3e experiment will search for the cLFV decay $\mu^+ \rightarrow e^+e^+e^-$ with an unprecedented Single Event Sensitivity (SES) of better than $1 \cdot 10^{-16}$. The work presented in this thesis focuses on contributions to the Mu3e SciFi detector consisting of 250 μm scintillating fibres coupled to SiPM arrays and its integration into the shared DAQ system.

ATTENUATION IN SCINTILLATING FIBRES Previous observations showed that the light attenuation in scintillating fibres is well-described by an exponential function only if the distance between the fibre end and light emission is sufficiently long. In this dissertation, the attenuation of light in thin scintillating and wavelength-shifting fibres over shorter distances was studied (see chapter 4). At distances on the scale of the Mu3e SciFi detector (30 cm), additional modes contribute to the transmission of light (see section 3.1.2). These modes have significantly shorter attenuation lengths than those stated in a single-exponential description. Since the SciFi detector is only 30 cm long, it is sensitive to differences in short-range attenuation.

Seven different fibre types were excited in different positions with side-induced fluorescence, and the radiant flux at the fibre end was measured with a PIN diode. All but one fibre type showed an attenuation component with a short attenuation length. A two-component exponential function is fit to the data to extract the long and short components of the attenuation length and was found to show short components between 10 and 35 cm. The measurements presented here enable detailed simulations of the fibre performance in small-scale applications such as Mu3e. In addition to the radiant flux measurements, a spectrometer was used to deduce the wavelength-dependent attenuation. Light with a short wavelength is attenuated significantly stronger than long-wave light. It was shown that UV light significantly deteriorates scintillating fibres. Illuminating a Kuraray SCSF-78 fibre with a 405 nm LED for 24 h

causes a more than 20 % decrease. At 260 nm, NOL-11 fibres deteriorate more than twice as fast as SCSF-78 fibres.

For Mu3e, this confirms the tentative choice to use the Kuraray SCSF-78 fibres. A short attenuation length is not problematic but must be understood, and the wavelength-dependent attenuation does not shift the spectrum significantly with respect to the photo-detection efficiency of our SiPM.

POSITRON IRRADIATION OF SiPMs Mu3e requires $2.5 \cdot 10^{15}$ muon decays on target to reach its Phase I goal of an SES of $2 \cdot 10^{-15}$ based on an efficiency of $\epsilon \approx 20\%$. Given the high number of muon decays and the small size of the detector, a positron fluence of 4 to $5 \cdot 10^{12} \text{ cm}^{-2}$ is expected in the Hamamatsu S13552 SiPMs of the SciFi detector. Under the NIEL hypothesis, this corresponds to an equivalent fluence of 1 MeV neutrons of 2 to $3 \cdot 10^{11} \text{ cm}^{-2}$. After irradiation with Michel positrons, the S13552 SiPMs showed increased leakage currents below and above breakdown.

The current through the SiPM, when biased with a voltage of 3.5 V above V_{BD} , increased by a factor of 200 ± 25 per 10^{10} cm^{-2} neutron-equivalent fluence. The irradiated SiPM showed a DCR of several 10^7 Hz per channel when operated at room temperature. Such a high rate exceeds the capabilities of the Mu3e readout, which cannot handle more than 450 kHz. Additionally, that high DCR makes single-photon resolution impossible due to baseline shifts. Extrapolated to the total fluence expected in Phase I, this would imply rates of about 100 MHz per channel. After irradiation, the SiPMs were heated to 80 °C for approximately five days, and a reduction of current of approximately 15 % was observed, showing evidence of thermal annealing (see section 5.6).

The findings in this thesis (and preliminary studies leading to them) triggered the development of a new coincidence clustering feature in the MuTRiG 2 ASIC. This feature suppresses data from signals unless they are validated by another signal in a proximate channel within a certain time window. Since dark counts occur independently from each other, a coincidence requirement lowers the data rate significantly. An estimation yields suppression by up to a factor 13, but measurements with the actual SiPMs and MuTRiG are pending.

Furthermore, the strong effect from positron irradiation stressed the importance of SiPM cooling. While the original cooling concept foresaw keeping the SiPMs at room temperature, the final design foresees a temperature below -10 °C [111]. Lowering the temperature of a SiPM by 10 °C halves the DCR so that the revision causes a reduction by a factor of eight. Introducing components below the freezing point of water increases operation complexity and imposes stringent requirements on the gas purification of the helium cooling plant.

On paper, the combined mitigations allow sufficient suppression. The real-world operability can only be investigated with additional hardware.

OUTLOOK The Mu3e subdetectors are currently being readied for complete construction and commissioning in 2023, and engineering runs in 2024. Therefore, the focus is shifting away from design to finding operation modes. In the following, I would like to present future topics of interest based on the work in this thesis.

Since light propagation is based on total internal reflection, the surrounding medium affects part of the light rays. The measurements herein were conducted on single fibres in air. Some, but not all, transmission modes in double-clad fibres rely on the reflection on the interface between the outer cladding and the surrounding medium. In order to further improve the understanding of the attenuation in fibre ribbons, systematic measurements of fibres potted in epoxy may be performed.

Working with single fibres has shown that minor defects can lead to a significant amount of light loss. Furthermore, light of 405 nm is part of the emission spectrum of fluorescent tubes, as found in many laboratories. It is therefore advised to store scintillating fibres with additional protection and minimize the exposure time when handling. Additionally, UV-free illuminants should be considered when working with fibres.

The efficiency of suppressing DCR with the new coincidence clustering feature of MuTRiG 2 has been tested in laboratory measurements. However, the effect on the final detector setup has yet to be evaluated. This will be an important measurement for the SciFi group.

DCR does not only lead to spurious digital signals but also analog noise. A high DCR leads to a baseline shift, complicating single-photon resolution or making it impossible. This must be considered when developing the procedures for detector tuning. With unirradiated SiPMs, thresholds can be tuned based on the distinct rates per photoelectron (see, for example, figure 5.4 in [40]). The tuning concept must also work when the rate increases gradually rather than in individual steps. In the case that any single-photon resolution is lost, the thresholds could be determined by the absolute rate rather than by the identification of the individual photoelectron levels. The target rate is then limited by the rate capability of MuTRiG.

Additionally, a high DCR will affect time resolution and efficiency. The effect on the timing performance must be evaluated using testbeam data with a precise time reference. A figure of merit that should be determined is the limiting DCR at which the specifications of the SciFi detector are no longer fulfilled.

These evaluations are possible with the SiPMs I irradiated, either by soldering the irradiated SiPMs to a new SMB or by designing an adapter that utilizes the connectors on the carrier board (see figure 5.2). Testing the operation with the final readout hardware might reveal unforeseen behaviour of MuTRiG. Therefore, it is strongly advised to test the irradiated SiPMs with the next generation of the SMB.

Annealing of radiation damage in SiPMs has been previously observed and could be seen in the positron-irradiated SiPMs as well. Based on literature research, an even

larger effect can be expected. I recommend investigating the possibilities of in-situ annealing by heating the SiPMs while installed in the detector as far as possible without damaging other detector components. Alternatives are slower annealing at ambient temperature during detector downtime and annealing at even higher temperatures outside of the detector. If all other mitigations fail, the SiPMs can be replaced thanks to the modular structure of the SciFi detector.

A high bias voltage allows for good time resolution, pulse height and efficiency but also increases the DCR. While the DCR measurements presented here were performed with a bias voltage of 55 V, it was previously suggested to operate the SiPMs with a bias voltage of 59 V (chapter 14.1 in [40]). Solely based on the DCR, I would recommend a bias voltage as low as possible. The new results show the necessity to update the earlier suggestion. A trade-off between efficiency, time resolution, and DCR has to be found in the final setup. The optimal tuning and calibration procedures must be determined in situ with the real beam.



Attenuation

A.1 EXPLANATION

This appendix contains data from all scintillating fibres under test in chapter 4. The data are presented grouped by fibre type first, and then per individual fibre. For most fibre samples, the following figures are shown:

SPECTRA An overlay of the spectral flux obtained from stimulating the fibre at different distances from the spectrometer. The magnitude can be compared within the same figure only.

SPECTRAL FLUX RELATIVE TO THE CLOSEST ONE The spectral flux data are smoothed and divided by the closest spectrum, which is typically the brightest one, i.e. the one with the largest spectral flux. In some cases, these ratios are larger than unity, indicating that exciting the fibre further away from the spectrometer results in more light yield. One possible reason is improper alignment in the setting leading to bending losses at these short distances.

DIFFERENTIAL ATTENUATION LENGTH All spectra have the same binning in λ . This is exploited to determine $\frac{\partial\Phi_e(z)}{\partial\lambda}$ for a single point in λ . A single exponential is fit to each set of $\frac{\partial\Phi_e(z)}{\partial\lambda}$ and the result of this fit is presented as a function of λ . The points indicate the attenuation length of the best fit, while the “standard error” is shown as a coloured band (see also chapter 4).

INTEGRATED FLUX AS A FUNCTION OF DISTANCE The aforementioned spectra are integrated numerically to obtain the total radiant flux Φ_e at each distance z . A single exponential is fit to the values $\Phi_e(z)$ and the attenuation length Λ and its standard error are determined.

RADIANT FLUX The last figure contains radiant flux data recorded using a PIN photodiode and a power meter. The blue data points are used to determine the long component of the attenuation length Λ_{long} in a first step. In a second step, a double-exponential function with a fixed Λ_{long} is fit to the yellow and blue data points to obtain Λ_{short} . The grey, struck through data points are not used in those fits.

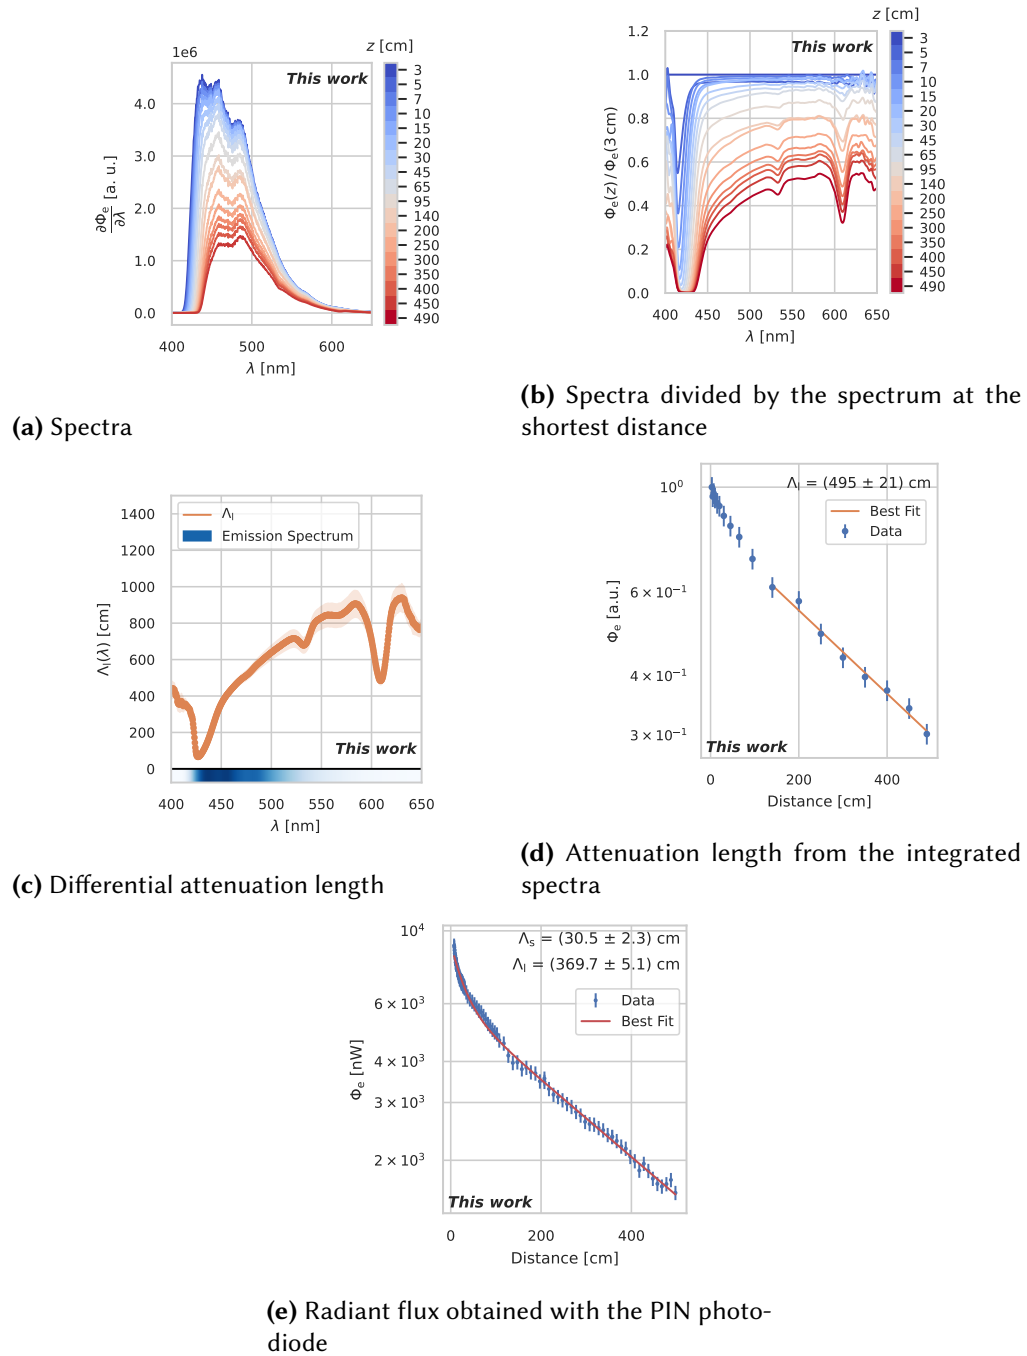





Figure A.1: SCSF-78 Fibre 1




(a) Spectra




(b) Spectra divided by the spectrum at the shortest distance



(c) Differential attenuation length




(d) Attenuation length from the integrated spectra




(e) Radiant flux obtained with the PIN photodiode


Figure A.2: SCSF-78 Fibre 2




(a) Spectra




(b) Spectra divided by the spectrum at the shortest distance



(c) Differential attenuation length




(d) Attenuation length from the integrated spectra




(e) Radiant flux obtained with the PIN photodiode


Figure A.3: SCSF-81 Fibre 1




(a) Spectra




(b) Spectra divided by the spectrum at the shortest distance



(c) Differential attenuation length



(d) Attenuation length from the integrated spectra



(e) Radiant flux obtained with the PIN photodiode

Figure A.4: SCSF-81 Fibre 2

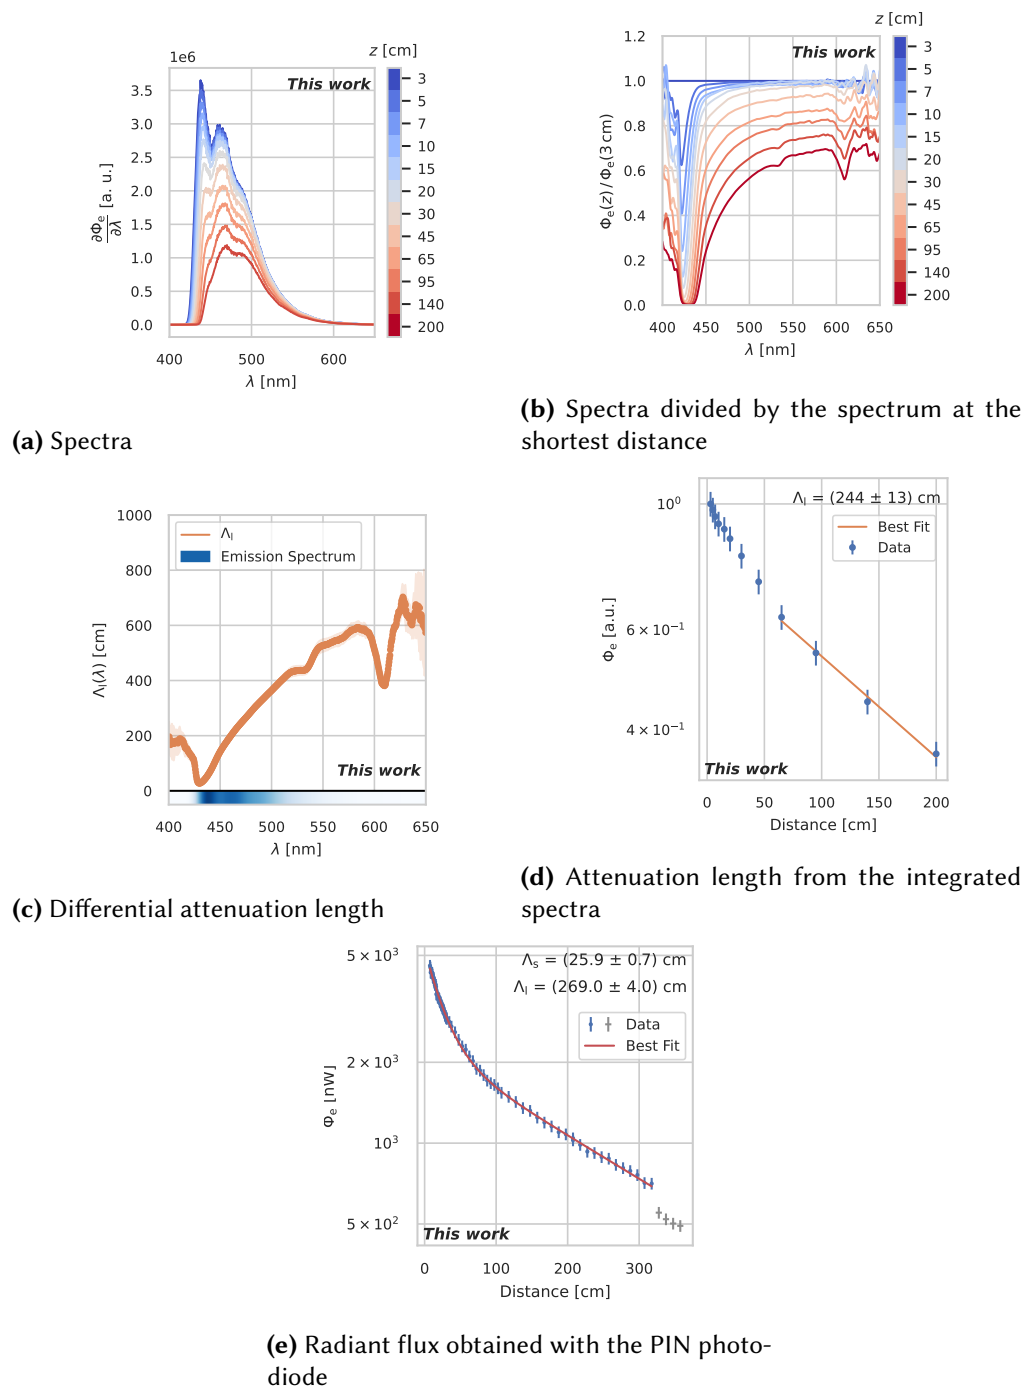


Figure A.5: SCSF-81 Fibre 3

Prototype fibres based on Nanostructured Organosilicon Luminophores (NOLs) 11 [58]. The samples are stimulated with a Thorlabs LED385L operated at 10 mA.

