

DISSERTATION

submitted to the

Combined Faculty of Mathematics, Engineering and Natural Sciences

of Heidelberg University, Germany

for the degree of

Doctor of Natural Sciences

Put forward by

Hannah Kl ingenmeyer

born in Heidelberg

Oral examination: 15.05.2024

Design and construction
of the Mu3e Tile Detector

Referees: Prof. Dr. Hans-Christian Schultz-Coulon
Jun.-Prof. Dr. Loredana Gastaldo

Abstract

The Mu3e experiment is designed to search for the charged lepton flavour violating decay $\mu \rightarrow eee$ with an ultimate single event sensitivity of 10^{-16} – four orders of magnitude better than the current best exclusion limit. In order to reach this sensitivity, excellent vertex, momentum, and timing measurements of the decay products are needed for the background suppression. With a timing resolution < 100 ps, the Tile Detector is designed to provide the most accurate timing information in the Mu3e experiment and is crucial for the suppression of combinatorial background. It consists of two barrels built in a modular approach and utilises plastic scintillator tiles, which are individually coupled to silicon photomultipliers read out by a custom ASIC.

This thesis describes the design, construction, and thermal simulation of the Tile Detector from the prototype stage to production readiness. The construction of the first technical prototype proves the functionality of the detector concept, yielding a single-channel time resolution of ≈ 47 ps. With the technical prototype serving as a baseline, the Tile Detector design is adapted and optimised in order to simplify the construction on the one hand, and on the other hand to facilitate the mechanical integration into the experiment. In parallel, thermal simulations of the intermediate design stages of the cooling circuit ensure that the cooling capability is suitable for the Tile Detector requirements. Based on the finalised detector design, which was implemented in this work, a production line with custom tools and dedicated quality control procedures is developed and evaluated during the construction of two preproduction modules. Measurement campaigns are conducted to commission and tune the modules using first calibration procedures. After excluding all known and fixable errors, more than 96% of the tested channels are fully functional. As a result, production readiness is achieved for the Tile Detector, paving the way for the full detector construction, installation and commissioning foreseen in 2024/2025.

Zusammenfassung

Das $\text{Mu}3e$ -Experiment ist für die Suche nach dem Lepton-Flavour verletzenden Zerfall $\mu \rightarrow eee$ konzipiert. Die angestrebte finale Sensitivität des Experiments beträgt dabei 10^{-16} , was eine Verbesserung der aktuellen Obergrenze um vier Größenordnungen darstellt. Um diese Sensitivität zu erreichen, wird eine hervorragende Vertex-, Impuls- und Zeitauflösung der Zerfallsprodukte zur Unterdrückung des Untergrundes benötigt. Der *Kacheldetektor* ist mit einer Zeitauflösung unter 100 ps darauf ausgelegt, innerhalb des $\text{Mu}3e$ -Experimentes die präziseste Zeitinformation bereitzustellen, und ist daher essentiell für die Unterdrückung von kombinatorischem Untergrund. Der Detektor besteht aus zwei zylindrischen Einheiten, die jeweils modular aufgebaut sind. Die Grundbestandteile sind Kacheln aus einem Plastiksintillator-Material, die einzeln an Silizium-Photomultiplier gekoppelt sind. Die Auslese erfolgt durch einen eigens dafür entwickelten ASIC.

Die vorliegende Dissertation beschreibt die Planung, den Bau und die thermische Simulation des Kacheldetektors von der Prototypphase bis zur Produktionsreife. Der Bau des ersten technischen Prototyps zeigt die Funktionalität des Detektorkonzepts, mit dem eine hervorragende Einzelkanal-Zeitauflösung von ≈ 47 ps erzielt werden kann. Der technische Prototyp fungiert daher als Grundlage für verschiedene Anpassungen und Optimierungen des Kacheldetektor-Designs, die sowohl dazu dienen, den Detektorbau zu erleichtern, als auch den Einbau in das $\text{Mu}3e$ -Experiment zu ermöglichen. Parallel dazu werden fortlaufend thermische Simulationen von neuen Entwürfen des Kühlkreislaufes durchgeführt um sicherzustellen, dass die Kühlfähigkeit trotz der Anpassungen den Anforderungen des Detektors entspricht. Des Weiteren wird auf Grundlage des finalen Detektorkonzepts, das in dieser Arbeit realisiert wurde, eine Produktionsstraße mit eigens entwickelten Vorrichtungen zum Detektorbau sowie Maßnahmen zur Qualitätssicherung aufgebaut und anhand des Baus zweier Vorserien-Module evaluiert. Die Module werden auf dedizierten Messkampagnen in Betrieb genommen und anhand erster Kalibrationsalgorithmen abgestimmt. Alle bekannten und korrigierbaren Fehlerquellen ausgenommen, sind über 96 % der getesteten Kanäle voll funktionsfähig und die finale Produktion des Kacheldetektors kann dementsprechend beginnen. Der Bau, die Installation und die Inbetriebnahme des Detektors sind für die Jahre 2024/2025 geplant.