A.4.1 FIBRE 1

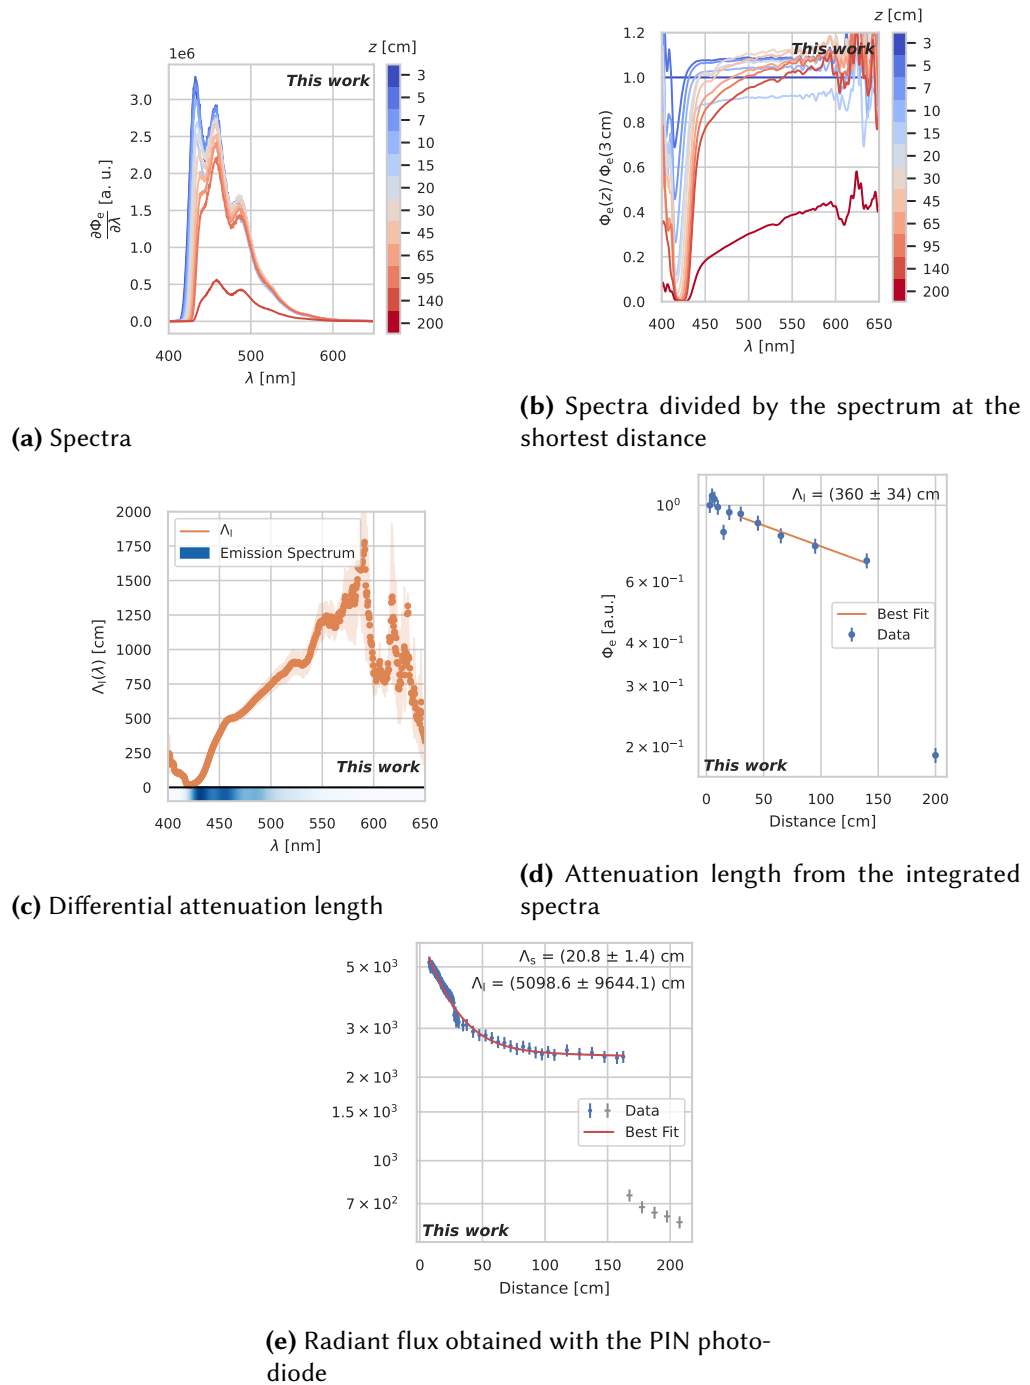





Figure A.6: NOL-11 Fibre 1




(a) Spectra




(b) Spectra divided by the spectrum at the shortest distance



(c) Differential attenuation length




(d) Attenuation length from the integrated spectra




(e) Radiant flux obtained with the PIN photo-diode


Figure A.7: NOL-11 Fibre 2




(a) Spectra




(b) Spectra divided by the spectrum at the shortest distance



(c) Differential attenuation length



(d) Attenuation length from the integrated spectra



(e) Radiant flux obtained with the PIN photodiode

Figure A.8: NOL-11 Fibre 3, the fibre was damaged at the 2 m mark as can be seen in (e), and in the last spectrum

A.5 BCF-12 ROUND

Round multi-clad, blue-emitting, scintillating fibre with a diameter of 250 μm .

A.5.1 FIBRE 1

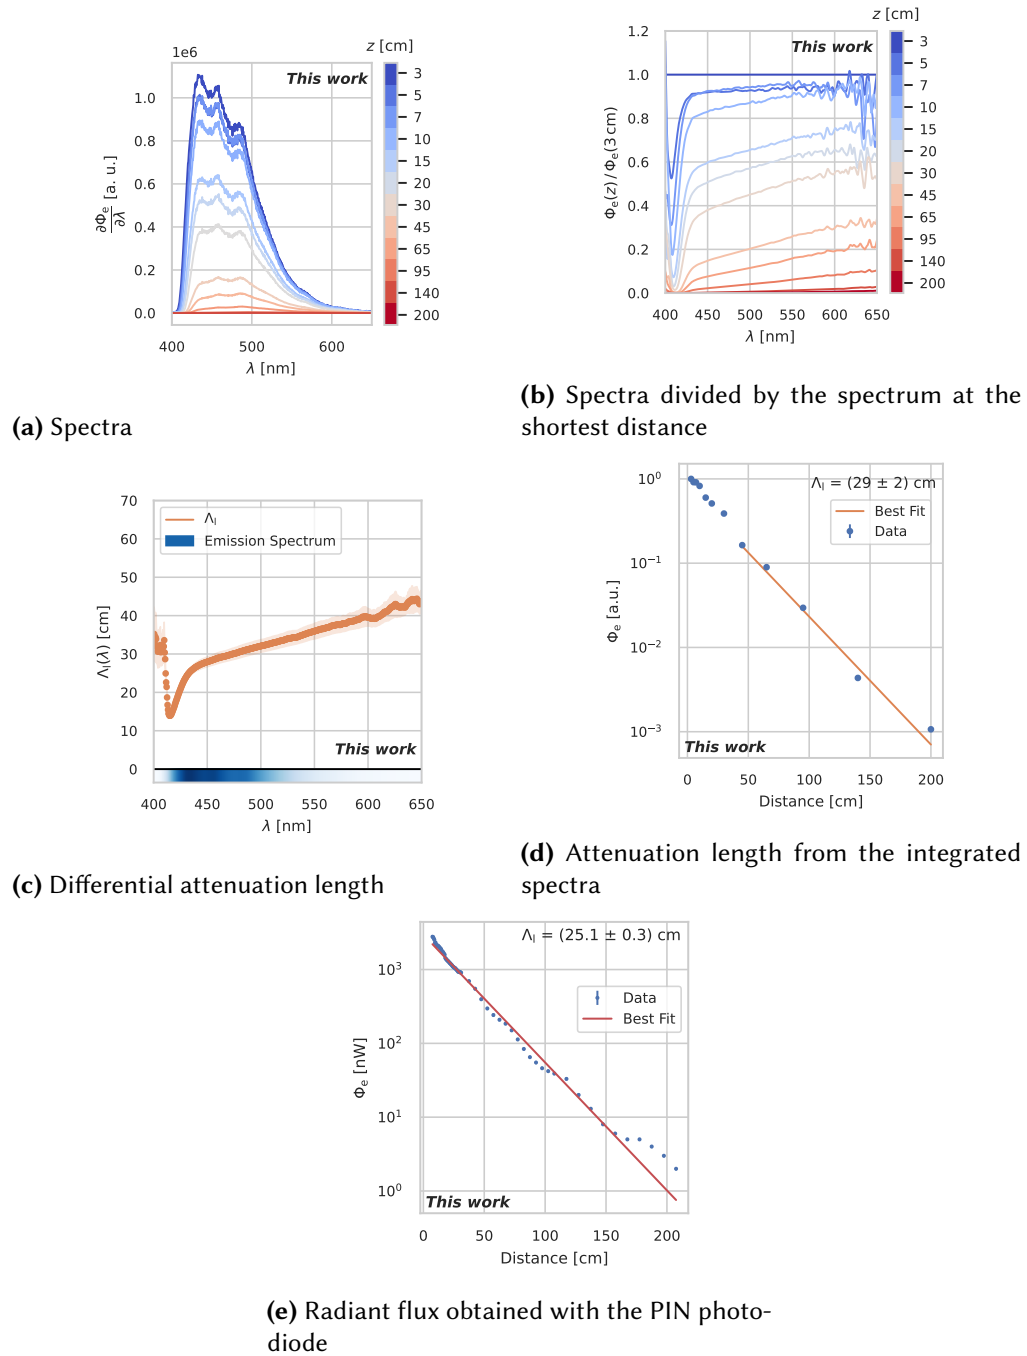


Figure A.9: BCF-12R Fibre 1

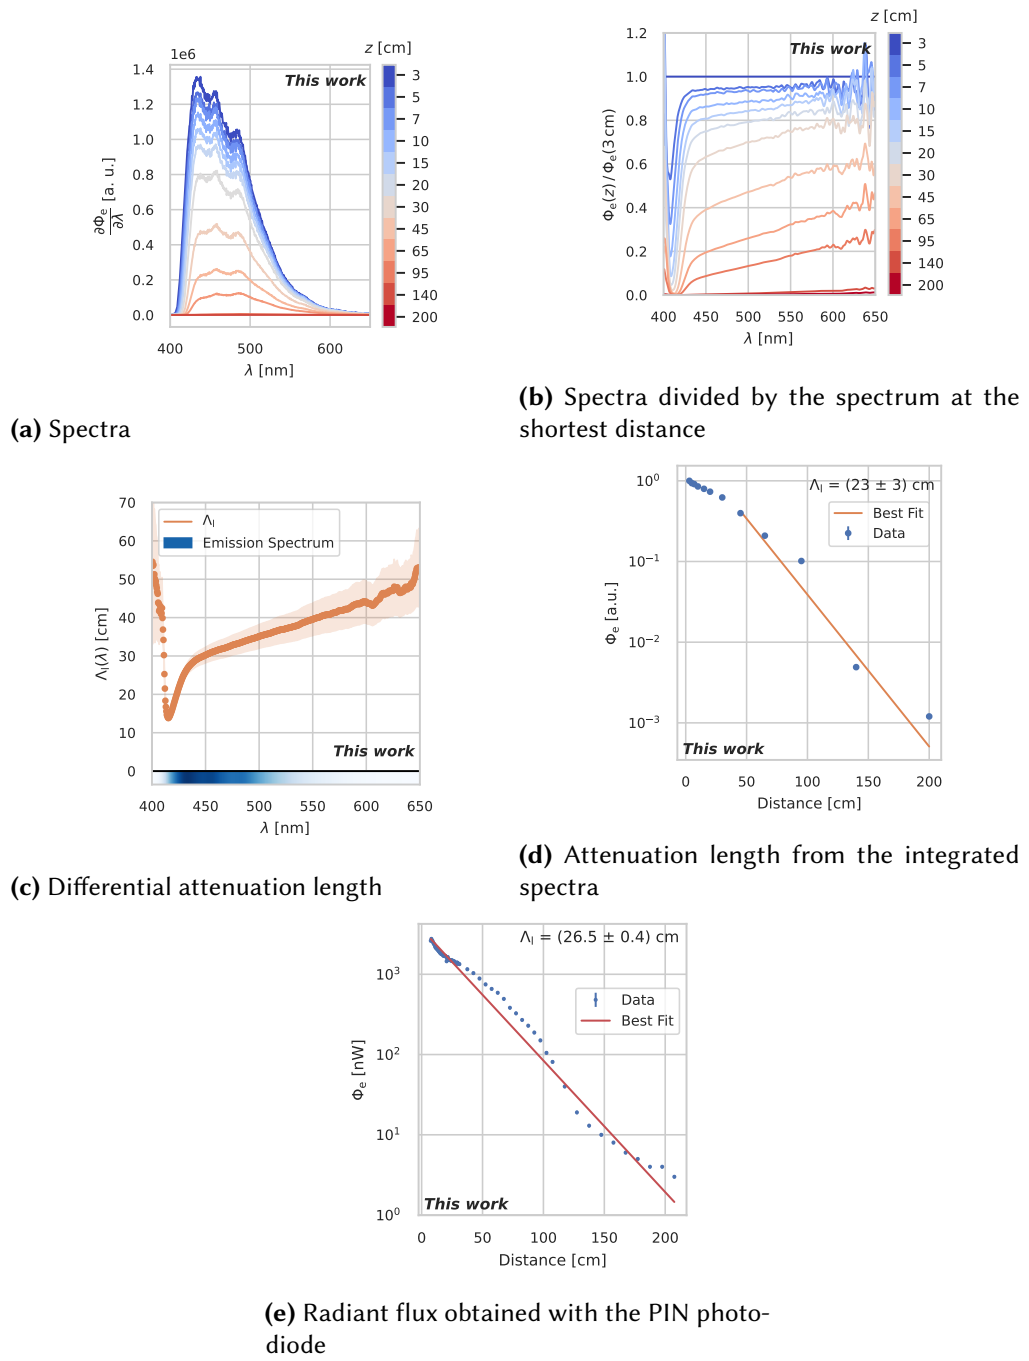


Figure A.10: BCF-12R Fibre 2

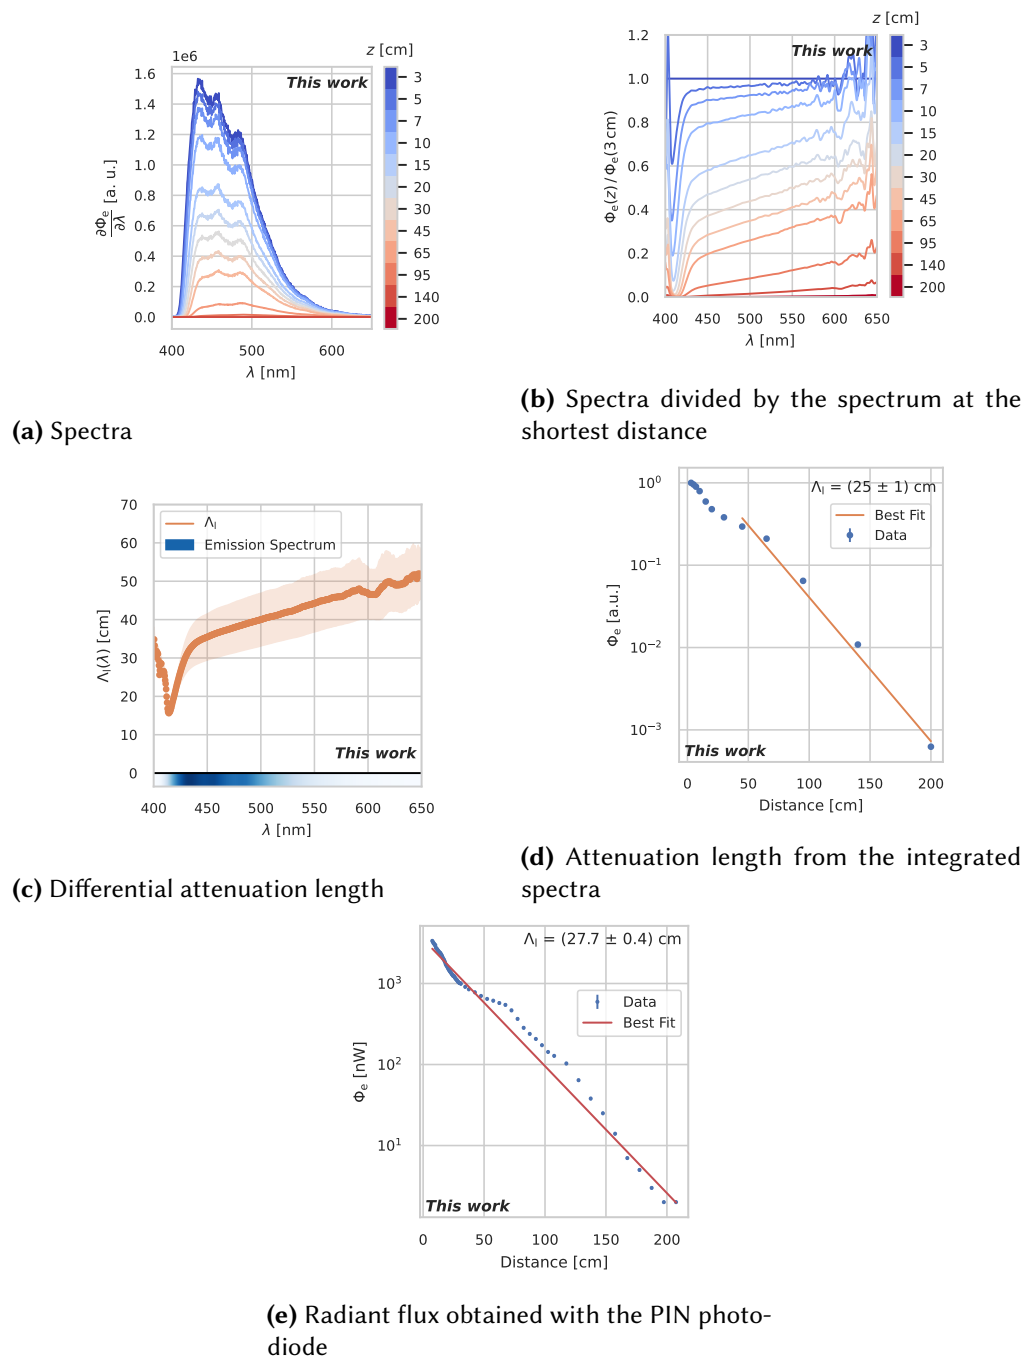


Figure A.11: BCF-12R Fibre 3

A.6 BCF-12 SQUARE

Square multi-clad, blue-emitting, scintillating fibre with a diameter of 250 μm .