Contents

1	Introduction	1
2	The Standard Model and beyond	5
2.1	The Standard Model of particle physics	5
2.2	Charged lepton flavour violation	10
2.3	Lepton flavour violation experiments with muons	11
3	The Mu3e experiment	15
3.1	Signal decay and background processes	15
3.2	Design of the Mu3e experiment	18
3.3	Introduction to the Tile Detector	27
4	Tile Detector construction	37
4.1	Updated detector design	37
4.2	Scintillating tile production	46
4.3	Development of a tile matrix production line	51
4.4	Module assembly and evaluation	59
5	Development and evaluation of quality control procedures	67
5.1	Quality control of the scintillating tiles	68
5.2	The SiPM matrix quality control system	73
5.3	Development of a light yield-based matrix quality control system	78
5.4	Outlook for the Tile Detector module QC	92
6	Detector cooling studies	97
6.1	Cooling simulation and validation of the technical prototype	97
6.2	Extension and adaptation of the simulation	102
6.3	Outlook on the final cooling system	109
7	Conclusion	115
	Appendices	119
A	Supplementary material for the Tile Detector dimensions	121
B	Supplementary material for the quality control	125
B.1	KLauS calibration	126
B.2	Light yield QC	139

Contents

C Lists	141
C.1 List of Figures	141
C.2 List of Tables	145
Acronyms	147
Bibliography	151
Acknowledgements	159

Chapter 1

Introduction

Since the discovery of the electron at the end of the 19th century, the field of particle physics has continuously grown and evolved, with a host of new particles being discovered over the past century. The current understanding of the fundamental particles and their interactions culminates in the *Standard Model of particle physics*, which has been remarkably successful in the prediction of new particles and the description of experimental data over the past decades. Nonetheless, there are still open questions concerning the nature of the Universe which the Standard Model cannot answer, suggesting that there may well be *physics beyond the Standard Model*.

One phenomenon, which could shed light on the nature of this new physics, is the observation of lepton flavour violation. In the Standard Model, the lepton flavour is assumed to be conserved; however, the discovery of neutrino oscillations [1, 2] shows that the conservation of lepton flavour is violated in the neutral sector. The observation of lepton flavour in the charged sector, *charged lepton flavour violation (cLFV)*, would be a clear sign of physics beyond the Standard Model and is therefore a promising candidate for experimental searches. One of these experiments is Mu3e [3], which is designed to search for the cLFV decay $\mu \rightarrow eee$. It is based on state-of-the-art semiconductor sensors for precise tracking and vertexing, and two scintillator-based timing systems providing an excellent time resolution.

One of the timing systems is the *Tile Detector*, which is designed to provide the most accurate timing information in the Mu3e experiment with a timing resolution better than 100 ps. Over the course of this work, the Tile Detector was developed from the prototype stage to production readiness. The advancements in the detector design, the detector construction and the simulation of the cooling circuit conducted during this period of time is the subject of this thesis.

Thesis outline

The first three chapters of this thesis serve as an introduction, with the Standard Model and cLFV decays explained in Chapter 2 and a discussion of the Mu3e experiment following in Chapter 3. Chapter 3 also includes a short introduction to the Tile Detector and discusses the first technical prototype, which was developed and constructed at the very beginning of this work. The main work conducted in the context of this thesis is described in Chapters 4–6. The discussion of the Tile Detector construction is split into two chapters: Chapter 4 focuses on the updated detector design and the development

of a manufacturing line for the detector assembly. Chapter 5 then discusses the quality control (QC) measures implemented in the manufacturing line to ensure the optimal detector performance. Chapter 6 is concerned with the Tile Detector cooling studies, which were developed using the first technical prototype and then adapted and extended to the current design. Lastly, Chapter 7 concludes this thesis with a summary and outlook.

Author's contribution

The work covered in this thesis was conducted in the frame of the Mu3e collaboration, and within the collaboration as part of the Tile Detector group at the Kirchhoff Institute for Physics (KIP), Heidelberg University. In the following paragraphs, the individual contributions by the author to the topics discussed in this thesis are outlined.

Technical prototype and concept validation The concept and baseline design of the Tile Detector were developed in [4]. The design of the first technical prototype is based on this work and included for the first time the full detector front-end with the active sensor part, the read-out electronics, and the cooling circuit. At the very beginning of the author's work leading to this thesis, the author was played a key role in the construction of the first technical prototype presented in Section 3.3.2. This included contributions to the design and implementation of tooling for the detector assembly, as well as the assembly itself and subsequent tests of the prototype. In particular, the author co-organised and participated in two testbeam campaigns of the Tile Detector group at the testbeam facility of the Deutsches Elektron-Synchrotron (DESY), which proved the validity of the full Tile Detector design for the first time. Some of the results were published in a peer reviewed journal as part of a proceeding paper by the author [5].

Detector design and construction Starting from the technical prototype as a verified baseline design, several iterations, which culminated in the finalised detector design presented in Section 4.1, were overseen and implemented by the author. Accordingly, the author also contributed to a large fraction of the Tile Detector chapter published in [3]. The author was involved in the development of the new tile production method presented in Section 4.2 and evaluated the new tiles in a testbeam campaign at DESY. For the development of the tile matrix production line in Section 4.3, the author evaluated the assembly tools, which were manufactured by the KIP workshops based on the 3D-printed prototypes co-developed by the author. The author implemented the first version of the production line and commissioned it by assembling 62 tile matrices, which were used for the assembly of the two preproduction modules in Section 4.4 by the author. The author greatly profited from the work conducted by the Tile Detector group on the read-out electronics of the preproduction modules, which were evaluated in a testbeam campaign at DESY conducted by the Tile Detector group including the author.

Quality control The measurements for the QC of the scintillating tiles in Section 5.1 were performed by the KIP mechanical workshop. The author contributed to the eval-

uation of the measurements and to the improvement of the new tile production. The QC system for the silicon photomultiplier (SiPM) matrices in Section 5.2 was developed by the Tile Detector group, and evaluated and improved by the author during the QC of the 62 preproduction matrices. Similarly, the concept of the light yield QC system for the tile matrices in Section 5.3 was originally developed by the Tile Detector group and reworked by the author, with the author performing the calibration of the KLauS chip in Section 5.3.1. Making use of the author's work, the development of a common matrix QC system for both the SiPM QC and the light yield QC steps was conducted and evaluated by the Tile Detector group. Equally, the Tile Detector group has been, and still is, developing the concepts and systems for the module QC in Section 5.4.

Cooling studies The thermal simulation of the cooling circuit of the technical prototype in Section 6.1 was performed by the author. The author conducted the setup and the data taking of the laboratory measurements, and tuned the simulation to fit the data. Accompanying the design changes implemented by the author, the author adapted and extended the thermal simulation of the cooling circuits accordingly (Section 6.2). Some results of these simulations are also published in the proceeding paper by the author [5]. Furthermore, the author contributed to the setup of the final Tile Detector cooling tests using the test system at the KIP mechanical workshop (Section 6.3). The finalisation of the Tile Detector cooling system, including the design of the distribution manifold, was implemented by the Tile Detector group, with the author contributing in a consulting position.

Chapter 2

The Standard Model and beyond

This chapter provides the theoretical foundation for the following chapters, and gives an introduction to the theory of charged lepton flavour violation, which is the primary motivation for the Mu3e experiment. The starting point of this endeavour is the Standard Model, introduced in Section 2.1; its limitations lead into the field of charged lepton flavour violation in Section 2.2, with a selection of important lepton flavour violation experiments using muons introduced in Section 2.3. This chapter is largely based on [6] and [7], with other sources given where appropriate.

2.1 The Standard Model of particle physics

The Standard Model (SM) of particle physics is a Quantum Field Theory (QFT) which summarises the current knowledge of the fundamental constituents of the Universe and their interactions. It comprises 17 *elementary particles* and their interactions, the *forces*, which are described by their respective QFT. A summary of the particles in the SM is shown in the form of a sketch in Figure 2.1.

2.1.1 Fundamental particles and interactions

Based on their spin, the fundamental particles of the SM can be divided into two groups: the twelve *fermions*, which carry half-integer spins and behave according to Fermi-Dirac statistics; and the five *bosons* with integer spin, which follow the Bose-Einstein statistics.

Fermions

The fermions interact with each other through the fundamental forces. Based on their interactions, they can be further sub-divided into six *quarks* and six *leptons*. In addition, they can also be grouped into three *generations*: the first row in Figure 2.1, or first generation, consists of the up and down quarks, the electron, and the electron neutrino. These four particles constitute most of the matter in the low-energy Universe. The second and third generation particles, which are essentially copies of the first generation with similar properties but higher masses, occur typically only at higher energies, for example at particle colliders or in cosmic rays.

The six quarks partake in all three fundamental interactions present in the SM: they carry a fractional electromagnetic charge e of either $+\frac{2}{3}e$ (up-type quarks) or $-\frac{1}{3}e$ (down-




Figure 2.1: Visualisation of the Standard Model particles (sketch based on [6] and [8]). The fundamental particles of the SM are grouped into fermions, carrying half-integer spins, and bosons, carrying integer spins. The four gauge bosons are mediators of the three fundamental interactions formalised in the SM: the gluon is the mediator of the strong force, the photon is the mediator of the electromagnetic force, and the W and Z bosons mediate the weak force. The Higgs boson is responsible for the mass generation. The twelve fermions are separated into six quarks and six leptons, with three generations each. While the quarks take part in all three fundamental interactions, the leptons do not interact strongly, and in the case of the three neutrinos, only interact weakly. From the first to the third generation, the masses of the fermions increase.

type quarks), and therefore interact via the *electromagnetic* force; they take part in *weak* interactions; and they are the only elementary fermions carrying a *colour charge*, which allows them to interact through the *strong* force. Due to the asymptotic freedom of the strong force, quarks are only observed in the form of bound states called *hadrons*; two prominent examples of these hadrons are the proton and the neutron, with a quark composition of $p(uud)$ and $n(udd)$.

All of the six leptons interact through the weak force, but only three of them are electrically charged and partake in electromagnetic interactions as well. These three charged leptons are the electron, the muon, and the tau-lepton. The other three leptons – the neutrinos – are electrically neutral, and therefore only interact via the weak force.