A.6.1 FIBRE 1

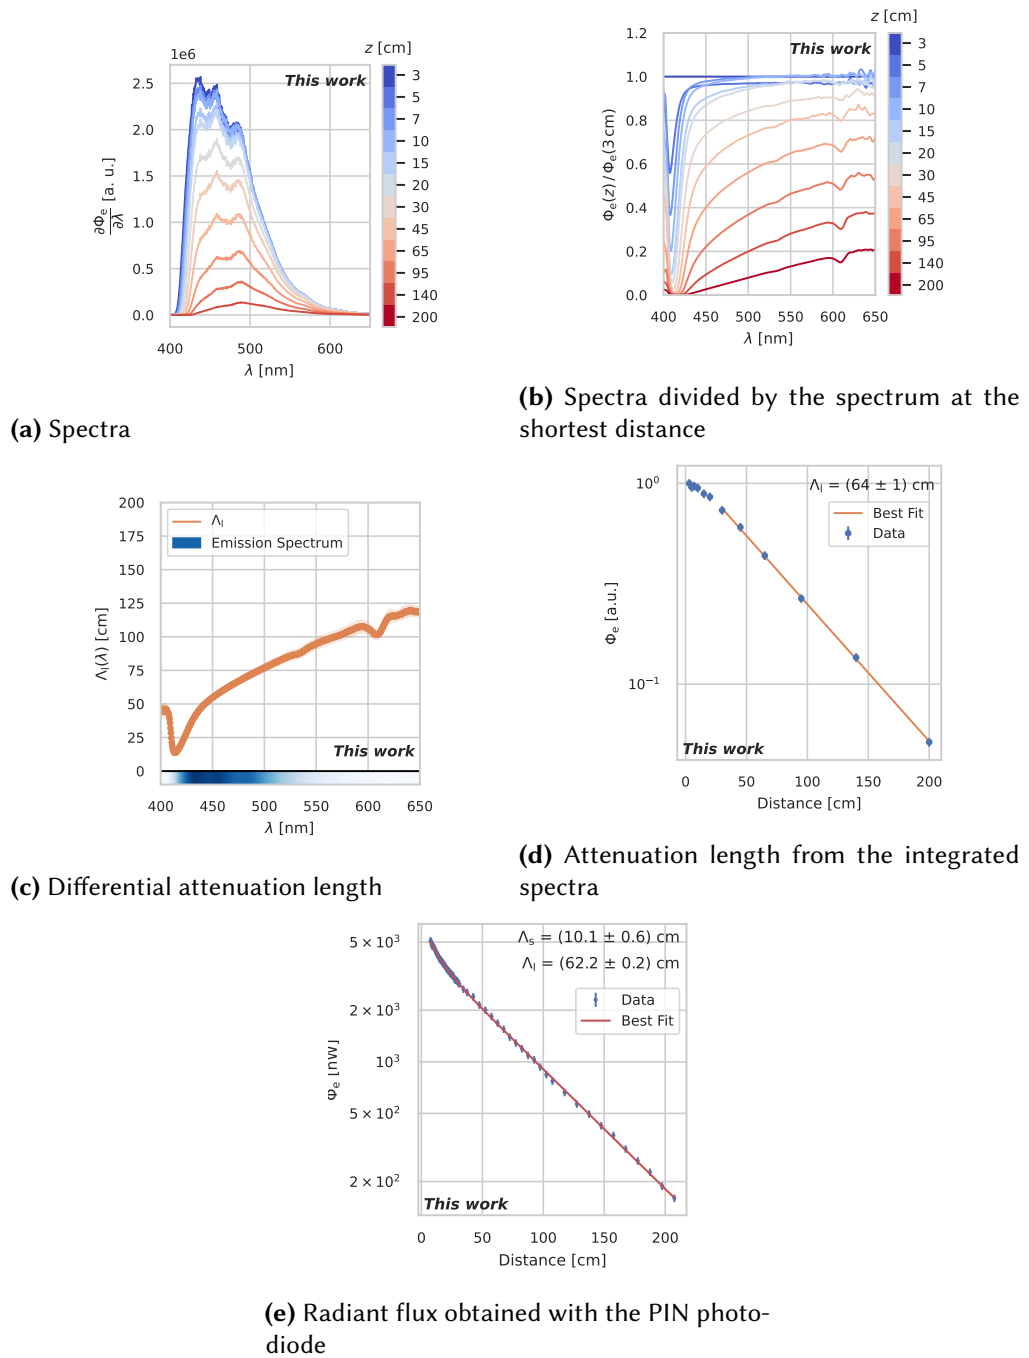


Figure A.12: BCF-12 Fibre 1

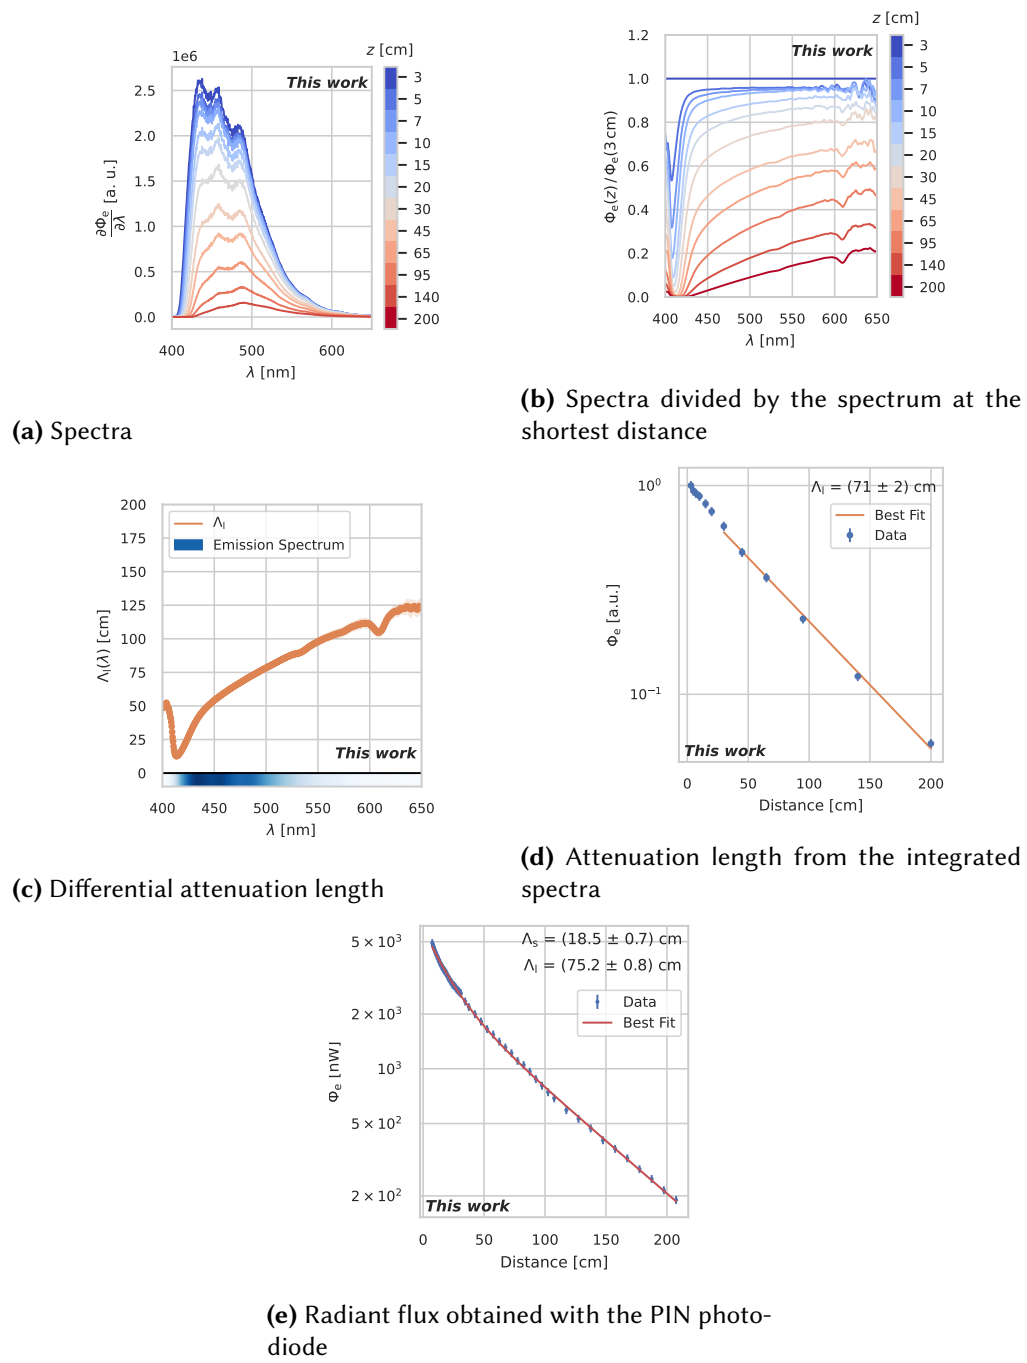





Figure A.13: BCF-12 Fibre 2




(a) Spectra




(b) Spectra divided by the spectrum at the shortest distance



(c) Differential attenuation length



(d) Attenuation length from the integrated spectra



(e) Radiant flux obtained with the PIN photodiode

Figure A.14: BCF-12 Fibre 3

Square multi-clad, green-emitting, scintillating fibre with a diameter of 250 μm .

A.7.1 FIBRE 1

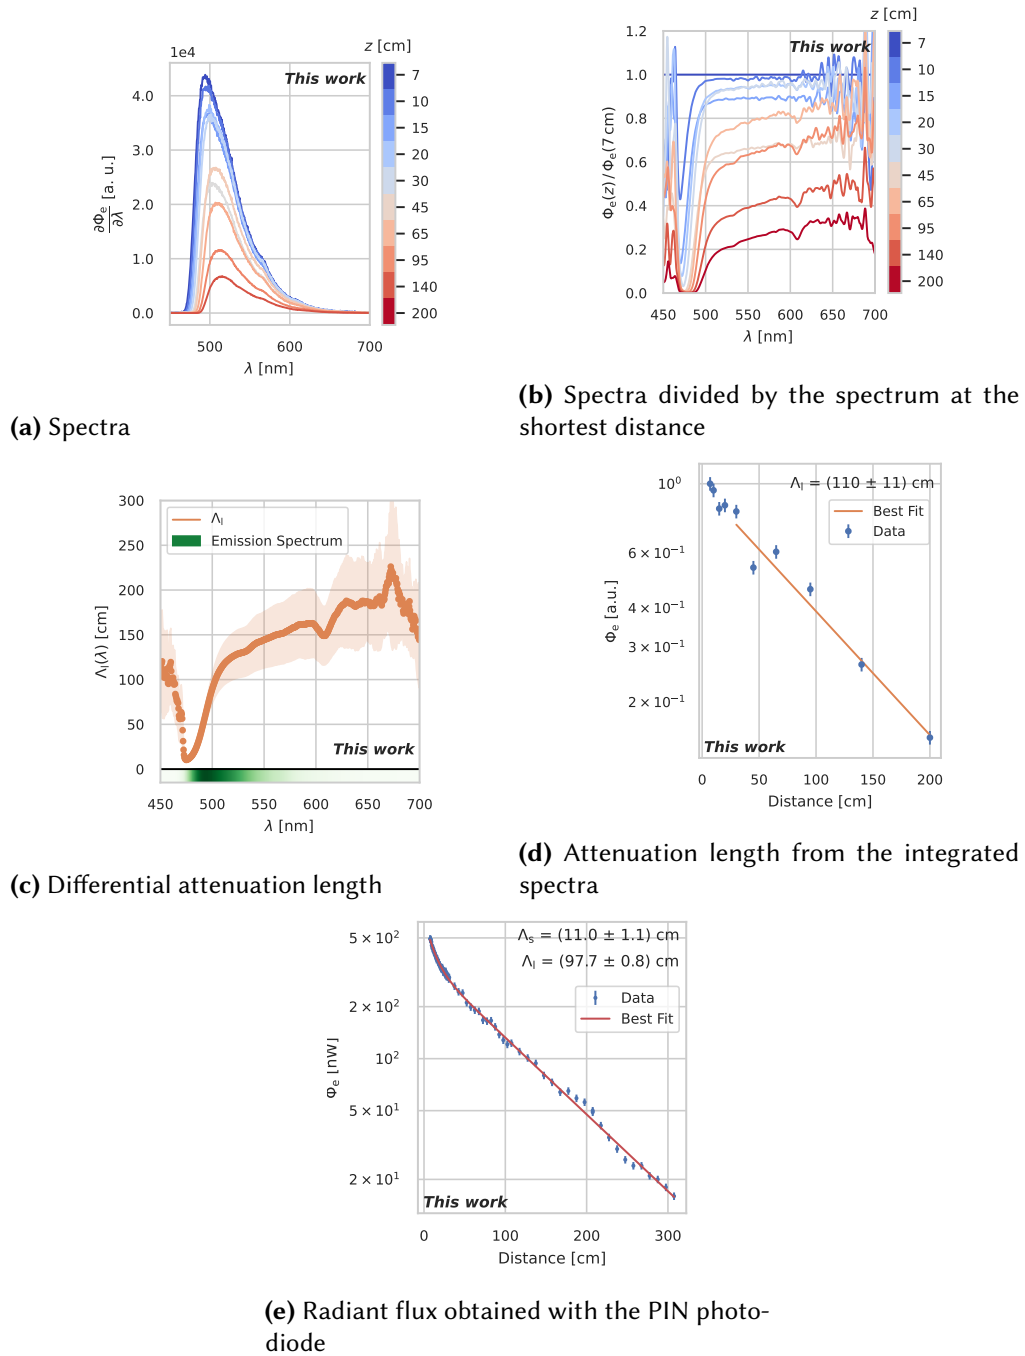


Figure A.15: BCF-20 Fibre 1

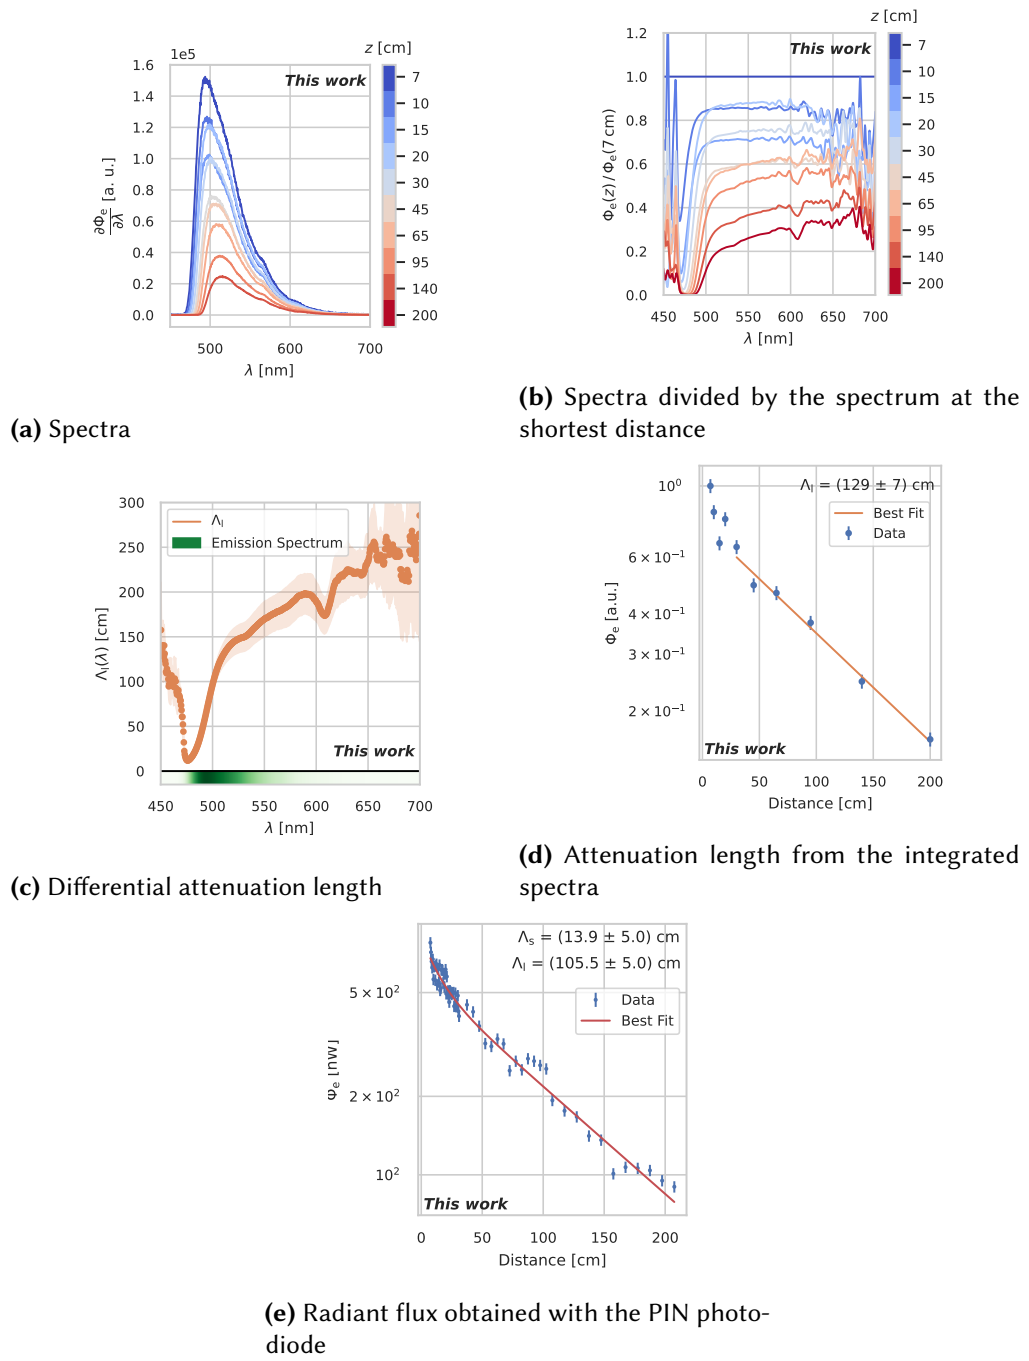


Figure A.16: BCF-20 Fibre 2

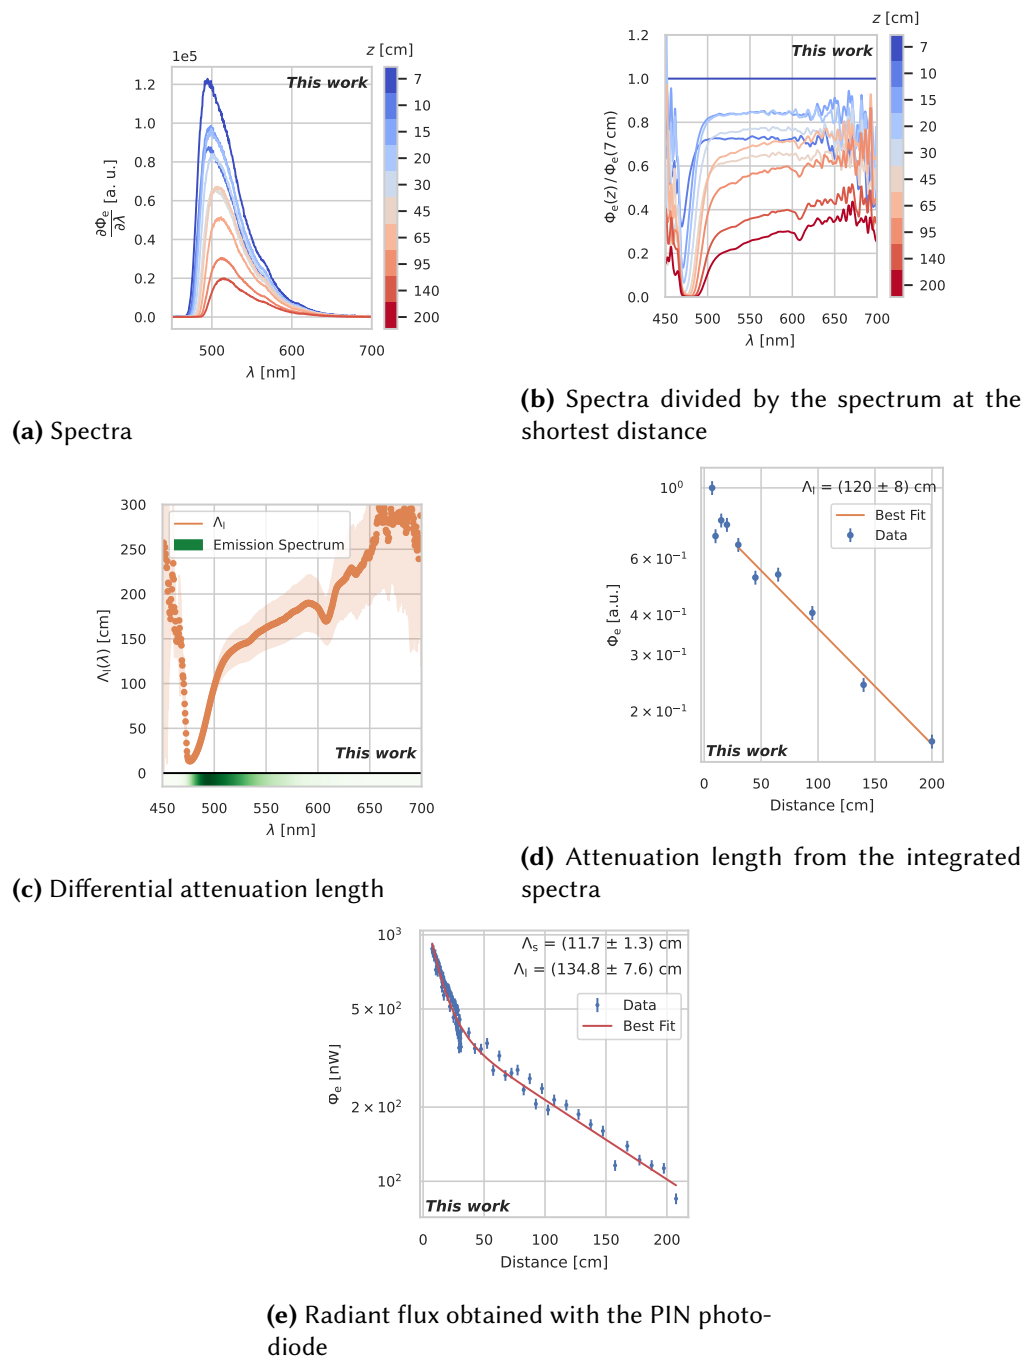


Figure A.17: BCF-20 Fibre 3

Square multi-clad, green-emitting, wavelength-shifting fibre with a diameter of 250 μm .