As a consequence of their relativistic nature, all twelve fundamental fermions have an antiparticle counterpart. These antiparticle states have the same mass as the particle state, but carry the opposite electromagnetic charge. Unless explicitly denoted for clarity, the term “particle”, as well as all notations of fundamental fermions, implicitly include the antiparticle state as well.

Gauge bosons

In the SM, interactions between particles are understood as the exchange of spin-1 force-carrier particles, the *gauge bosons*. The three fundamental forces described by the SM are mediated by four gauge bosons: the photon for the electromagnetic interaction, the gluon for the strong interaction, and the W^\pm and Z bosons for the weak interaction.

In the QFT framework of the electromagnetic force, Quantum Electrodynamics (QED), a virtual photon is exchanged between particles which interact electromagnetically. As QED corresponds to a $U(1)$ local gauge symmetry, there is only one massless gauge boson, i.e. the photon. Quantum Chromodynamics (QCD), on the other hand, is the QFT describing strong interactions and invariant under $SU(3)$ transformations. Therefore, there are eight gauge bosons in QCD, the gluons, which are associated with the eight generators of the $SU(3)$ symmetry group. The conserved electromagnetic charge of QED corresponds to the colour charge of QCD. In contrast to the photon, which mediates the electromagnetic interactions but is itself electrically neutral, the gluons are colour-charged and therefore participate in the interactions they mediate. This behaviour gives rise to the QCD-inherent phenomena of asymptotic freedom and confinement.

Lastly, the weak interaction is mediated by the weak gauge bosons W^\pm and Z , which correspond to the generators of the $SU(2)$ symmetry group of the electroweak theory. Unlike the massless photon and gluons, the W and Z bosons are heavy, with masses of $m_W = 80.4 \text{ GeV}$ ¹ and $m_Z = 91.2 \text{ GeV}$. The W^\pm bosons carry an electromagnetic charge of $\pm 1e$ while the Z boson, also sometimes referred to as Z^0 , is electrically neutral.

¹This thesis follows the common convention of natural units in particle physics, that is $c = \hbar = 1$.

Higgs boson

The last fundamental particle left – and in fact, the last of the 17 particles to be discovered in 2012 by the ATLAS [9] and CMS [10] collaborations – is the *Higgs boson*. It has a mass of $m_H \approx 125$ MeV and, unlike the other bosons, carries a spin of 0 instead of 1, making it the only fundamental scalar particle discovered so far. Through the *Higgs mechanism* [11–13], and the associated Higgs boson, the weak gauge bosons and the fermions acquire their mass. The Higgs boson couples to massive particles, possibly also including itself.

2.1.2 Limitations of the Standard Model

The SM was, and still is, remarkably successful at predicting and describing the nature of particle physics. However, there are clear indications that the SM as a model is limited, and that there are phenomena which cannot be explained by it. For a start, only three out of the four forces are described by the SM; gravity is not included, suggesting that there might be a more fundamental underlying theory beyond the SM. For illustration purposes, a few more examples of open questions are listed below:

Dark matter

The most direct evidence of the existence of *dark matter* are velocity distributions of galaxies, which indicate that a large component of the galaxy mass is non-luminous (and thus “dark”). Furthermore, other observations such as precision measurements of the cosmic microwave background support evidence for the Λ_{CDM} cosmological model, which predicts the distribution of the energy-matter density of the Universe. In the Λ_{CDM} model, $\approx 5\%$ is made up by normal baryonic matter, $\approx 26.5\%$ consists of dark matter, and the remaining $\approx 68.5\%$ is in the form of *dark energy* [14].

It can also be deduced from cosmological models, and observations of the large scale structures in the Universe, that the main portion of the dark matter energy-mass density is non-relativistic. Therefore, neutrinos – which, due to their very small mass, remained relativistic throughout the expansion of the Universe – can be excluded as primary dark matter candidates. This gives rise to the assumption that there must be a new type of particle, summarised in the general term weakly interacting massive particle (WIMP). WIMPs are included in various new physics extensions to the SM, and WIMP searches are currently on-going, both at colliders like the Large Hadron Collider (LHC) and in dedicated dark matter experiments.

The nature of neutrinos

For a long time, neutrinos were assumed to be massless in the SM, until clear evidence for neutrino oscillations was observed by the Super-Kamiokande and SNO experiments at the end of the 1990s and the beginning of the 21st century [1, 2]. This suggests that neutrinos indeed do have mass, and that their *flavour eigenstates* (ν_e, ν_μ, ν_τ) which are

observed in weak interactions are superpositions of their *mass eigenstates* (ν_1, ν_2, ν_3). The masses of ν_1, ν_2 and ν_3 are not (yet) known; only an upper bound of 0.8 eV can be given through precision measurements of the Tritium β decay spectrum [15].

There are different ways to extend the SM in order to include a mass generation term for the neutrinos. One option would be to introduce a mass term similar to the charged fermion masses with a Yukawa coupling constant for the neutrinos. However, this raises the question why the neutrinos should be so much lighter, and their Yukawa coupling so much smaller, than that of the charged fermions if their mass is generated via the same mechanism. Another possibility is the addition of a *Majorana* mass term, which would allow the generation of very small neutrino masses via the *seesaw mechanism* by introducing a very massive neutrino state as a counterpart. This would in consequence make the neutrinos their own antiparticles, which could be detected e.g. by experiments searching for neutrino-less double beta decays.