A.8.1 FIBRE 1

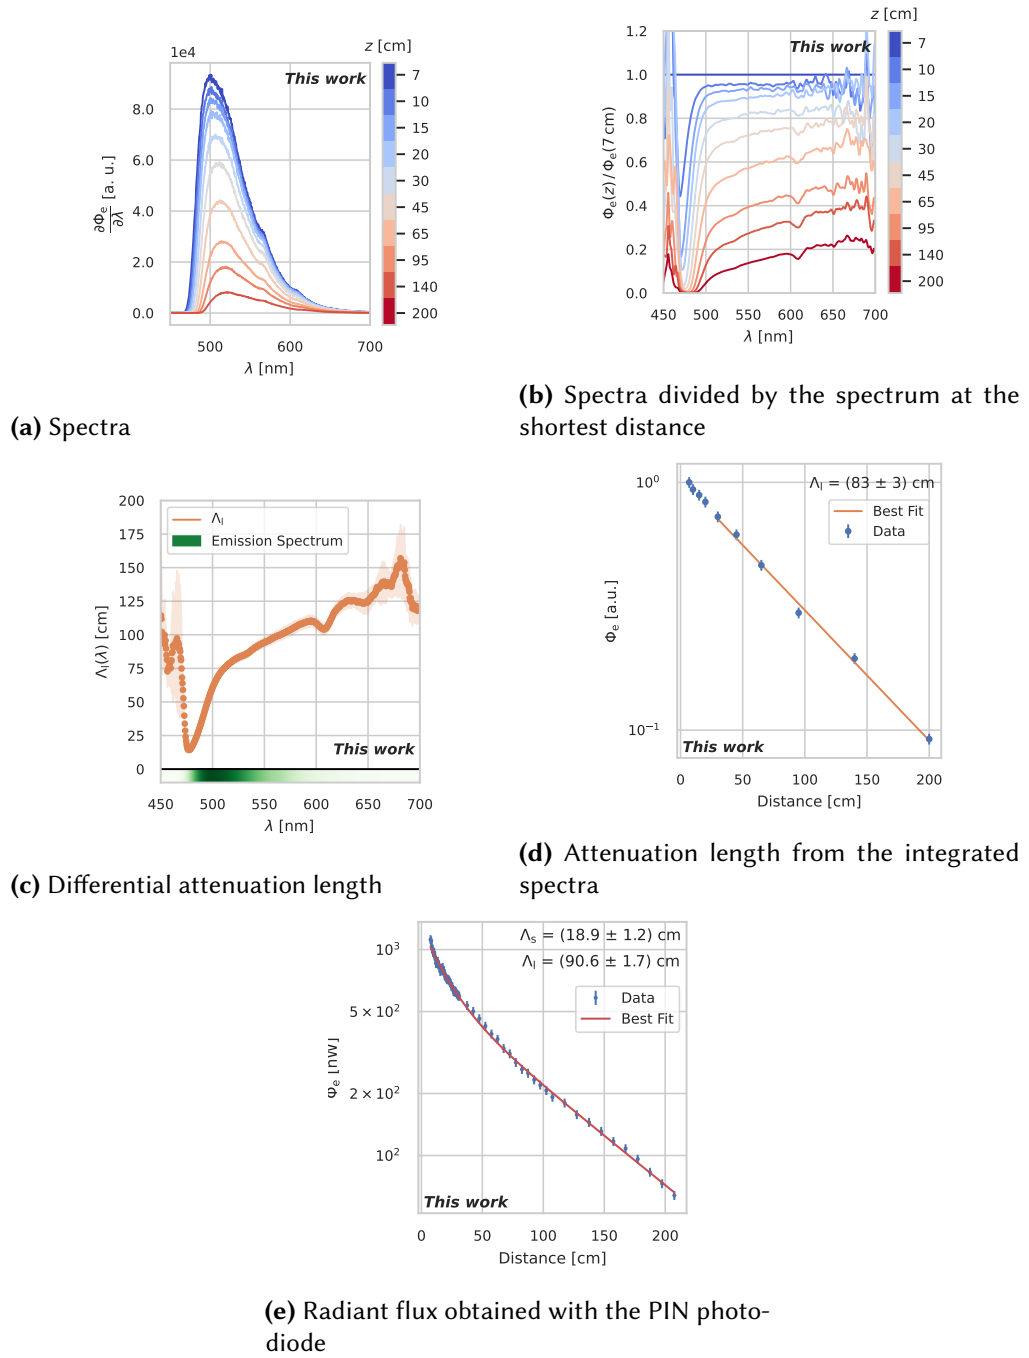


Figure A.18: BCF-92 Fibre 1

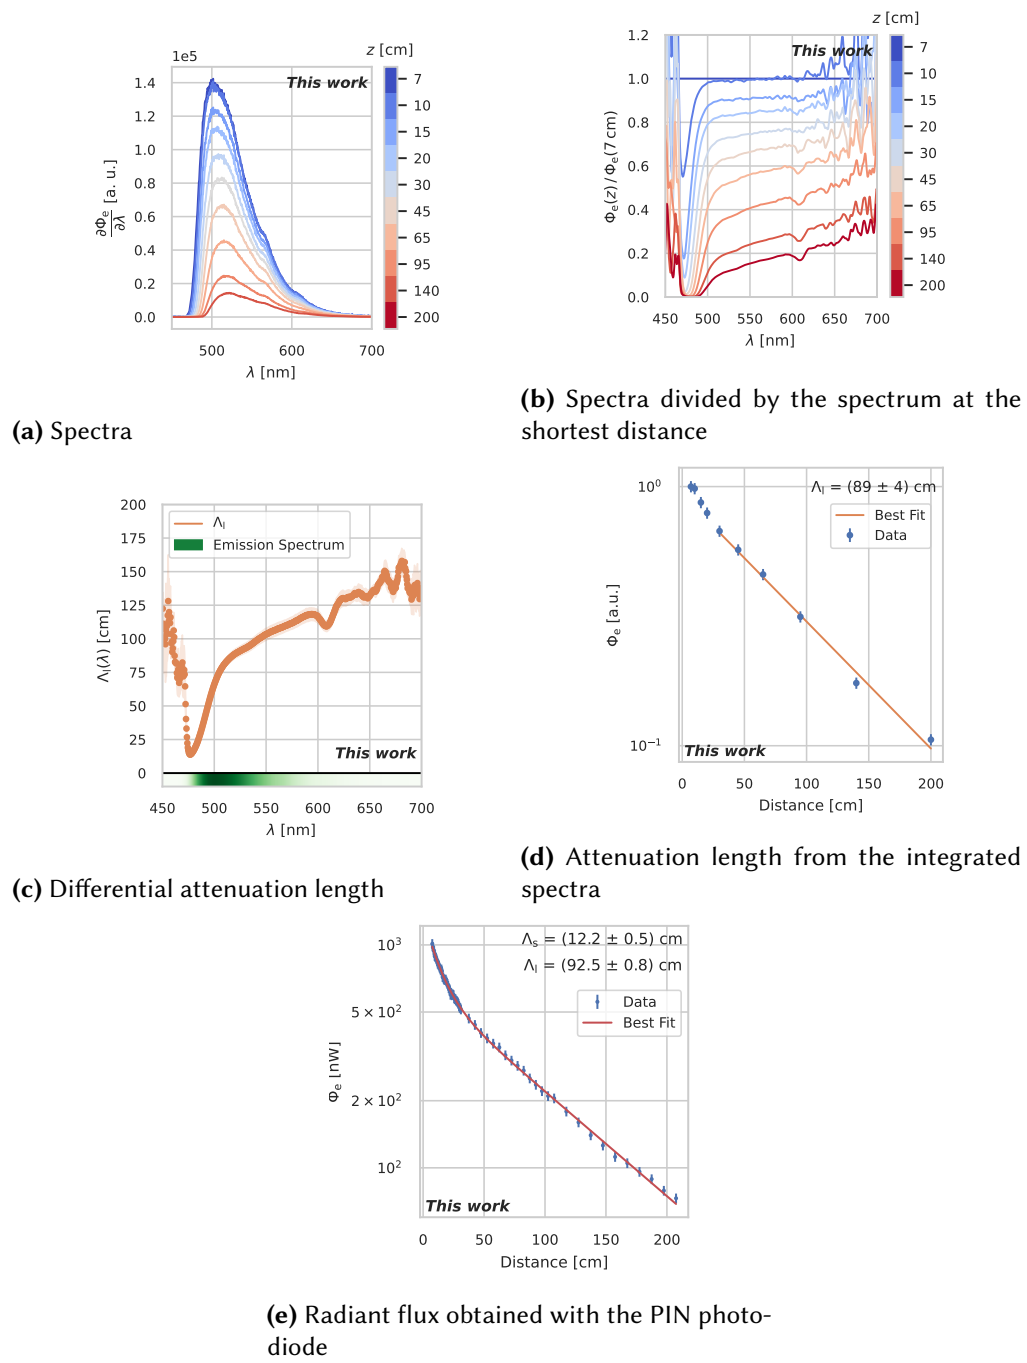





Figure A.19: BCF-92 Fibre 2




(a) Spectra




(b) Spectra divided by the spectrum at the shortest distance



(c) Differential attenuation length



(d) Attenuation length from the integrated spectra



(e) Radiant flux obtained with the PIN photodiode

Figure A.20: BCF-92 Fibre 3

B

Irradiation

B.1 SIMULATION RESULTS FOR MU3E

B.1.1 TOTAL FLUENCE

Table B.1: Simulated total lepton fluence per scintillating fibre SiPM array for $2.5 \cdot 10^{15}$ muon stops on target. The equivalent fluence of neutrons with an energy of 1 MeV is calculated under the NIEL hypothesis with the help of [98].

SiPM	$\Phi_{\text{eq}} [10^{11} \text{ cm}^{-2}]$	$\Phi_{e^+} [10^{12} \text{ cm}^{-2}]$	$\Phi_{e^-} [10^{11} \text{ cm}^{-2}]$
US 0	2.536 ± 0.016	5.000 ± 0.031	5.837 ± 0.089
US 1	2.650 ± 0.017	5.165 ± 0.033	5.876 ± 0.090
US 2	2.662 ± 0.016	5.228 ± 0.032	5.734 ± 0.085
US 3	2.780 ± 0.017	5.450 ± 0.034	6.330 ± 0.091
US 4	2.706 ± 0.016	5.342 ± 0.032	5.824 ± 0.086
US 5	2.789 ± 0.017	5.481 ± 0.033	6.083 ± 0.090
US 6	2.659 ± 0.016	5.237 ± 0.032	5.845 ± 0.083
US 7	2.708 ± 0.017	5.311 ± 0.033	5.744 ± 0.086
US 8	2.619 ± 0.016	5.156 ± 0.031	6.028 ± 0.087
US 9	2.805 ± 0.017	5.492 ± 0.034	5.861 ± 0.088
US 10	2.628 ± 0.016	5.216 ± 0.032	6.018 ± 0.087
US 11	2.799 ± 0.017	5.478 ± 0.033	5.871 ± 0.093
DS 0	1.912 ± 0.014	3.905 ± 0.028	3.551 ± 0.065
DS 1	1.778 ± 0.014	3.663 ± 0.027	3.310 ± 0.063
DS 2	1.917 ± 0.014	3.902 ± 0.028	3.476 ± 0.065
DS 3	1.772 ± 0.014	3.657 ± 0.027	3.522 ± 0.065
DS 4	1.978 ± 0.014	4.040 ± 0.028	3.584 ± 0.065
DS 5	1.793 ± 0.014	3.727 ± 0.027	3.305 ± 0.061
DS 6	1.972 ± 0.014	4.029 ± 0.029	3.537 ± 0.063
DS 7	1.779 ± 0.014	3.654 ± 0.027	3.420 ± 0.065
DS 8	1.940 ± 0.014	3.958 ± 0.028	3.716 ± 0.066
DS 9	1.793 ± 0.014	3.691 ± 0.027	3.394 ± 0.060
DS 10	1.987 ± 0.015	4.070 ± 0.029	3.467 ± 0.065
DS 11	1.824 ± 0.014	3.758 ± 0.027	3.263 ± 0.061

B.1.2 NEUTRON-EQUIVALENT FLUENCE BY DECAY LOCATION

Table B.2: Simulated neutron-equivalent fluence due to Michel by decay location of the muon. Positrons coming from the target region are split in incoming and outgoing particles, defined by the scalar product of the momentum and location vectors. Most muons decaying outside of the target region decay inside the beampipe

SiPM	$\Phi_{\text{eq}} [10^{10} \text{ cm}^{-2}]$ (not from target)	$\Phi_{\text{eq}} [10^{10} \text{ cm}^{-2}]$ (from target, incoming)	$\Phi_{\text{eq}} [10^{10} \text{ cm}^{-2}]$ (from target, outgoing)
US 0	$(1.073 \pm 0.010) \cdot 10^1$	6.515 ± 0.090	5.760 ± 0.067
US 1	$(1.018 \pm 0.011) \cdot 10^1$	6.603 ± 0.093	7.323 ± 0.069
US 2	$(1.145 \pm 0.011) \cdot 10^1$	6.639 ± 0.091	6.243 ± 0.070
US 3	$(1.156 \pm 0.012) \cdot 10^1$	6.785 ± 0.093	6.931 ± 0.067
US 4	$(1.194 \pm 0.011) \cdot 10^1$	6.631 ± 0.090	6.145 ± 0.068
US 5	$(1.103 \pm 0.012) \cdot 10^1$	7.023 ± 0.096	7.329 ± 0.070
US 6	$(1.114 \pm 0.011) \cdot 10^1$	6.782 ± 0.092	6.270 ± 0.068
US 7	$(1.049 \pm 0.011) \cdot 10^1$	6.900 ± 0.095	7.321 ± 0.071
US 8	$(1.120 \pm 0.010) \cdot 10^1$	6.638 ± 0.090	5.881 ± 0.066
US 9	$(1.160 \pm 0.012) \cdot 10^1$	6.828 ± 0.094	7.178 ± 0.070
US 10	$(1.133 \pm 0.011) \cdot 10^1$	6.518 ± 0.090	5.979 ± 0.069
US 11	$(1.121 \pm 0.012) \cdot 10^1$	6.678 ± 0.092	7.620 ± 0.071
DS 0	2.994 ± 0.048	8.19 ± 0.11	6.096 ± 0.072
DS 1	2.851 ± 0.048	7.93 ± 0.10	5.296 ± 0.068
DS 2	3.037 ± 0.049	8.15 ± 0.11	6.226 ± 0.073
DS 3	2.956 ± 0.049	8.07 ± 0.10	4.953 ± 0.067
DS 4	3.198 ± 0.051	8.24 ± 0.11	6.436 ± 0.074
DS 5	2.877 ± 0.048	8.13 ± 0.10	5.212 ± 0.068
DS 6	3.097 ± 0.050	8.37 ± 0.11	6.441 ± 0.075
DS 7	2.919 ± 0.048	8.07 ± 0.10	5.071 ± 0.067
DS 8	3.154 ± 0.050	8.12 ± 0.11	6.257 ± 0.073
DS 9	3.022 ± 0.050	7.99 ± 0.10	5.202 ± 0.068
DS 10	3.086 ± 0.050	8.67 ± 0.11	6.262 ± 0.073
DS 11	2.909 ± 0.048	8.11 ± 0.10	5.573 ± 0.070

B.1.3 NEUTRON-EQUIVALENT FLUENCE BY PARTICLE ORIGIN

Table B.3: Simulated neutron-equivalent fluence for $2.5 \cdot 10^{15}$ muon stops on target, grouped by the origin of the particles. The lower digit marks the particle type (1: positron, 2: electron) and the higher digit marks the process (1: Michel decay, 2: radiative muon decay, 4: photon conversion, 5: Bhabha scattering, 8: Compton scattering)