Matter-antimatter asymmetry

The Universe – or at least, its observable part – is made primarily of matter, and there are no astrophysical observations to indicate the existence of an antimatter region somewhere in the Universe to counterbalance this. This gives rise to the assumption that some mechanism favoured matter over antimatter during the formation of the Universe. The only process in the SM to do so is CP violation, which has as of now only been observed in the weak interaction of quarks. A new generation of neutrino oscillation experiments could in the future detect CP violation in the leptonic sector. Still, it seems likely that the CP violation incorporated in the SM alone is not enough to explain the observed matter-antimatter asymmetry. Physics beyond the SM could be a solution to this open issue, possibly providing new and large CP-violating mechanisms.

The above examples show that physics beyond the SM is expected, but no clear evidence has been found yet. A popular new physics extension of the SM on the theory side is supersymmetry (SUSY), which pairs each SM particle with a super-symmetric partner particle, a *sparticle*. SUSY could potentially solve some issues of the SM, such as the *hierarchy problem*, where the Higgs mass is driven out of the electroweak mass scale by quantum loop corrections at very high mass scales. In addition, it introduces the neutralinos $\tilde{\chi}_i^0$ as partner sparticles to the Higgs and gauge bosons, of which the lightest, $\tilde{\chi}_1^0$, could be a WIMP candidate for dark matter.

On the experimental side, searches for beyond-SM physics are on-going in different directions: high energy searches, with the LHC at the forefront, search for new heavy particles directly generated at high energies. Complementary to the high energy experiments, on the other end of the experimental spectrum, lie the high intensity and high precision experiments. Among these is the Mu3e experiment, which searches for charged lepton flavour violation (cLFV) in decays of muons. The concept of cLFV and a selection of experimental searches for cLFV decays will be the focus of the remainder of this chapter.

2.2 Charged lepton flavour violation

In the SM, flavour changes are only allowed via the weak charged interaction, at least at tree level. In order to determine which flavour changing reactions are allowed and which are not, a conservation quantity was introduced for the leptons: the *lepton number*. It is defined as $L = +1$ for l^- , ν_l and $L = -1$ for l^+ , $\bar{\nu}_l$ (with $l = e, \mu, \tau$), while $L = 0$ for all other particles. This explains why a process like

$$\nu + n \rightarrow p^+ + e^-, \quad \Delta L = 0,$$

(an inverse beta decay, in a sense), is allowed, but the analogous reaction with an antineutrino

$$\bar{\nu} + n \rightarrow p^+ + e^-, \quad \Delta L = 2,$$

does not occur. As of now, the lepton number seems to be a conserved quantity in the SM. At this point, it is interesting to note that if neutrinos are found to be Majorana particles, lepton number conservation would be violated [16].

If only lepton number conservation (as well as charge and energy conservation) is considered, there seems to be no reason why a decay such as

$$\mu^+ \rightarrow e^+ + \gamma$$

should not be allowed, as $\Delta L = 0$. However, despite searches on-going since the late 1940s [17, 18], this process has not been experimentally observed yet. This led to the introduction of another quantity: the *lepton flavour* L_l . It is defined for each lepton generation $l = e, \mu, \tau$ as summarised in Table 2.1. In a decay like $\mu^+ \rightarrow e^+ + \gamma$, the lepton number may be conserved, but the lepton flavour is not, as both $\Delta L_\mu \neq 0$ and $\Delta L_e \neq 0$. Assuming that lepton flavour conservation holds, this process would be forbidden and therefore unobservable.

Table 2.1: Definition of the lepton flavour L_l , $l = e, \mu, \tau$ [16].

	L_e	L_μ	L_τ
e^\mp	± 1	0	0
$\nu_e/\bar{\nu}_e$	± 1	0	0
μ^\mp	0	± 1	0
$\nu_\mu/\bar{\nu}_\mu$	0	± 1	0
τ^\mp	0	0	± 1
$\nu_\tau/\bar{\nu}_\tau$	0	0	± 1




Figure 2.2: Feynman diagrams of the $\mu \rightarrow eee$ decay (created using [22]). In (a), a diagram for the $\mu \rightarrow eee$ process in the extended SM including neutrino masses is drawn. (b) shows a possible $\mu \rightarrow eee$ decay in SUSY. A neutralino $\tilde{\chi}^0$ mediates the interaction, with sleptons ($\tilde{\mu}$ and \tilde{e}) running in the loop.

It is evident from the observation of neutrino oscillations, however, that lepton flavour is clearly not conserved, at least not in the neutral sector. This makes charged lepton flavour violating decays such as $\mu \rightarrow e\gamma$ or $\mu \rightarrow eee$ possible via loop diagrams with neutrinos oscillating in the loop; a corresponding example of a Feynman diagram for the $\mu \rightarrow eee$ process is shown in Figure 2.2a. However, such decays are strongly suppressed with branching ratios of $\mathcal{O}(10^{-54})$ [16, 19, 20]. If models for new physics beyond the SM such as SUSY are taken into account, predictions for the branching ratios are enhanced to an experimentally accessible level. An example of a $\mu \rightarrow eee$ decay via SUSY slepton mixing is shown in Figure 2.2b. The SM background (including neutrino oscillations) of such a beyond-SM process is completely negligible, making cLFV searches ideal for probing beyond-SM physics [16, 21].

2.3 Lepton flavour violation experiments with muons

Already in 1948, muons were studied in the search for cLFV [17, 18]. Since then, muons have been consistently used in cLFV experiments, achieving ever higher sensitivities over the years, as illustrated in Figure 2.3. Muons are very well suited for cLFV searches for a number of reasons [23, 24]:

- Muons, in the form of cosmic rays, are available in abundance and independent of accelerator structures, making the early start of cLFV searches at the end of the 1940s possible.
- The muon lifetime of $2.2\mu\text{s}$ is relatively long, allowing the production of high intensity beams at meson production facilities.
- The muon has essentially only one decay mode, the Michel decay $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$ (and its radiative derivatives, with branching ratios $< 10^{-5}$). Moreover, the final

states are clear and well detectable, providing a clean environment for new physics searches.

In Table 2.2, the most relevant SM parameters of the muon are summarised.

Table 2.2: Key Standard Model parameters of the muon [25].

Parameter		Value
mass	$m_\mu[\text{MeV}]$	$105.6583755 \pm 0.0000023$
mean life time	$\tau_\mu[\mu\text{s}]$	2.1969811 ± 0.0000022
decay modes		
Michel decay	$\mathcal{B}(\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu)$	$\approx 100\%$
radiative decay	$\mathcal{B}(\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu \gamma)$	$(6.0 \pm 0.5) \cdot 10^{-8}$
internal conversion decay	$\mathcal{B}(\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu e^+ e^-)$	$(3.4 \pm 0.4) \cdot 10^{-5}$

The three main processes of cLFV involving muons are the anomalous muon decays $\mu \rightarrow e\gamma$ and $\mu \rightarrow eee$, and the muon-to-electron conversion in a nuclear field $\mu N \rightarrow eN$ [24]. Figure 2.3 shows a selection of past, present, and future experiments searching for cLFV in these three processes. Over the seven decades since the first searches at the end of the 1940s, the sensitivity at 90% confidence level (CL) of the experiments has increased by approximately twelve orders of magnitude. The current limit of the $\mu \rightarrow e\gamma$ decay is provided by the MEG experiment at PSI, with a branching ratio of $\mathcal{B}(\mu \rightarrow e\gamma) < 4.2 \cdot 10^{-13}$ at 90% CL [26]. The best limit for the $\mu \rightarrow eee$ decay was set by the SINDRUM experiment in 1988 at $\mathcal{B}(\mu \rightarrow eee) < 1.0 \cdot 10^{-12}$ (90% CL) [27], while the SINDRUM II collaboration obtained the best limit for the $\mu \rightarrow e$ conversion at $7 \cdot 10^{-13}$ (90% CL) in 2006, using muonic gold [28].

The open symbols in the bottom right corner of Figure 2.3 represent the sensitivities expected to be reached by future experiments. The MEG II experiment, an upgrade of MEG at PSI, will search for the process $\mu \rightarrow e\gamma$ with a projected sensitivity of $\approx 6.0 \cdot 10^{-14}$. MEG II already finished construction and is taking data since 2021. The MEG II collaboration has recently combined the results from the first physics run with the limit obtained by MEG to put the most stringent limit to date on the $\mu \rightarrow e\gamma$ decay, $\mathcal{B}(\mu \rightarrow e\gamma) < 3.1 \cdot 10^{-13}$ at 90% CL [30]. The Mu2e experiment [32] at Fermilab in the U.S. and the COMET experiment [33] at J-PARC, Japan, aim to set unprecedented limits of $\mathcal{O}(10^{-17} - 10^{-18})$ [31] on the $\mu N \rightarrow eN$ process. And lastly, the Mu3e experiment will search for the decay $\mu \rightarrow eee$ with a foreseen sensitivity of $2 \cdot 10^{-15}$ in phase I and 10^{-16} in phase II. The phase I Mu3e experiment, including the Tile Detector, will be the topic of the following chapter.




Figure 2.3: History of lepton flavour violation experiments using muons (adapted from [29], with data from [24, 26, 30, 31]). Over seven decades, the sensitivity of the experiments has increased by about twelve orders of magnitude. The open symbols on the bottom right are the expected sensitivities of future experiments currently under construction, or already taking data.

Chapter 3

The Mu3e experiment

The aim of the Mu3e experiment [34, 35] is the search for the process $\mu^+ \rightarrow e^+e^-e^+$ ¹ with a sensitivity of $2 \cdot 10^{-15}$ in phase I, and ultimately 10^{-16} in phase II, thus improving the current limit set by the SINDRUM experiment [27] by four orders of magnitude. In order to understand the experimental challenges, the signal and background processes of the $\mu \rightarrow eee$ decay are discussed in Section 3.1, followed by the resulting experimental design in Section 3.2. Section 3.3, lastly, addresses the Tile Detector in greater detail in preparation of the main body of this thesis in Chapters 4, 5 and 6.