SIPM	$\Phi_{\text{eq}}^{11} [10^{11} \text{ cm}^{-2}]$	$\Phi_{\text{eq}}^{21} [10^9 \text{ cm}^{-2}]$	$\Phi_{\text{eq}}^{52} [10^9 \text{ cm}^{-2}]$	$\Phi_{\text{eq}}^{82} [10^9 \text{ cm}^{-2}]$	$\Phi_{\text{eq}}^{42} [10^9 \text{ cm}^{-2}]$	$\Phi_{\text{eq}}^{41} [10^9 \text{ cm}^{-2}]$
US 0	2.301 ± 0.015	7.08 ± 0.27	6.78 ± 0.21	4.94 ± 0.18	2.34 ± 0.12	2.35 ± 0.14
US 1	2.412 ± 0.016	7.03 ± 0.26	6.83 ± 0.22	5.01 ± 0.17	2.68 ± 0.15	2.30 ± 0.14
US 2	2.433 ± 0.016	6.67 ± 0.25	6.55 ± 0.21	4.93 ± 0.17	2.34 ± 0.12	2.32 ± 0.12
US 3	2.528 ± 0.017	7.57 ± 0.28	6.90 ± 0.21	5.61 ± 0.18	2.57 ± 0.13	2.65 ± 0.13
US 4	2.471 ± 0.016	7.17 ± 0.26	6.99 ± 0.22	5.01 ± 0.17	2.11 ± 0.11	2.17 ± 0.13
US 5	2.540 ± 0.017	7.63 ± 0.28	7.36 ± 0.22	4.92 ± 0.17	2.54 ± 0.14	2.56 ± 0.14
US 6	2.420 ± 0.016	7.77 ± 0.27	6.49 ± 0.19	4.86 ± 0.16	2.41 ± 0.13	2.39 ± 0.12
US 7	2.472 ± 0.016	7.53 ± 0.28	6.61 ± 0.20	4.75 ± 0.17	2.44 ± 0.14	2.36 ± 0.14
US 8	2.372 ± 0.015	7.33 ± 0.26	6.75 ± 0.21	5.18 ± 0.17	2.84 ± 0.14	2.56 ± 0.14
US 9	2.562 ± 0.017	7.65 ± 0.28	6.99 ± 0.22	4.74 ± 0.16	2.50 ± 0.14	2.50 ± 0.14
US 10	2.382 ± 0.015	7.31 ± 0.26	7.03 ± 0.22	4.97 ± 0.16	2.53 ± 0.13	2.69 ± 0.15
US 11	2.552 ± 0.017	7.71 ± 0.29	6.70 ± 0.23	5.25 ± 0.19	2.56 ± 0.14	2.54 ± 0.14
DS 0	1.729 ± 0.014	9.06 ± 0.29	4.82 ± 0.17	2.37 ± 0.12	1.166 ± 0.091	0.983 ± 0.091
DS 1	1.607 ± 0.013	8.31 ± 0.26	4.56 ± 0.17	2.00 ± 0.11	1.112 ± 0.086	1.035 ± 0.089
DS 2	1.743 ± 0.014	8.36 ± 0.27	5.02 ± 0.18	2.09 ± 0.10	1.113 ± 0.085	0.930 ± 0.081
DS 3	1.598 ± 0.013	8.32 ± 0.27	4.89 ± 0.18	2.165 ± 0.010	1.135 ± 0.095	0.929 ± 0.074
DS 4	1.788 ± 0.014	9.42 ± 0.29	4.86 ± 0.17	2.32 ± 0.11	1.149 ± 0.087	1.235 ± 0.099
DS 5	1.622 ± 0.013	8.60 ± 0.28	4.56 ± 0.16	2.072 ± 0.099	0.908 ± 0.080	1.027 ± 0.092
DS 6	1.792 ± 0.014	8.95 ± 0.28	4.91 ± 0.17	2.25 ± 0.11	0.975 ± 0.076	1.043 ± 0.079
DS 7	1.606 ± 0.013	8.16 ± 0.27	4.78 ± 0.18	2.27 ± 0.11	1.114 ± 0.091	0.944 ± 0.084
DS 8	1.754 ± 0.014	9.22 ± 0.29	5.11 ± 0.18	2.28 ± 0.11	1.132 ± 0.088	0.975 ± 0.082
DS 9	1.622 ± 0.013	8.40 ± 0.27	4.70 ± 0.16	2.084 ± 0.098	9.23 ± 0.75	0.926 ± 0.085
DS 10	1.803 ± 0.014	9.22 ± 0.28	5.10 ± 0.18	2.05 ± 0.11	9.99 ± 0.83	1.070 ± 0.089
DS 11	1.659 ± 0.013	8.12 ± 0.27	4.26 ± 0.16	2.12 ± 0.10	9.49 ± 0.82	0.895 ± 0.067

B.1.4 TOTAL IONIZING DOSE PER SiPM

Table B.4: Simulated total ionizing doses in the bulk silicon and silicon dioxide of each silicon photomultiplier array. The total ionizing dose corresponds to $2.5 \cdot 10^{15}$ muon stops on target

SiPM	D_{Si} [kGy]	D_{SiO_2} [kGy]
US 0	1.473 ± 0.011	1.219 ± 0.077
US 1	1.523 ± 0.011	1.80 ± 0.53
US 2	1.525 ± 0.011	2.21 ± 0.49
US 3	1.607 ± 0.011	2.62 ± 0.88
US 4	1.555 ± 0.011	1.314 ± 0.088
US 5	1.614 ± 0.011	1.148 ± 0.037
US 6	1.533 ± 0.011	1.43 ± 0.16
US 7	1.563 ± 0.011	1.304 ± 0.070
US 8	1.518 ± 0.010	1.53 ± 0.18
US 9	1.606 ± 0.011	2.07 ± 0.72
US 10	1.543 ± 0.011	1.253 ± 0.067
US 11	1.615 ± 0.012	1.211 ± 0.053
DS 0	1.1090 ± 0.0090	0.882 ± 0.059
DS 1	1.0499 ± 0.0088	0.816 ± 0.043
DS 2	1.1218 ± 0.0094	0.989 ± 0.080
DS 3	1.0530 ± 0.0090	0.881 ± 0.091
DS 4	1.1607 ± 0.0094	1.52 ± 0.58
DS 5	1.0732 ± 0.0093	0.881 ± 0.067
DS 6	1.1464 ± 0.0092	0.884 ± 0.046
DS 7	1.0545 ± 0.0092	1.33 ± 0.40
DS 8	1.1563 ± 0.0097	1.21 ± 0.28
DS 9	1.0592 ± 0.0090	8.93 ± 0.90
DS 10	1.1661 ± 0.0095	0.946 ± 0.050
DS 11	1.0767 ± 0.0090	0.972 ± 0.067

B.2 TIME STAMPS FOR SiPM SAMPLES

Table B.5: Insertion and extraction times for all SiPMs in Unix time. The *side* refers to the two possible positions inside the setup, *Aare* and *Berg* are landmark references at PSI. The integral of the proton current multiplied by the shutter status S is proportional to the received dose. In addition to the six samples, a populated PCB missing the SiPM was irradiated as a control sample

SiPM	Time stamp in	Time stamp out	Side	$\int S I_p dt$ [mAh]
1	1596814221	1597402808	Aare	263.2
2	1596814221	1597402808	Berg	263.2
3	1597759202	1598022002	Aare	88.9
4	1597759202	1597849202	Berg	27.6
5	1597849202	1598022002	Berg	61.3
6	1598022002	1598263201	Berg	98.2
control	1598022002	1598263201	Aare	98.2

B.3 IV CURVES

In the following, per-channel IV curves for the SiPMs arrays in chapter 5 are presented. Each line represents one channel.




Figure B.1: IV curves for SiPM 1 before (blue) and after irradiation (red).




Figure B.2: IV curves for SiPM 2 before (blue) and after irradiation (red).




Figure B.3: IV curves for SiPM 3 before (blue) and after irradiation (red).




Figure B.4: IV curves for SiPM 4 before (blue) and after irradiation (red).




Figure B.5: IV curves for SiPM 5 before (blue) and after irradiation (red).




Figure B.6: IV curves for SiPM 6 before (blue) and after irradiation (red).






Figure B.7: IV curves for SiPM 7 before (blue) and “after irradiation” (red). This SiPM was treated like the others, but was not irradiated and serves as a control sample.

B.4 CURRENT CHANGES

B.4.1 SiPM 1




(a) $V = V_{BD} + 3.5 \text{ V}$




(b) $V = 50 \text{ V} < V_{BD}$

Figure B.8: Per-channel ratio of the current in SiPM 1 before and after irradiation. The top shows the operational voltage recommended by the vendor, the bottom shows a fixed voltage below breakdown.

B.4.2 SiPM 2




(a) $V = V_{BD} + 3.5 \text{ V}$




(b) $V = 50 \text{ V} < V_{BD}$

Figure B.9: Per-channel ratio of the current in SiPM 2 before and after irradiation. The top shows the operational voltage recommended by the vendor, the bottom shows a fixed voltage below breakdown.

B.4.3 SiPM 3




(a) $V = V_{BD} + 3.5 \text{ V}$




(b) $V = 50 \text{ V} < V_{BD}$

Figure B.10: Per-channel ratio of the current in SiPM 3 before and after irradiation. The top shows the operational voltage recommended by the vendor, the bottom shows a fixed voltage below breakdown.

B.4.4 SiPM 4




(a) $V = V_{BD} + 3.5 \text{ V}$




(b) $V = 50 \text{ V} < V_{BD}$

Figure B.11: Per-channel ratio of the current in SiPM 4 before and after irradiation. The top shows the operational voltage recommended by the vendor, the bottom shows a fixed voltage below breakdown.

B.4.5 SiPM 5




(a) $V = V_{BD} + 3.5 \text{ V}$




(b) $V = 50 \text{ V} < V_{BD}$

Figure B.12: Per-channel ratio of the current in SiPM 5 before and after irradiation. The top shows the operational voltage recommended by the vendor, the bottom shows a fixed voltage below breakdown.

B.4.6 SiPM 6




(a) $V = V_{\text{BD}} + 3.5 \text{ V}$




(b) $V = 50 \text{ V} < V_{\text{BD}}$

Figure B.13: Per-channel ratio of the current in SiPM 6 before and after irradiation. The top shows the operational voltage recommended by the vendor, the bottom shows a fixed voltage below breakdown.

B.4.7 SiPM 7



(a) $V = V_{BD} + 3.5 \text{ V}$




(b) $V = 50 \text{ V} < V_{BD}$

Figure B.14: Per-channel ratio of the current in SiPM 7 before and after irradiation. This SiPM is not irradiated and is used as a control sample. The top shows the operational voltage recommended by the vendor, the bottom shows a fixed voltage below breakdown.


B.5 BREAKDOWN VOLTAGE

In this section, the breakdown voltages for each SiPM is presented per-channel using both the relative derivative (section 5.5.3) and tangent (section 5.5.3) models.

B.5.1 SiPM 1




(a) Relative derivative method




(b) Tangent method

Figure B.15: Breakdown voltage of each channel in SiPM 1

B.5.2 SiPM 2




(a) Relative derivative method




(b) Tangent method

Figure B.16: Breakdown voltage of each channel in SiPM 2

B.5.3 SiPM 3




(a) Relative derivative method




(b) Tangent method

Figure B.17: Breakdown voltage of each channel in SiPM 3

B.5.4 SiPM 4




(a) Relative derivative method




(b) Tangent method

Figure B.18: Breakdown voltage of each channel in SiPM 4

B.5.5 SiPM 5




(a) Relative derivative method




(b) Tangent method

Figure B.19: Breakdown voltage of each channel in SiPM 5

B.5.6 SiPM 6




(a) Relative derivative method




(b) Tangent method

Figure B.20: Breakdown voltage of each channel in SiPM 6

B.5.7 SiPM 7




(a) Relative derivative method




(b) Tangent method

Figure B.21: Breakdown voltage of each channel in SiPM 7


B.6 CURRENT DURING IRRADIATION




(a) SiPM 1




(b) SiPM 2




(c) SiPM 3



(d) SiPM 4



(e) SiPM 5



(f)

Figure B.22



MuTRiG Details

C.1 MuTRiG CONFIGURATION SETTINGS

The following tables are based on [40], tooltips in the Mu3e DAQ software written by Konrad Briggel, code documentation by the KIP SciTile group and discussions with Yonathan Munwes. This reference is written for MuTRiG 2, although only minor changes are expected for MuTRiG 3. Variable names are printed as used in August 2022. See also <https://bitbucket.org/mu3e/online/wiki/MuTRiG%20Settings>

Table C.1: DAQ Settings, no direct effect on the MuTRiG configuration pattern.

Name in web config	Description	Comments
DUMMY CONFIG	If enabled, ignore different replies from ASICs when configuring via SPI	
DUMMY DATA	Enable dummy data generator on FPGA	
DUMMY FAST	Produce short events in dummy data generator on FPGA	
DUMMY NUM	Number of hits to be generated per ASIC using the FPGAs dummy data generator	
PRBS decoder disable	Bypass PRBS decoder on the FPGA	
Wait for all	Data receivers are only allowed to produce frames when all enabled receivers are locked. Helps to keep synchronization and avoid buffer overruns at start of a run.	
Wait for all sticky	Once a receiver has locked, do not stop producing data even if the link is lost. Ready bits are reset at start of run.	
Reset Blocks	Reset various parts of the DAQ chain.	Triggers the corresponding RPC in the switching frontend.

Table C.2: Global Settings. These settings are used for all MuTRiGs.

Name in web config	Description	Comments
Num Asics	Number of ASICs	
EXT TRG OFFSET	Configuration of the external trigger.	
EXT TRG ENDT	End time of the trigger matching window.	
DEBUG		
PRBS_SINGLE	Generate only one PRBS event per frame.	In the PRBS_DEBUG mode, if PRBS_SINGLE is selected, there will be only one 48-bit hit in each data frame.
EXT TRG	External trigger mode.	
EXT TRG ENDT SIGN	Sign of the external trigger matching window.	1: end time is after the trigger; 0: end time is before the trigger
GEN IDLE	Enables K28.5 comma word in LVDS data stream	While no data is sent, K28.5 is constantly repeated, this helps with clock recovery and debugging.
DIS COARSE SHORT	Enables the short event mode, only T-timestamps and E-flags are transmitted.	Other names: Fast mode, SciFi mode
PRBS	Enables the generation of pseudorandom bit sequence in the digital part of MuTRiG.	The output data will be 48-bit PRBS data. If PRBS_SINGLE is not selected, there will be 255 PRBS <i>hits</i> in each frame.
SYNC_CH_RST	Synchronous PLL reset	1: Use internally synchronized (2FF-synchronizer) version of the reset signal for PLL reset. '0': Feed reset input to pll directly.
PLL COARSE		
PLL ENVMON		

Table C.3: Settings per single MuTRiG.

Name in web config	Description	Comments
PLL Settings		
VNPFC		
VNCnt		
VNVCOBuffer		
VND2C		
VNPCP		
VNHitLogic		
VNCntBuffer		
VNVCODelay		
LATCHBIAS		
MS Limits		
MS Polarity		
Monitor Driver Settings		Not accessible on SMB2
AMON	Bias current for the analog monitor.	
DMON 1	Bias current and enable for Q/T digital monitor. Higher value means higher current.	Default value: 220
DMON 2	Bias current and enable for flag/E digital monitor. Higher value means higher current.	Default value: 220
LVDS Driver Settings		
LVDS VCM	LVDS driver common mode voltage.	A typical value is 155, corresponding to 1.1 V (0: 0 V, 255: 1.8 V)
LVDS BIAS	Additional bias for the LVDS transmitter.	The LVDS is already biased with a tail current of 4.5 mA. This DAC is used in case the LVDS TX need higher driving power. 0: 0 A, 63: 4.5 A. The default value is 0.

Table C.4: Settings per single channel.

Name in web config	Description	Comments
TTHRESH	Timing threshold setting.	Range is 0 to 63. Lower values correspond to higher thresholds, a threshold of zero (noise) is around a value of 32.
TTHRESH Scale	Timing threshold current scaling.	A value of 0 corresponds to 1:1 scaling, 4 corresponds to 1:2 (1:2 means higher maximum threshold). The dynamic range at is maximized for a value of 0 [40].
ETHRESH	Energy threshold, see above.	A value of 0 corresponds to a threshold of 1.8 V; 255 corresponds to 0 V.
MASK	Mask a channel	
SIPM	Bias voltage at the SiPM input terminal.	Shifts the potential of the SiPM anode, thus changing the bias voltage.
INPUTBIAS	Bias current at the input stage.	Higher values mean higher current.
POLE	Pole-zero cancellation.	
EBIAS	Discharge current for the energy signal.	A lower current produces a higher time-over-threshold.
AMP COM	Tuning the bias current of amplifiers, to change the gain of the amplifier and the hysteresis of the comparator.	The higher the value, the larger the gain and the smaller the hysteresis.
CML	Current Mode Logic	
AMON CTRL	Analog monitor selection	
COMP SPI	Change the tail compensation effect.	
TDCTEST_N	Disable TDC test input to this channel	
S SWITCH	Single-ended or differential signals	In Mu3e, the SiPMs are connected single-ended.
DELAY		
ENERGY_C_EN	Enable the capacitor in the low-pass filter before the E-Comparator.	
ENERGY_R_EN	Enable the Resistor in the low-pass filter before the E-Comparator.	
SENSING_HIGH_R	Enables the high impedance source in the feedback path.	
EDGE	Select polarity of E-branch comparator	See also EDGE CML
EDGE CML	Invert polarity of digital signals to hitlogic in channel.	
DMON ENA	Enable digital monitor for this channel.	Should only be enabled for one channel per ASIC. Cannot be read out with the SciFi Module Board.
DMON SW	Digital monitor output mode.	0: Show Q and flag output (as it goes to the TDC), 1: show raw T & E flags
RECV ALL	Send the hit structure of “long” events, but the <i>Energy timestamp</i> is actually the next T timestamp	
POLE_EN_N	0: enable dac for pole zero cancellation.	Only controls the dac, not the feature itself
AMON_EN_N	1: disable analog monitor output from this channel.	The analog monitor should only be enabled for up to one channel in a chip.

Table C.5: Coincidence matrix settings. Not accessible via the web frontend. Lower and upper crossbar refer to the connections to neighbouring chips at the respective end of the chip.

Name	Description	Comments
COIN_WND	Coincidence window	0: 2 coarse counters, 1: three coarse counters
XBAR_lower_TX_ena	Enable transmission for the lower crossbar	
XBAR_lower_RX_ena	Enable reception for the lower crossbar	
XBAR_lower_TX_I	Additional current.	Nominal value at 0, can increase current.
XBAR_lower_TX_V	Common mode voltage.	From 0 to 1.8 V, nominal value of 1.1 V at 155
XBAR_upper_TX_ena	Enable transmission for the upper crossbar	
XBAR_upper_RX_ena	Enable reception for the upper crossbar	
XBAR_upper_TX_I	Additional current.	Nominal value at 0, can increase current.
XBAR_upper_TX_V	Common mode voltage.	From 0 to 1.8 V, nominal value of 1.1 V at 155
COIN_MAT_XBL	Channels whose hit information is sent to the neighbouring MuTRiG	3 bit, for channels 0, 1, and 2. The distance information is lost and hits are treated as a single channel 32 on the other MuTRiG
COIN_MAT_XBH	Channels whose hit information is sent to the neighbouring MuTRiG	3 bit, for channels 29, 30, and 31. The distance information is lost and hits are treated as a single channel -1 on the other MuTRiG
COIN_MAT_CHn	Channels whose hits validate hits in channel n	6 bit, three closest neighbours in either direction. Enabling channels < -1 and > 32 has no effect.

Table C.6: Configuration byte map for MuTRiG 2. Data are sent from top left to bottom right. The numbers in parentheses denote bit indices of a parameter. A parameter that is set to 8 would be represented as 1000, with a one at index (3) and zeroes at indices (2:0). Scale parameters are denoted by the suffix *sc*, and “enable” flags by the suffix *en*.