3.1 Signal decay and background processes

As discussed in Section 2.2, the $\mu \rightarrow eee$ decay is forbidden in the SM at tree level, but possible via higher order corrections. A corresponding Feynman diagram for this process is shown in Figure 2.2a. The neutrino oscillations in the loop introduce a suppression proportional to $(\Delta m_{ij}/m_W)^4$, with the neutrino mass-squared differences Δm_{ij}^2 and the W boson mass m_W , resulting in a branching ratio of $\mathcal{B}(\mu \rightarrow eee) \approx 10^{-54}$ [21] which would not be observable. This clean probe makes searches for the $\mu \rightarrow eee$ decay sensitive to potential processes induced by physics beyond the SM.

Due to the strong suppression of the neutrino-mixing SM process, there is effectively no SM process with the same final state as the $\mu \rightarrow eee$ decay. The relevant background processes for $\mu \rightarrow eee$ mimic the signal, but can be suppressed by exploiting energy and momentum conservation. Broadly speaking, two categories of background can be distinguished: *internal conversion* background and *combinatorial* background. The following paragraphs will illustrate the challenges of signal detection and means of suppressing the different backgrounds, allowing us to draw conclusions regarding the experimental design.

Signal

In the $\mu \rightarrow eee$ process, the three decay electrons are produced promptly, thus originating from the same vertex, i.e. the same point in both space and time. This is illustrated in the sketch in Figure 3.1a. If the decaying muons are at rest, the vectorial sum of the

¹Henceforth shortened to $\mu \rightarrow eee$ for simplicity, unless the discussion requires explicit charges.

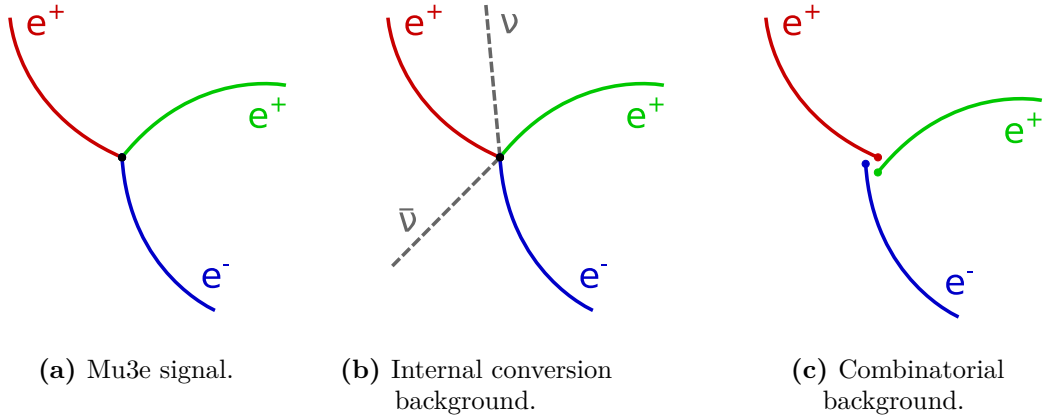


Figure 3.1: Sketches of the Mu3e signal and background. (a) The signal is characterised by three decay electrons originating from the same vertex, with their energies adding up to the muon mass. (b) The signature of the internal conversion background is very similar to the signal, but some of the energy is carried away by the neutrinos invisible to the particle detector. (c) Combinatorial background events are accidental overlays of several processes, resulting in a signature mimicking that of the signal. However, as these are different events, they are not necessarily coincident in space and time, and their energies typically do not add up to the muon mass.

decay particle momenta must vanish and their energy must add up to the muon mass [3]:

$$|\vec{p}_{\text{tot}}| = \left| \sum_{i=1}^3 \vec{p}_i \right| = 0, \quad E_{\text{tot}} = \sum_{i=1}^3 E_i = m_{\mu} \quad (3.1)$$

From the requirement on the momentum conservation, it follows that the momentum vectors of all decay particles must lie in a plane. The energies of the decay particles range from the electron mass, $E_e^{\text{min}} \approx 0.5 \text{ MeV}$, to half the muon mass, $E_e^{\text{max}} \approx 53 \text{ MeV}$. However, the exact decay dynamics depend on the underlying cLFV process and are therefore not known. A general parameterised Lagrangian, as presented in [36], can be used to simulate different signal types, depending on the relative strengths of different coupling types such as dipole, scalar- and vector-like couplings. One challenge for a potential $\mu \rightarrow eee$ experiment arising from these unknown kinematics is the need of a high signal acceptance in order to cover as much phase space as possible in the search for new physics.