7	6	5	4	3	2	1	0
	LVDS BIAS (0:5)						LVDS VCM (0)
LVDS VCM (1:6)						DMON 2 (0)	
DMON 2 (1:6)						DMON 2 en	
DMON 1 (0:7)							
DMON 1 en	AMON (0:6)						
AMON (7)	AMON_EN_N	COIN_WND	XBAR_upper_TX_I (0:4)				
XBAR_upper_TX_I (5)	XBAR_upper_TX_V (0:6)						
XBAR_upper_TX_V (7)	XBAR_upper_TX_ena	XBAR_upper_RX_ena	COIN_MAT_XBH (0:2)			COIN_MAT_CH31 (0:1)	
COIN_MAT_CH31 (2:5)				COIN_MAT_CH30 (0:3)			
⋮							
COIN_MAT_CH0 (0:5)					COIN_MAT_XBL (0:1)		
COIN_MAT_XBL (2)	XBAR_lower_TX_I (0:5)						XBAR_lower_TX_V (0)
XBAR_lower_TX_V (1:7)						XBAR_lower_TX_ena	
XBAR_lower_RX_ena	LATCHBIAS (11:5)						
LATCHBIAS (4:0)				VNPFC (0:2)			
VNPFC (3:5)			VNPFC Offset (0:1)		VNPFC Scale	VNHitLogic (0:1)	
VNHitLogic (2:5)				VNHitLogic Offset		VNHitLogic Scale	VNVCOBuffer (0)
VNVCOBuffer (1:5)					VNVCOBuffer Offset (0:1)		VNVCOBuffer Scale
VNVCODelay (5:0)						VNVCODelay Offset (0:1)	
VNVCODelay Scale	VNPCP (0:5)						VNPCP Offset (0)
VNPCP Offset (1)	VNPCP Scale	VNCnt (0:5)					

VNCnt OFFSET (0:1)		VNCnt SCALE	VNCntBuffer (0:4)				
VNCntBuffer (5)	VNCntBuffer OFFSET		VNCntBuffer SCALE	VND2C (0:3)			
VND2C (4:5)		VND2C offset (0:1)	VND2C scale	RECV_ALL_CH31	CHANNEL_MASK_CH31	POLE_EN_N_CH31	
DELAY_CH31	CML_CH31 (0:3)			POLE_CH31 (0:2)			
POLE_CH31 (3:5)		POLE_SC_CH31	EBIAS_CH31 (0:1)		ETHRESH_CH31 (0)		
ETHRESH_CH31 (1:7)					INPUTBIAS_CH31 (0)		
INPUTBIAS_CH31 (1:5)				INPUTBIAS_SC_CH31	AMPCOM_CH31 (0:1)		
AMPCOM_CH31 (2:5)			AMPCOM_SC_CH31 (0:1)		TTHRESH_CH31 (0:1)		
TTHRESH_CH31 (4:5)			TTHRESH_SC_CH31 (0:2)		SIPM_CH31 (0)		
SIPM_CH31 (4:5)				SIPM_SC_CH31	COMP_SPI_CH31 (0:1)		
AMON_CTRL_CH31 (0)	AMON_CTRL_CH31 (1:2)		TDCTEST_N_CH31	DMON SW CH31	DMON ENA CH31	CMLSCALE_CH31	EDGE_CML_CH31
EDGE_CH31	AMON_CH31	SENSING_HIGH_R_CH31	S_SWITCH_CH31	ENERGY_R_CH31	ENERGY_C_CH31	RECV_ALL_CH30	CHANNEL_MASK_CH30

⋮

SENSING_HIGH_R_CH0	S_SWITCH_CH0	ENERGY_R_CH0	ENERGY_C_CH0	RECV_ALL_CH0	CHANNEL_MASK_CH0	POLE_EN_N_CH0	DELAY_CH0
CML_CH0 (3:0)			POLE_CH0 (3:0)				
POLE_CH0 (5:4)		POLE_SC_CH0	EBIAS_CH0 (2:0)		ETHRESH_CH0 (1:0)		
ETHRESH_CH0 (7:2)					INPUTBIAS_CH0 (1:0)		
INPUTBIAS_CH0 (5:2)				INPUTBIAS_SC_CH0	AMPCOM_CH0 (2:0)		
AMPCOM_CH0 (5:3)			AMPCOM_SC_CH0 (1:0)		TTHRESH_CH0 (2:0)		
TTHRESH_CH0 (5:3)			TTHRESH_SC_CH0 (2:0)		SIPM_CH0 (1:0)		
SIPM_CH0 (5:2)				SIPM_SC_CH0	COMP_SPI_CH0 (1:0)		AMON_CTRL_CH0 (0)
AMON_CTRL_CH0 (2:1)		TDCTEST_N_CH0	DMON SW CH0	DMON ENA CH0	CMLSCALE_CH0	EDGE_CML_CH0	EDGE_CH0
AMON_CH0	SENSING_HIGH_R_CH0	S_SWITCH_CH0	ENERGY_R_CH0	ENERGY_C_CH0	Channel Event Counter	PLL_LOL_dbg	DIS COARSE
PLL ENVMON	PLL COARSE	SHORT	PRBS SINGLE	PRBS	MS DEBUG	MS SWITCH SEL	MS LIMITS (4:4)
MS LIMITS (0:3)				EXT TRIG ENDT (3:0)			
EXT TRIG OFFSET (n/3:0)				EXT TRIG ENDT SIGN	EXT TRIG	SYNC_CH_RST	GEN IDLE

C.2 MuTRiG CONFIGURATION IN DEPTH

The configuration bit pattern of MuTRiG is sent via an SPI interface on the FEB. While the *chip select* line is active¹, the MOSI line is connected to the beginning of a shift register; the MISO line is connected to the last flip flop. The data are shifted through the register by the SPI *clock*. That shift register is 2358 bits long in MuTRiG 1 and 2719 bits in MuTRiG 2. Only when the *chip select* line becomes inactive, the content changes the actual settings inside the chip.

The design of this shift register allows detecting transmission errors only to a certain degree. After a complete pattern is sent, this pattern is stored in the shift register. By sending further data on the MOSI line, the received pattern is streamed by the MuTRiG on the MISO line. If the communication works, the same pattern can be sent twice, and during the second iteration, the data on both lines should match.

Since the NIOS II SPI interface sends complete bytes, additional caution is required because the length of the shift register is not divisible by eight. Before streaming data, they need to be grouped in octets and padded in the beginning, as those bits are the first ones to *drop out of the shift register*. This needs to be considered when comparing the response to the data that are sent.

C.3 MuTRiG EVENTS

Table C.7: MuTRiG Event data. For short events, energy information is omitted, resulting in 27-bit events, long events have a length of 48 bits. The length is denoted in number of bits.

	Channel No.	Bad	time Coarse	Fine	Bad	energy Coarse	Fine	E-Flag
length [bit]	5	1	15	5	1	15	5	1

¹Active when low

C.4 MuTRiG FRAMES

Table C.8: A MuTRiG data frame. The data are 8b/10b encoded, so in addition to 8-bit words, additional symbols are available. The first symbol is an 8b/10b control symbol to indicate the start of a frame, followed by a 16-bit frame ID number, frame flags (see below) and the number of events in this frame (the frame length). The payload consists of the event data, which are followed by a trailer that allows the detection of transmission errors. Idle signals (K.28.5) before and after the frame allow synchronisation and clock detection. The length is stated in bits. CRC refers to cyclic redundancy check.

* The maximum payload length is 7144 b for short events and 12 240 b for long events.

	Header				Payload	Trailer	
	K.28.0	Frame ID	Flags	Length	Events	CRC	K.28.4
length [bit]	8	16	6	10	*	16	8

Table C.9: MuTRiG frame flags. Certain settings can be marked, the generation of K.28.5 symbols as *idle signals* between frames, the event mode (see above), whether PRBS debug “data” is sent and if on-chip memory (L2 FIFO) is full.

Bit position	Flag Definition
5	generate idle signals
4	short event mode
3	PRBS debug mode
2	Only one PRBS event per frame
1	L2 FIFO full
0	unused

Acknowledgments

First and foremost, I would like to thank Christoph Grab and Rainer Wallny for their guidance and continuous support throughout my doctorate in any form necessary. I am grateful to Angela Papa for her support throughout all stages of my doctorate and for agreeing to be my third examiner. I personally, and the Mu3e collaboration owe Frank Meier a great deal of of gratitude for his efforts as our technical coordinator, collaborator, and friend.

My first colleague at ETH Simon Corrodi has taught me a lot and introduced me to lab work and testbeams, thus enabling everything that came afterwards. Additionally, it was a great pleasure to work with my fellow colleagues Antoaneta Damyanova, Yannick Demets, Felix Berg, Konrad Briggel, Cristina Martín Pérez and Yifeng Wang. Aldo Antognini supported me greatly during the design of my irradiation chamber and taught me plenty of CAD. I am grateful to Diogo Di Calafiori for the freedom to work on electronics under his guidance and his lessons that surely took more effort than if he had just built everything on his own.

In addition to their guidance, Angela Papa and Alessandro Bravar made my measurements possible by supplying fibres and SiPMs and letting me characterize them in Geneva and at PSI.

I am grateful to the numerous groups whose outstanding efforts have made my experiments possible, namely Hallendienst, the vacuum, water, electrics and radiation protection groups at PSI, and the ETH workshops. I want to express special gratitude to Thomas Rauber, Manuel Schwarz and Davide Reggiani who helped me leave no stone unturned when things did not work as expected in the beamline.

None of the past years would have been possible without the loving support of my friends and family, for which I am forever grateful.

Part of this work was performed at the High Intensity Proton Accelerator using the Swiss Infrastructure for Particle Physics at the Paul Scherrer Institute, Villigen, Switzerland.

I acknowledge the generous funding support from the Swiss National Science Foundation grant number 200021_172519.

Acronyms

SM	Standard Model of particle physics	VCO	Voltage Controlled Oscillator
LFV	Lepton Flavour Violation	SPI	Serial Peripheral Interface
cLFV	Charged Lepton Flavour Violation	LVDS	Low Voltage Differential Signalling
PSI	Paul Scherrer Institute	CRC	Cyclic Redundancy Check
SES	Single Event Sensitivity	FPGA	Field Programmable Gate Array
NOL	Nanostructured Organosilicon Luminophore	PCIe	Peripheral Component Interconnect Express
SiPM	Silicon Photomultiplier	MIDAS	Maximum Integrated Data Acquisition System
PCB	Printed Circuit Board	RPC	Remote Procedure Call
PLL	Phase Locked Loop	ODB	Online Database
DAQ	Data Acquisition System	LED	Light Emitting Diode
ASIC	Application-Specific Integrated Circuit	CMBL	Compact Muon Beam Line
NIEL	non-ionizing energy loss	FEB	Front-End Board
PKA	Primary Knock-on Atom	SWB	Switching Board
PMT	Photomultiplier Tube	DAB	Detector Adapter Board
APD	Avalanche Photodiode	CSCS	Swiss National Supercomputing Centre
GAPD	Geiger-mode Avalanche Photodiode	GPU	Graphics Processing Unit
SPAD	Single-Photon Avalanche Diode	CPU	Central Processing Unit
MPPC	Multi-Pixel Photon Counter	SMB	SciFi Module Board
PDE	Photon Detection Efficiency	HIPA	High-Intensity Proton Accelerator
DCR	Dark Count Rate	EPICS	Experimental Physics and Industrial Control System
TDC	Time-to-Digital Converter		

APPENDIX C. MUTRIG DETAILS

TID	Total Ionizing Dose	CCD	Charge-Coupled Device
PKA	Primary Knock-on Atom	SciFi	Scintillating Fibre
PS	polystyrene	SciTile	Scintillating Tile
IQR	interquartile range	HIMB	High-Intensity Muon Beams
PMMA	poly(methyl methacrylate)	UV	Ultraviolet

List of Figures

1.1	A Feynman diagram showing the decay $\mu \rightarrow eee$ via neutrino oscillation. This process is forbidden in the SM and strongly suppressed in the SM with massive neutrinos.	7
1.2	Historical and projected upper limits for the three golden cLFV channels, adapted from [28], originally from [29]. The predictions are upper limits at 90 % if the experiments do not observe an event with the SES listed in table 1.3.	9
2.1	Schematic of the Mu3e Phase I detector in lateral and transverse view. Due to a homogeneous magnetic field, positrons and electrons (red and blue) describe helical trajectories that are reconstructed from hits in the four layers of the pixel tracker.	14
2.2	When the particle momentum is deduced by the deflection Ω , multiple scattering by θ_{MS} limits the momentum resolution. For half a turn, the effect of θ_{MS} vanishes in the first order. Figures from [1].	15
2.3	CAD renderings of the SciFi detector [1]. Each of the twelve ribbons in (a) is connected to two SiPMs and SciFi Module Board (SMB) as shown in (b) , one per side.	16
2.4	Prototypes of the main SciFi detector components	17
2.5	Numbering scheme in the scintillating fibre detector	18
2.6	Rendering of the PiE5 secondary beamline in the Compact Muon Beam Line (CMBL) configuration used for the Mu3e experiment [1]	21
3.1	Comparison of the capture efficiency between single- and double-clad scintillating fibres. Only total internal reflection on the inner interfaces are considered. The second cladding increases the critical angle θ_c and thus the capture efficiency. Adapted from [41].	25
3.2	Attenuation spectrum of PS optical fibres. Figure adapted from figure 2.176 in [61]. The two curves show two independent measurements, both of which show a peak in attenuation for wavelengths around 610 nm.	26
3.3	Two different ways of propagation light under total internal reflection. The left part shows a cross-section of a round fibre, while the right part shows a lateral view. Helical light rays have greater path lengths and undergo more reflections than meridional rays.	29

LIST OF FIGURES

3.4	Layout of a silicon APD in reverse reach-through configuration. Adapted from [65]	31
3.5	Hamamatsu S13552 SiPM column array [42]. Each of the 128 channels has an effective photosensitive area of $230 \mu\text{m} \times 1625 \mu\text{m}$ and consists of 4×26 individual Geiger-mode APDs (see blow-up region). One APD has an area of $57.5 \mu\text{m} \times 62.5 \mu\text{m}$	32
3.6	Particles crossing a thin surface with thickness δ . The path length of a perpendicular particle (left) is δ , but under an angle of incidence θ the path length is $\delta / \cos(\theta)$	34
3.7	Displacement damage cross-section of neutrons, protons, pions and electrons in silicon as a function of the kinetic energy E . The function is normalized to 95 MeV mb, which corresponds to a neutron with $E = 1$ MeV [67]. An electron with $E = 15$ MeV, for example, causes a factor of 20 less damage than a 1 MeV neutron. The inset shows a zoomed-in part of the electron curve, covering the range of energies in Mu3e (up to 53 MeV). Adapted from [75], originally summarized in [76], based on data from [77] (also published in [78]), [79–82].	36
4.1	Fibre characterization setup. The fibre is excited with a UV LED in different positions. The end of the fibre is coupled to either a PIN photodiode (a) or a spectrometer (b)	39
4.2	Photograph of the characterization setup	39
4.3	Radiant flux at one fibre end as a function of distance to the side-induced fluorescence. A double-exponential function is fit to the data in blue; grey data points are excluded from the fit. The two subfigures show data from two fibre samples of the same fibre type.	42
4.4	For individual fibres, the data are not well-described by just two components of the attenuation length. Fitting a model with an additional ultra-short component Λ_{us} improves the agreement between model and data significantly.	44
4.5	Emission spectra taken from two different samples stimulated at multiple distances to the fibre end coupled to the spectrometer (colour coded)	47
4.6	The <i>differential</i> attenuation length of different fibres. For each wavelength on the abscissa in e.g. figure 4.5, the intensities of all distances are fit to determine the attenuation length at this particular wavelength. The statistical uncertainty from the fit algorithm is shown as a light orange band. This does not include any systematic uncertainties (see also section 4.3).	48
4.7	Comparison of emission spectra and differential attenuation length between blue- and green-emitting scintillating fibres by Saint-Gobain	49

4.8	Determination of the long component of the attenuation length by integrating the spectra (see figure 4.5) and fitting a single exponential function in the range covered by the solid line. Each point in this figure corresponds to one integrated wavelength spectrum. The result is used to cross-check the attenuation length determined by measuring the radiant flux.	50
4.9	Attenuation lengths for $\varnothing 1$ mm scintillating fibres published by the manufacturer, derived from the figures titled “Transmission Loss” in [41]	51
4.10	The differential attenuation length of different fibre samples from this work. Superimposed are the data from Borshchev et al. [58]. Both NOL-11 fibres studied here show slightly more attenuation than their measurement (36 % and 11 %, respectively). The differential attenuation of the first SCSF-78 sample in (c) is in good agreement to their measurement, while the second fibre shows 40 % stronger attenuation	54
4.11	The relative light intensity (radiant flux) detected from fibres during long-term excitation with UV light from different LEDs. The number indicates the wavelength in nanometres.	56
4.12	Relative emission spectra of Kuraray SCSF-78 fibre 1 at 3 and 30 cm, the size of the Mu3e detector. Overlaid is the Photon Detection Efficiency (PDE) of the Hamamatsu S13552 SiPM (data from [42]).	58
5.1	Open base of the irradiation chamber: The green boards to the left and right serve as electrical feedthroughs for the SiPMs. “Aare” and “Berg” denote two landmarks at PSI which are used as references. The “Aare” side lies at positive x values, the “Berg” side at negative x . Seen in $+z$ direction.	63
5.2	Rendering of the SiPM carrier board used for irradiation. Apart from the two connectors, the backside is flat.	64
5.3	Schematic of the lowpass filter and bias of the SiPM on the carrier board. The two halves of the SiPM can be biased individually or combined if a zero-ohm link (R3) is used.	65
5.4	PiE5 beam line and elements with irradiation chamber. Adapted from [50].	66
5.5	Displacement damage cross-section of electrons in silicon as a function of the electron kinetic energy E . This function is used to weigh the fluences in Mu3e and the irradiation setup. Based on [98].	67

5.6	Vertex position for Michel positrons reaching the SciFi SiPMs in simulation. The muon beam is generated upstream (left) and a lead collimator at $z = -400$ mm reduces the beam diameter. Due to the helical tracks of the muons, some muons pass the collimator and are stopped in a narrowing of the aluminium beampipe at $z = -135$ mm where the beam pipe changes diameter from 30 mm to 20 mm. Most muons reaching the SiPMs decay on the stopping target around $z = 0$ or the first layer of pixel sensors. The SciFi SiPMs are located at approximately $r \approx 60$ mm and $z \approx \pm 152$ mm. Muons stopped in the narrowing are responsible for a significant part of the radiation in the upstream SiPMs.	68
5.7	Simulated spectral fluence through the scintillating fibre SiPMs upstream and downstream for $2.5 \cdot 10^{15}$ muon stops on target. Shown is the fluence of all electrons and positrons regardless of the decay type.	69
5.8	Simulated spectral fluence of Michel positrons through the scintillating fibre SiPMs for $2.5 \cdot 10^{15}$ muon stops on target, the Mu3e Phase I goal. The fluence is subdivided in positrons originating from the target region (target and first pixel layer) that are either approaching the z axis (Target, in) or moving away from it (Target, out). The third category of positrons stems from muons decaying in places away from the target (Not target). Michel positrons make up the largest share of the total positron fluence shown in figure 5.7	70
5.9	Kinetic energy of incoming (per equation 5.1) Michel positrons hitting a SiPM vs their azimuth angle after the muon decay. The range of angles under which a SiPM is hit is the largest for positrons with $E \approx 18$ MeV, thus explaining the peak in figure 5.8	71
5.10	Cumulative distribution of the neutron-equivalent fluence due to Michel positrons by decay location. Positrons from the target region are split into incoming and outgoing particles as defined in equation 5.1. Most muons decaying outside the target region decay in the beampipe (see also figure 5.6).	72
5.11	Cumulative distribution of the neutron-equivalent fluence per SiPM broken down by particle type and origin. Positrons from Michel decays, and radiative muon decays are responsible for most of the deposited non-ionizing energy, which is converted to neutron-equivalents using [98].	72
5.12	Simulated neutron-equivalent fluence per SiPM array inside the Mu3e scintillating fibre detector. The fluence corresponds to $2.5 \cdot 10^{15}$ muon stops on target and is calculated from positron and electron flux under the NIEL hypothesis using [98]. The SiPM index corresponds to the number in figure 2.5.	73

5.13	Simulated total ionizing dose (TID) in each SiPM array for $2.5 \cdot 10^{15}$ muon stops on target. A silicon dioxide layer of 200 nm is simulated on top of the 300 μm thick bulk silicon. The SiPM index corresponds to the number in figure 2.5.	75
5.14	Visualization of the radiation setup and simulated particle trajectories during 64 ns. Muons pass through layers of polyimide foil before they are stopped in a polyethylene target and decay. The spectral fluence of the decay positrons as well as secondary particles through the SiPMs and the tile matrix is counted and used to estimate the equivalent fluence of 1 MeV neutrons using [98]. This figure shows the simulated decay of 315 muons generated 1 cm away from the target (green lines in the centre of the figure).	76
5.15	With a pill scan, the beam shape and total muon rate $R_{2.2\text{mA}}$, normalized to the HIPA proton beam current are determined. These results are used in a Geant4 simulation to determine the spectral flux in the irradiated samples, normalized to the proton beam current.	77
5.16	The muon beam spot in PiE5. The measured muon rate and $1\sigma, 2\sigma, 3\sigma$ contour lines of a bivariate Gaussian distribution fitted to data are shown. Note that the coordinates refer to the internal coordinates of the pill scanner. The origin of the irradiation setup lies at $x = 7.4\text{ mm}, y = 0$	78
5.17	Simulated spectral flux in both positions of the irradiation chamber. This flux is normalized to $I_p = 2.2\text{ mA}$	79
5.18	Simulated effect of the muon polarization on the neutron-equivalent flux at the nominal beam current of $I_p = 2.2\text{ mA}$, as a function of the muon polarization.	80
5.19	Monitoring variables during irradiation. In the middle of the campaign there was a beam interrupt during which no irradiation took place	88
5.20	Two-dimensional distribution of the change in neutron-equivalent flux $\dot{\Phi}_{\text{eq}}$ due to variations of the beam position and width (colour coded). The parameters μ_x, μ_y, σ_x and σ_y are varied by $\pm 5\text{ mm}$ in steps of 1 mm. The results are used as systematic uncertainties for the dose measurements. (All <i>This work</i>)	89
5.21	Currents of each SiPM during irradiation. SiPMs 1 and 2 were irradiated at a lower temperature, and the bias voltage was adjusted twice. All other SiPMs were irradiated with a constant bias voltage.	90
5.22	IV curves for SiPM 6 before (blue) and after irradiation (red). Different curves represent individual channels. The approximate region of the breakdown voltages V_{BD} is marked. The current increases after irradiation above and below breakdown.	91

LIST OF FIGURES

5.23	Temperature differences during IV measurements before and after irradiation for each channel. The measurements before irradiation were conducted at warmer temperatures for almost every individual channel.	92
5.24	Per-channel ratio of the current in SiPM 6 after and before irradiation at voltages above and below breakdown. Each data point resembles the current ratio $I_{\text{after irradiation}}/I_{\text{before irradiation}}$ in an individual channel. (a) shows the operational voltage recommended by the vendor, while (b) shows a fixed voltage below breakdown.	93
5.25	Change in current of all SiPMs as a function of neutron-equivalent fluence	95
5.26	Determination of the breakdown voltage of one SiPM channel using the relative derivative $\frac{d}{dV} \log(I(v)) = I^{-1} \frac{dI}{dV}$. The voltage at the maximum V_{peak} has an offset with respect to V_{BD} that only depends on the type of SiPM. Therefore, a change in V_{peak} is equivalent to a change in V_{BD}	97
5.27	Example of how the tangent method is used to determine the breakdown voltage. The current I , a second-order polynomial fit of the baseline and the tangent are shown. The vertical line V_{BD} marks the intersection of the baseline and the tangent.	97
5.28	Temperature-corrected breakdown voltage for each channel of SiPM 6 before and after irradiation. The bottom shows the change after irradiation.	98
5.29	Box plots of the change in breakdown voltage, as determined by two different methods. The median of the difference for each SiPM is indicated by an orange line, surrounded by a box marking the first and third quartile. The whiskers cover the central 90% of the data points. For all SiPMs except number 7, the width of the box shows the uncertainty of the neutron-equivalent fluence (SiPM 7 is not irradiated and has therefore zero fluence). A linear function is fitted to SiPMs 3–7	100
5.30	Shown are heat maps (overlays) of 100 waveforms with of 200 ns length each. The irradiated SiPMs 1 and 6 show a significant amount of dark counts and baseline shift. SiPM 7 (figure 5.30c) is not irradiated, and most waveforms do not contain any signal.	102
5.31	Two methods of determining the dark count rate of the same SiPM channel using waveforms	104
5.32	Per-channel dark count rate for three different SiPM arrays, determined by counting events in windows of 200 ns and fitting a Poisson distribution to the result	105

5.33	Current during annealing as a function of time. The left half shows the first ten hours, while the right shows the current over the course of five days. The dark solid lines show the result of a double-exponential fit with a constant baseline shown as a dotted line.	107
5.34	Calculated hit rates due to dark pulses that are validated by a coincident signal	112
6.1	Overview of the Mu3e DAQ topology. The hardware of the front-end boards is generic, but the firmware is subdetector-specific. Adapted from [1].	117
6.2	Flow of SciFi data (a) and MuTRiG configuration data (b) in the opposite direction. One SWB serves all twelve FEBs with two links to an SMB each. Each MuTRiG is configured independently using a unique address.	118
6.3	Diagram of a single MuTRiG Channel. A silicon photomultiplier can be connected differentially or single-ended. The analog signal is split into two branches for <i>timing</i> (T) and <i>energy</i> (E) discrimination. The hit logic unit combines the two trigger signals so that a single TDC can convert information from both branches to digital timestamps. The timestamps are then further processed and stored in on-chip memory. All stages can be configured via SPI [1, 43, 116]	120
6.4	Working principle of the two branches inside MuTRiG. The top shows the analog SiPM signal and the two thresholds. The outputs of the two thresholds and the exclusive disjunction are shown in the middle. The state of the two counters (bottom) determines the timing and energy timestamps. Adapted from [1, 43].	121
6.5	SciFi Module Board prototypes. The first two versions were used in test beam campaigns and proved the design. Only the third iteration will fulfil the spatial constraints inside the Mu3e detector.	123
6.6	SciFi Detector Adapter Board (DAB) prototypes. One Front-End Board (FEB) can read out up to two SciFi Module Boards (SMBs), one per <i>bank</i>	125
6.7	A Front-End Board (FEB). The connection to the DAB is realized via the connectors CON2 and CON3, serving one SMB each. Data is sent to a Switching Board (SWB) via two optical transceivers (not installed here). The clock signal can either be supplied via the optical transceivers or a differential electrical signal (top left). A JTAG port allows debugging and control of the Arria V front-end FPGA. Picture adapted from [1].	126
6.8	Switching board, developed under the name PCIe40. The heatsink obscures the Arria 10 FPGA. Figure from [1].	127

LIST OF FIGURES

6.9	Message sequence chart of MuTRiG configuration. A bit pattern containing all settings is generated in software and written to the front-end FPGA on the Front-End Board (FEB). Via a <i>configure</i> command, transmission to the MuTRiG via SPI is initiated. The MuTRiG automatically responds with the previous bit pattern, so sending the same pattern twice allows the verification. The result of this verification is stored in a status register on the FEB that can be read out by the user	129
6.10	SciFi hardware during the 2021 integration run	133
6.11	A distribution of the interarrival times between subsequent MuTRiG timestamps during the 2021 integration run. The MuTRiG receives a periodic signal expected to generate a very narrow distribution. A bug mangled up the data, so the data are random.	134

List of Tables

1.1	Quark and lepton flavours	6
1.2	Left-chiral leptons (weak eigenstates) and selected properties [12] . .	6
1.3	Current and planned upper branching ratio limits of the three golden channels	8
2.1	Time schedule of Mu3e as of December 2021 [33]	12
4.1	Uncertainties for photodiode and power meter	40
4.2	Fit results for the short and long components of the attenuation length for each fibre.	43
4.3	Summary of literature comparison between the measurements in this work (on 250 μm fibres) and selective other measurements	52
5.1	Fit results for the beam spot measurements	78
5.2	“Summary of beam profile parameters at a standard pill measurement position DS [downstream] the SML/collimator system [...] for the G4BL simulation starting from TgE.” [50] The quoted values are used to estimate the uncertainty of the beam spot size due to the distance between the stopping target and where the beam spot was measured	82
5.3	Systematic uncertainties of the beam parameters and their effect on the neutron-equivalent flux	83
5.4	Systematic uncertainty of neutron-equivalent flux. The uncertainties are treated as uncorrelated, so the total uncertainties for each side are obtained by summation in quadrature.	84
5.5	Integrated proton current for all SiPMs	85
5.6	Time scale and relative recovery of high-temperature annealing . . .	107
5.7	Overview of partial annealing studies in SiPMs	108
6.1	Comparison of numbers of channels and data rates of the two MuTRiG-based subdetectors. Each ASIC has 32 channels. Although the SciFi detector has fewer readout channels, it produces more data due to hit multiplicity and noise [1]	116

LIST OF TABLES

B.1	Simulated total lepton fluence per scintillating fibre SiPM array for $2.5 \cdot 10^{15}$ muon stops on target. The equivalent fluence of neutrons with an energy of 1 MeV is calculated under the NIEL hypothesis with the help of [98].	162
B.2	Simulated neutron-equivalent fluence due to Michel by decay location of the muon. Positrons coming from the target region are split in incoming and outgoing particles, defined by the scalar product of the momentum and location vectors. Most muons decaying outside of the target region decay inside the beampipe	163
B.3	Simulated neutron-equivalent fluence for $2.5 \cdot 10^{15}$ muon stops on target, grouped by the origin of the particles. The lower digit marks the particle type (1: positron, 2: electron) and the higher digit marks the process (1: Michel decay, 2: radiative muon decay, 4: photon conversion, 5: Bhabha scattering, 8: Compton scattering)	164
B.4	Simulated total ionizing doses in the bulk silicon and silicon dioxide of each silicon photomultiplier array. The total ionizing dose corresponds to $2.5 \cdot 10^{15}$ muon stops on target	165
B.5	Insertion and extraction times for all SiPMs in Unix time. The <i>side</i> refers to the two possible positions inside the setup, <i>Aare</i> and <i>Berg</i> are landmark references at PSI. The integral of the proton current multiplied by the shutter status <i>S</i> is proportional to the received dose. In addition to the six samples, a populated PCB missing the SiPM was irradiated as a control sample	166
C.1	DAQ Settings, no direct effect on the MuTRiG configuration pattern.	188
C.2	Global Settings. These settings are used for all MuTRiGs.	189
C.3	Settings per single MuTRiG.	190
C.4	Settings per single channel.	191
C.5	Coincidence matrix settings. Not accessible via the web frontend. Lower and upper crossbar refer to the connections to neighbouring chips at the respective end of the chip.	192
C.6	Configuration byte map for MuTRiG 2. Data are sent from top left to bottom right. The numbers in parentheses denote bit indices of a parameter. A parameter that is set to 8 would be represented as 1000, with a one at index (3) and zeroes at indices (2:0). Scale parameters are denoted by the suffix <i>sc</i> , and “enable” flags by the suffix <i>en</i>	193
C.7	MuTRiG Event data. For short events, energy information is omitted, resulting in 27-bit events, long events have a length of 48 bits. The length is denoted in number of bits.	195
C.8	A MuTRiG data frame	196

C.9 MuTRiG frame flags. Certain settings can be marked, the generation of K.28.5 symbols as *idle signals* between frames, the event mode (see above), whether PRBS debug “data” is sent and if on-chip memory (L2 FIFO) is full. 196

Bibliography

- [1] K. Arndt *et al.*, “Technical design of the phase I Mu3e experiment,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 1014, 165679, Oct 2021. doi:10.1016/j.nima.2021.165679.
- [2] N. Amos, A. Bross, and M. Lundin, “Optical attenuation length measurements of scintillating fibers,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 297, no. 3, pp. 396–403, 1990. doi:https://doi.org/10.1016/0168-9002(90)91321-2.
- [3] E. Garutti and Y. Musienko, “Radiation damage of SiPMs,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 926, pp. 69–84, 2019, silicon Photomultipliers: Technology, Characterisation and Applications. doi:https://doi.org/10.1016/j.nima.2018.10.191.
- [4] F. Hönniger, “Radiation damage in silicon. defect analysis and detector properties,” Ph.D. dissertation, 2007.
- [5] C. Inguibert, P. Arnolda, T. Nuns, and G. Rolland, “Effective NIEL in Silicon: Calculation Using Molecular Dynamics Simulation Results,” *IEEE Transactions on Nuclear Science*, vol. 57, no. 4, pp. 1915–1923, 2010. doi:10.1109/TNS.2010.2049581.
- [6] T. Nuns *et al.*, “Experimental Study of the NIEL Scaling for Silicon Devices,” in *RADECS 2018*, Göteborg, Sweden, Sep. 2018.
- [7] C. Leroy and P.-G. Rancoita, *Principles of Radiation Interaction in Matter and Detection*, 4th ed. WORLD SCIENTIFIC, 2016. doi:10.1142/9167.
- [8] A. Kuonen, O. Girard, and G. Haefeli, “Characterisation of the Hamamatsu MPPC multichannel array for LHCb SciFi Tracker v.9.2014,” April 2015, LHCb internal note LPHE Note 2015-01.
- [9] M. Köppel, “Mu3e Integration Run 2021,” 2022. doi:10.48550/ARXIV.2203.07855.
- [10] A. M. F. Halzen, *Quarks and Leptons*. John Wiley & Sons, 1984.
- [11] M. Thomson, *Modern particle physics*. New York: Cambridge University Press, 2013.

BIBLIOGRAPHY

- [12] P. Zyla *et al.* [Particle Data Group], “Review of Particle Physics,” *PTEP*, vol. 2020, no. 8, 083C01, 2020. doi:10.1093/ptep/ptaa104.
- [13] Y. Fukuda *et al.* [Super-Kamiokande Collaboration], “Evidence for oscillation of atmospheric neutrinos,” *Phys. Rev. Lett.*, vol. 81, pp. 1562–1567, 1998.
- [14] F. An *et al.* [Daya Bay Collaboration], “Observation of electron-antineutrino disappearance at Daya Bay,” *Phys.Rev.Lett.*, vol. 108, 171803, 2012. doi:10.1103/PhysRevLett.108.171803.
- [15] Q. R. Ahmad *et al.* [SNO Collaboration], “Measurement of the charged current interactions produced by B-8 solar neutrinos at the Sudbury Neutrino Observatory,” *Phys. Rev. Lett.*, vol. 87, 071301, 2001.
- [16] G. Hernández-Tomé, G. L. Castro, and P. Roig, “Erratum to: Flavor violating leptonic decays of τ and μ leptons in the standard model with massive neutrinos,” *The European Physical Journal C*, vol. 80, no. 5, may 2020. doi:10.1140/epjc/s10052-020-7935-5.
- [17] R. R. Crittenden, W. D. Walker, and J. Ballam, “Radiative decay modes of the muon,” *Phys. Rev.*, vol. 121, pp. 1823–1832, Mar 1961. doi:10.1103/PhysRev.121.1823.
- [18] 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, 2018.
- [19] R. Aaij *et al.* [LHCb], “Search for the lepton flavour violating decay $b^+ \rightarrow k^+ \mu^- \tau^+$ using B_{s2}^{*0} decays,” *Journal of High Energy Physics*, vol. 2020, 129, 2020.
- [20] A. M. Sirunyan *et al.* [CMS], “Search for the lepton flavor violating decay $\tau \rightarrow 3\mu$ in proton-proton collisions at $\sqrt{s} = 13$ TeV,” *JHEP*, vol. 01, 163, 2021. doi:10.1007/JHEP01(2021)163.
- [21] E. P. Hincks and B. Pontecorvo, “Search for gamma-radiation in the 2.2-microsecond meson decay process,” *Phys. Rev.*, vol. 73, pp. 257–258, Feb 1948. doi:10.1103/PhysRev.73.257.
- [22] A. M. Baldini *et al.* [MEG], “Search for the lepton flavour violating decay $\mu^+ \rightarrow e^+ \gamma$ with the full dataset of the MEG experiment,” *Eur. Phys. J. C*, vol. 76, no. 8, 434, 2016. doi:10.1140/epjc/s10052-016-4271-x.
- [23] W. Bertl *et al.*, “A search for μ -e conversion in muonic gold,” *The European Physical Journal C-Particles and Fields*, vol. 47, no. 2, pp. 337–346, 2006.

- [24] L. Bartoszek *et al.* [Mu2e], “Mu2e Technical Design Report,” 10 2014. doi:10.2172/1172555.
- [25] R. Abramishvili *et al.* [COMET collaboration], “COMET Phase-I technical design report,” *Progress of Theoretical and Experimental Physics*, vol. 2020, no. 3, 033C01, 2020.
- [26] U. Bellgardt *et al.* [SINDRUM Collaboration], “Search for the Decay $\mu^+ \rightarrow e^+e^+e^-$,” *Nucl.Phys.*, vol. B299, 1, 1988. doi:10.1016/0550-3213(88)90462-2.
- [27] A. o. Baldini, “The design of the MEG II experiment,” *The European Physical Journal C*, vol. 78, no. 5, pp. 1–60, 2018.
- [28] M. Aiba *et al.*, “Science Case for the new High-Intensity Muon Beams HIMB at PSI,” 2021. <https://arxiv.org/abs/2111.05788>. doi:10.48550/ARXIV.2111.05788.
- [29] R. Bernstein and P. Cooper, “Charged lepton flavor violation: An experimenter’s guide,” *Physics Reports*, vol. 532, 07 2013. doi:10.1016/j.physrep.2013.07.002.
- [30] L. Michel, “Interaction between four half-spin particles and the decay of the μ -meson,” *Proceedings of the Physical Society. Section A*, vol. 63, no. 5, 514, 1950.
- [31] Y. Kuno and Y. Okada, “Muon decay and physics beyond the standard model,” *Rev. Mod. Phys.*, vol. 73, pp. 151–202, 2001.
- [32] A. Knecht, “High Intensity Muon Beam (HiMB) at PSI,” in *Snowmass CLFV with high intensity muon factories*, Dec 2020.
- [33] A. Schöning, “Current Mu- \rightarrow eee experiment,” 2021, cLFV White Paper Discussion.
- [34] W. Bertl *et al.* [SINDRUM Collaboration], “Search for the decay $\mu^+ \rightarrow e^+e^+e^-$,” *Nucl. P.*, vol. B 260, no. 1, pp. 1 – 31, 1985. doi:DOI: 10.1016/0550-3213(85)90308-6.
- [35] G. Pruna, A. Signer, and Y. Ulrich, “Fully differential NLO predictions for the rare muon decay,” *Phys. Lett. B*, vol. 765, pp. 280–284, 2017. doi:10.1016/j.physletb.2016.12.039.
- [36] M. Fael and C. Greub, “Next-to-leading order prediction for the decay $\mu \rightarrow e (e^+e^-) \nu \bar{\nu}$,” *JHEP*, vol. 01, 084, 2017. doi:10.1007/JHEP01(2017)084.
- [37] T. Rudzki *et al.*, “The Mu3e experiment: Toward the construction of an HV-MAPS vertex detector,” *arXiv preprint arXiv:2106.03534*, 2021.

BIBLIOGRAPHY

- [38] I. Perić, “A novel monolithic pixelated particle detector implemented in high-voltage CMOS technology,” *Nucl.Instrum.Meth.*, vol. A582, 876, 2007. doi:10.1016/j.nima.2007.07.115.
- [39] T. Rudzki, “The Mu3e vertex detector - construction, cooling, and first prototype operation,” PhD Thesis, Heidelberg University, 2022.
- [40] S. Corrodi, “A Timing Detector based on Scintillating Fibres for the Mu3e Experiment,” PhD Thesis, ETH Zürich, 2018.
- [41] L. Kuraray Co. Plastic scintillating fibers. https://www.kuraray.com/uploads/5a717515df6f5/PR0150_psf01.pdf (accessed 2022-01-19).
- [42] Hamamatsu Photonics K.K. “S13552 Datasheet”. Cat. No. KAPD1083E01 Apr. 2021. https://www.hamamatsu.com/resources/pdf/ssd/s13552_kapd1083e.pdf (accessed 2022-01-19).
- [43] H. Chen, “A Silicon Photomultiplier Readout ASIC for the Mu3e Experiment,” Ph.D. dissertation, Heidelberg University, 2018.
- [44] P. Eckert, “The Mu3e Tile Detector,” PhD thesis, Heidelberg University, 2015.
- [45] Kuraray Co., Ltd. (2021, Jul) Fast timing plastic scintillator ej-228, ej-230. https://eljentechnology.com/images/products/data_sheets/EJ-228_EJ-230.pdf (accessed 2022-07-25).
- [46] Hamamatsu Photonics K.K. (2016, Aug.) “S13360 series Datasheet”. Cat. No. KAPDB0324EA. http://www.hamamatsu.com/resources/pdf/ssd/s13360_series_kapd1052e.pdf (accessed 2022-01-19).
- [47] J. Grillenberger, C. Baumgarten, and M. Seidel, “The High Intensity Proton Accelerator Facility,” *SciPost Phys. Proc.*, 2, 2021. doi:10.21468/SciPostPhysProc.5.002.
- [48] A. Pifer, T. Bowen, and K. Kendall, “A High Stopping Density μ^+ Beam,” *Nucl.Instrum.Meth.*, vol. 135, pp. 39–46, 1976. doi:10.1016/0029-554X(76)90823-5.
- [49] C. Grupen, *Particle Detectors*. Cambridge University Press, 1997.
- [50] 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.
- [51] F. Berg *et al.*, “Target studies for surface muon production,” *Phys. Rev. Accel. Beams*, vol. 19, 024701, Feb 2016. doi:10.1103/PhysRevAccelBeams.19.024701.

- [52] Z. D. Hodge, “Production, Characterization, and Monitoring of Surface Muon Beams for Charged Lepton Flavor Violation Experiments,” Ph.D. dissertation, ETH Zürich, 2018.
- [53] G. F. Knoll, *Radiation detection and measurement*. John Wiley & Sons, 2010.
- [54] W. R. Leo, *Techniques for nuclear and particle physics experiments: a how-to approach; 2nd ed.* Berlin: Springer, 1994. doi:10.1007/978-3-642-57920-2.
- [55] S. A. Ponomarenko *et al.*, “Nanostructured organosilicon luminophores and their application in highly efficient plastic scintillators,” *Scientific reports*, vol. 4, no. 1, pp. 1–8, 2014.
- [56] L. Reiffel and N. S. Kapany, “Some considerations on luminescent fiber chambers and intensifier screens,” *Review of Scientific Instruments*, vol. 31, no. 10, pp. 1136–1142, 1960. doi:10.1063/1.1716826.
- [57] L. Collaboration, “LHCb Tracker Upgrade Technical Design Report,” CERN, Tech. Rep. CERN-LHCC-2014-001. LHCb-TDR-015, Feb 2014.
- [58] O. Borshchev *et al.*, “Development of a new class of scintillating fibres with very short decay time and high light yield,” *Journal of Instrumentation*, vol. 12, no. 05, pp. P05 013–P05 013, may 2017. doi:10.1088/1748-0221/12/05/p05013.
- [59] Plastic Scintillating Fibers. Saint-Gobain Crystals. <https://www.crystals.saint-gobain.com/sites/hps-mac3-cma-crystals/files/2021-11/Fiber-Product-Sheet.pdf> (accessed 2022-12-15).
- [60] Scintillating Optical Fibers. Saint-Gobain Crystals. <http://web.archive.org/web/20200222134709/https://www.crystals.saint-gobain.com/sites/imdf.crystals.com/files/documents/fiber-brochure.pdf> (accessed 2020-02-22).
- [61] O. Ziemann, J. Krauser, P. E. Zamzow, and W. Daum, *POF handbook*. Springer, 2008.
- [62] C.-A. Bunge, R. Kruglov, and H. Poisel, “Rayleigh and Mie Scattering in polymer optical fibers,” *Journal of Lightwave Technology*, vol. 24, pp. 3137 – 3146, Sep. 2006. doi:10.1109/JLT.2006.878077.
- [63] G. Mie, “Beiträge zur Optik trüber Medien, speziell kolloidaler Metallösungen,” *Annalen der Physik*, vol. 330, no. 3, pp. 377–445, 1908. doi:<https://doi.org/10.1002/andp.19083300302>.

BIBLIOGRAPHY

- [64] J. W. Strutt, *On the Light from the Sky, its Polarization and Colour*, ser. Cambridge Library Collection - Mathematics. Cambridge University Press, 2009, vol. 1, 87–103. doi:10.1017/CBO9780511703966.009.
- [65] Hamamatsu Photonics K.K. (2021, Jun.) “MPPC®”. Cat. No. KAPD9005E03. https://www.hamamatsu.com/resources/pdf/ssd/mppc_kapd9005e.pdf (accessed 2022-01-19).
- [66] V. S. Vavilov and N. A. Ukhin, *Radiation Effects in Semiconductors and Semiconductor Devices*. Boston, MA: Springer US, 1977. doi:10.1007/978-1-4684-9069-5_5.
- [67] [American Society for Testing and Materials] “Standard practice for characterizing neutron fluence spectra in terms of an equivalent monoenergetic neutron fluence for radiationhardness testing of electronics,” *ASTM E722 – 09*, 2009.
- [68] Q. Yan, X. Meng, D. Liu, Q. Zhang, and J. Zhu, “Evaluation of displacement damage in solids induced by fast positrons: Modeling and effect on vacancy measurement,” *Nuclear Materials and Energy*, vol. 27, 101022, 2021. doi:<https://doi.org/10.1016/j.nme.2021.101022>.
- [69] G. P. Summers *et al.*, “Energy dependence of proton displacement damage factors for bipolar transistors,” *IEEE Transactions on Nuclear Science*, vol. 33, no. 6, pp. 1282–1286, 1986. doi:10.1109/TNS.1986.4334593.
- [70] G. P. Summers *et al.*, “Correlation of particle-induced displacement damage in silicon,” *IEEE Transactions on Nuclear Science*, vol. 34, no. 6, pp. 1133–1139, 1987. doi:11.1109/TNS.1987.4337442.
- [71] C. Dale *et al.*, “High energy electron induced displacement damage in silicon,” *IEEE Transactions on Nuclear Science*, vol. 35, no. 6, pp. 1208–1214, 1988. doi:10.1109/23.25441.
- [72] R. Wunstorf, “A Systematic investigation of the radiation hardness of silicon detectors for high-energy physics experiments,” Ph.D. dissertation, Oct. 1992.
- [73] M. Huhtinen, “Simulation of non-ionising energy loss and defect formation in silicon,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 491, no. 1, pp. 194–215, 2002. doi:[https://doi.org/10.1016/S0168-9002\(02\)01227-5](https://doi.org/10.1016/S0168-9002(02)01227-5).
- [74] J. R. Srour and J. W. Palko, “A framework for understanding displacement damage mechanisms in irradiated silicon devices,” *IEEE Transactions on Nuclear Science*, vol. 53, no. 6, pp. 3610–3620, 2006. doi:10.1109/TNS.2006.885796.

- [75] E. Lavrik, “Development of quality assurance procedures and methods for the CBM Silicon Tracking System,” Ph.D. dissertation, 2017. doi:10.15496/publikation-20433.
- [76] M. Moll, “Radiation damage in silicon particle detectors. Microscopic defects and macroscopic properties,” Ph.D. dissertation, Universität Hamburg, 1999.
- [77] P. Griffin, SAND92-0094 (Sandia National Laboratory 93), private communications. See also [76, 78].
- [78] G. Lindström, “Radiation damage in silicon detectors,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 512, no. 1, pp. 30–43, 2003. doi:https://doi.org/10.1016/S0168-9002(03)01874-6.
- [79] A. Y. Konobeyev, Y. Korovin, and V. N. Sosnin, “Neutron displacement cross-sections for structural materials below 800 MeV,” *Journal of Nuclear Materials*, vol. 186, pp. 117–130, 1992.
- [80] M. Huhtinen and P. Aarnio, “Pion induced displacement damage in silicon devices,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 335, pp. 580–582, 1993.
- [81] G. Summers *et al.*, “Damage correlations in semiconductors exposed to gamma, electron and proton radiations,” *IEEE Transactions on Nuclear Science*, vol. 40, no. 6, pp. 1372–1379, 1993. doi:10.1109/23.273529.
- [82] M. Huhtinen and P. Aarnio, “Estimation of pion induced displacement damage in silicon,” in *HU-SEFT R 1993-02*, 1993.
- [83] Newport Corporation. <https://www.newport.com/> (accessed Nov, 2022).
- [84] Ocean Insight, USA. <https://www.oceaninsight.com> (accessed Nov, 2022).
- [85] Saint-Gobain. Organic scintillation materials and assemblies. <https://www.crystals.saint-gobain.com/sites/hps-mac3-cma-crystals/files/2021-09/Organics-Plastic-Scintillators.pdf> (accessed 2022-12-15).
- [86] E. L. Volker Blobel, *Statistische und numerische Methoden der Datenanalyse*. Deutsches Elektronen-Synchrotron (DESY), 2012.
- [87] P. T. Boggs, P. T. Boggs, J. E. Rogers, and R. B. Schnabel. (1992) User’s reference guide for odrpack version 2.01: Software for weighted orthogonal distance regression.
- [88] P. Virtanen *et al.*, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020. doi:10.1038/s41592-019-0686-2.

BIBLIOGRAPHY

- [89] K. E. Atkinson, *An introduction to numerical analysis*. John Wiley & Sons, 2008.
- [90] A. Antonelli, M. Antonelli, G. Barbiellini, M. Barone, S. Bertolucci, S. Bianco, C. Bini, C. Bloise, V. Bolognesi, F. Bossi *et al.*, “Measurements of light yield, attenuation length and time response of long samples of “blue” scintillating fibers,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 370, no. 2-3, pp. 367–371, 1996.
- [91] A. Damyanova, “Development of the Scintillating Fiber Detector for Timing Measurements in the Mu3e Experiment,” PhD thesis, University of Geneva, 2019.
- [92] M. Khazova and J. B. O’Hagan. (2008, Aug.) Optical radiation emissions from compact fluorescent lamps. <https://doi.org/10.1093/rpd/ncn234>. doi:10.1093/rpd/ncn234.
- [93] E. Aschenauer *et al.*, “Tests of a fiber detector concept for high rate particle tracking,” in *AIP Conference Proceedings*, vol. 450, no. 1. American Institute of Physics, 1998, pp. 249–256.
- [94] T. Zhong, “Irradiation Studies for the Mu3e Tile detector,” in *SiPM Radiation: Quantifying Light for Nuclear, Space and Medical Instruments under Harsh Radiation Conditions*, Apr 2022.
- [95] Alphasense oxygen sensors (o2). <https://www.alphasense.com/products/oxygen-sensors/> (accessed 6 May 2022).
- [96] Yoctopuce sarl. <https://www.yoctopuce.com> (accessed 6 May 2022).
- [97] Experimental Physics and Industrial Control System. <https://epics-controls.org/> (accessed 2021-01-17).
- [98] P. R. M.J. Boschini and M. Tacconi. (2014) SR-NIEL Calculator: Screened Relativistic (SR) Treatment for Calculating the Displacement Damage and Nuclear Stopping Powers for Electrons, Protons, Light- and Heavy- Ions in Materials (version 7.7.3), available at INFN sez. Milano-Bicocca, Italy. <http://www.sr-niel.org/> (accessed May, 2022).
- [99] M. J. Berger, J. Coursey, M. Zucker, and J. Chang, “Estar,” *PSTAR, and ASTAR: computer programs for calculating stopping-power and range tables for electrons, protons, and helium ions (version 1.2. 3)*, vol. 400, 2005.
- [100] G. Cowan, *Statistical data analysis*. Oxford University Press, USA, 1998.
- [101] A. M. Baldini *et al.* [MEG], “Muon polarization in the MEG experiment: predictions and measurements,” *Eur. Phys. J. C*, vol. 76, no. 4, 223, 2016. doi:10.1140/epjc/s10052-016-4047-3.

- [102] F. Hinterberger, *Physik der Teilchenbeschleuniger und Ionenoptik*, 2nd ed. Berlin Heidelberg: Springer-Verlag, 2008.
- [103] P. Horowitz, W. Hill, and I. Robinson, *The Art of Electronics*, 3rd ed. Cambridge university press Cambridge, 2015.
- [104] F. Nagy, G. Hegyesi, G. Kalinka, and J. Molnár, “A model based DC analysis of SiPM breakdown voltages,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 849, pp. 55–59, 2017. doi:<https://doi.org/10.1016/j.nima.2017.01.002>.
- [105] A. Kuonen, G. Haefeli, M. E. Stramaglia, and O. Girard [LHCb Collaboration], “Characterisation of the Hamamatsu MPPC multichannel array for the LHCb SciFi Tracker v.11.2016,” EPFL, Tech. Rep., 2017.
- [106] A. Kuonen, “Development and characterisation of silicon photomultiplier multichannel arrays for the readout of a large scale scintillating fibre tracker,” Ph.D. dissertation, 2018.
- [107] M. Calvi *et al.*, “Single photon detection with SiPMs irradiated up to $1 \cdot 10^{14} \text{ cm}^{-2}$ 1-MeV-equivalent neutron fluence,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 922, pp. 243–249, 2019. doi:<https://doi.org/10.1016/j.nima.2019.01.013>.
- [108] T. Tsang, T. Rao, S. Stoll, and C. Woody, “Neutron radiation damage and recovery studies of SiPMs,” *Journal of Instrumentation*, vol. 11, no. 12, pp. P12 002–P12 002, dec 2016. doi:[10.1088/1748-0221/11/12/p12002](https://doi.org/10.1088/1748-0221/11/12/p12002).
- [109] A. Heering *et al.*, “Effects of very high radiation on SiPMs,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 824, pp. 111–114, 2016, frontier Detectors for Frontier Physics: Proceedings of the 13th Pisa Meeting on Advanced Detectors. doi:<https://doi.org/10.1016/j.nima.2015.11.037>.
- [110] D. Pitzl *et al.*, “Type inversion in silicon detectors,” *Nucl. Instrum. Methods Phys. Res. A*, vol. 311, no. 1, pp. 98–104, 1992. doi:[https://doi.org/10.1016/0168-9002\(92\)90854-w](https://doi.org/10.1016/0168-9002(92)90854-w).
- [111] A. Bravar *et al.*, “Development of the Scintillating Fiber Timing Detector for the Mu3e Experiment,” 2022. doi:[10.48550/ARXIV.2208.09906](https://doi.org/10.48550/ARXIV.2208.09906).
- [112] H. Augustin *et al.*, “The Mu3e Data Acquisition,” *IEEE Transactions on Nuclear Science*, vol. 68, no. 8, pp. 1833–1840, 2021. doi:[10.1109/TNS.2021.3084060](https://doi.org/10.1109/TNS.2021.3084060).
- [113] [Swiss National Supercomputing Centre] (2018, Jul) CSCS will store petabyte data for the Paul Scherrer Institute. <https://www.cscs.ch/publications/press-releases/2018/589/> (accessed 2022-01-19).

BIBLIOGRAPHY

- [114] T. Harion, “The STiC ASIC: High Precision Timing with Silicon Photomultipliers,” Ph.D. dissertation, Heidelberg University, 2015.
- [115] T. Harion *et al.*, “STiC: a mixed mode silicon photomultiplier readout ASIC for time-of-flight applications,” *Journal of Instrumentation*, vol. 9, no. 02, C02003, 2014.
- [116] W. Shen, et al., “A silicon photomultiplier readout ASIC for time-of-flight application using a new time-of-recovery method.” *IEEE Transactions on Nuclear Science*, vol. 65, pp. 1196–1202, 2018.
- [117] A. X. Widmer and P. A. Franaszek, “A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code,” *IBM Journal of Research and Development*, vol. 27, 440, 1983.
- [118] P. Franaszek and A. Widmer, “Byte oriented DC balanced (0,4) 8B/10B partitioned block transmission code,” Dec. 4 1984. <http://www.google.com/patents/US4486739>
- [119] L. Noehte, “Microstrip Impedance Control in High Density Interconnects for the Mu3e Electrical Readout Chain,” Master Thesis, Heidelberg University, 2019.
- [120] Arria V FPGA and SoC FPGA. <https://www.intel.com/content/www/us/en/products/details/fpga/arria/v.html> (accessed 2021-12-30).
- [121] Nios® Soft Processor Series. <https://www.intel.com/content/www/us/en/products/details/fpga/nios-processor.html> (accessed 2021-01-07).
- [122] P. Durante, N. Neufeld, R. Schwemmer, U. Marconi, G. Balbi, and I. Lax, “100Gbps PCI-express readout for the LHCb upgrade,” in *Proceedings, 19th Real Time Conference (RT2014)*, 2014. doi:10.1109/RTC.2014.7097492.
- [123] P. Durante, N. Neufeld, R. Schwemmer, G. Balbi, and U. Marconi, “100 Gbps PCI-Express readout for the LHCb upgrade,” *JINST*, vol. 10, no. 04, C04018, 2015. doi:10.1088/1748-0221/10/04/C04018.
- [124] J. Cachemiche, “PCIE40 Status report,” in *LHCb Electronics upgrade WG*, 2015.
- [125] DE5a-Net Arria 10 FPGA Development Kit. <https://www.terasic.com.tw/cgi-bin/page/archive.pl?Language=English&CategoryNo=231&No=970&PartNo=1> (accessed 2021-01-06).
- [126] S. Ritt, P. Amaudruz, and K. Olchanski, “Maximum Integration Data Acquisition System,” 2001. <http://midas.psi.ch>