Internal conversion background

If a muon decays radiatively, $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu \gamma$, combined with an internal conversion of the photon $\gamma \rightarrow e^- e^+$, the resulting process is called *internal conversion decay*

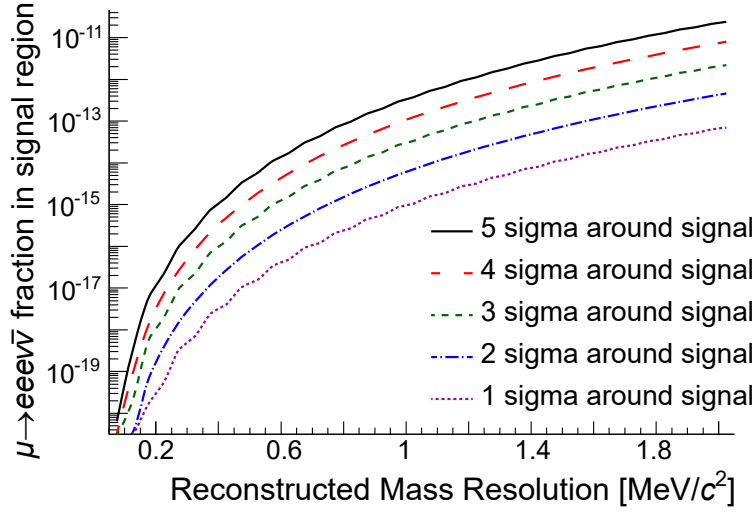


Figure 3.2: Contamination of the signal region with internal conversion events as a function of the reconstructed mass resolution of the three decay particles [3]. In order to reach the Mu3e phase I sensitivity goal of $2 \cdot 10^{-15}$ at a 2σ cut on the reconstructed mass, a resolution better than 1 MeV is needed.

$\mu^+ \rightarrow e^+e^-e^+\nu_e\bar{\nu}_\mu$. Its branching fraction is of the order of 10^{-5} (see Table 2.2). A sketch of the internal conversion decay is shown in Figure 3.1b. Its features are similar to those of the signal – three decay electrons coincident in space and time – with the addition of two neutrinos, which carry energy away undetected. Thus, the energy sum of the three observable decay electrons is smaller than the muon mass, and the vectorial sum of momenta does not vanish anymore:

$$|\vec{p}_{\text{tot}}| = \left| \sum_{i=1}^3 \vec{p}_i \right| \neq 0, \quad E_{\text{tot}} = \sum_{i=1}^3 E_i < m_\mu \quad (3.2)$$

Given an adequate energy and momentum resolution, it is possible to differentiate the internal conversion background from the signal. Figure 3.2 shows the fraction of internal conversion events in the signal region depending on the resolution of the reconstructed mass [3]. A momentum resolution better than 1 MeV is needed for a sensitivity of $2 \cdot 10^{-15}$ with a 2σ requirement on the reconstructed muon mass, which is the aim of phase I of the Mu3e experiment.

Combinatorial background

The second type of background is caused by an overlay of different processes resulting in three electron tracks which mimic the $\mu \rightarrow eee$ signature. One example would be the accidental combination of two Michel decays where one of the generated positrons

undergoes Bhabha scattering $e^+e^- \rightarrow e^+e^-$ with an electron in the detector material. Another type of event is a radiative muon decay combined with subsequent pair production by the photon in the target or detector material. Consequently, the vectorial sum of the three observable decay particles does not necessarily vanish, their energy typically does not add up to the muon mass, and they are usually not created coincident in time:

$$|\vec{p}_{\text{tot}}| = \left| \sum_{i=1}^3 \vec{p}_i \right| \neq 0, \quad E_{\text{tot}} = \sum_{i=1}^3 E_i \neq m_\mu, \quad t_1 \neq t_2 \neq t_3 \quad (3.3)$$

A corresponding sketch is shown in Figure 3.1c. There are different measures to counter the combinatorial background: a reduction of both the amount and the density of the target and innermost detector material helps to mitigate Bhabha scattering as well as pair production; an excellent vertex reconstruction allows the identification of the three production vertices; and a precise time resolution can distinguish the different muon decay times and is crucial for the suppression of this type of background, as discussed in [37] and [38].

3.2 Design of the Mu3e experiment

From the above considerations and requirements, the key points for the design of the Mu3e experiment can be derived:

1. In order to reach a sensitivity of $2 \cdot 10^{-15}$ (and ultimately, 10^{-16}), a high-intensity muon beam is needed. As the muons should decay at rest, the beam energy must be low. The $\pi E5$ beam line at PSI is currently the best choice to fulfil these conditions, providing muons at energies of ≈ 28 MeV at rates up to $1 \cdot 10^8$ Hz [3].
2. To study muon decays at rest, Mu3e is designed as a stopping target experiment. However, in order to reduce the amount of Bhabha scattering and pair production, the target design must find a compromise between stopping power and material minimisation.
3. A solenoidal magnetic field with a field strength of $\mathcal{O}(1 \text{ T})$ and very high homogeneity is needed for the precise momentum determination necessary to suppress internal conversion events.
4. As the maximum energy of the decay particles is limited to half the muon mass, the momentum resolution is limited by multiple Coulomb scattering. This puts stringent constraints on the design of the tracking detector, which must have a very low material budget while still providing excellent vertex and momentum resolution at a very high efficiency.
5. A timing system with excellent time resolution is needed to suppress combinatorial background and facilitate event reconstruction at very high muon rates.

6. In order to operate at the required intensities and to successfully suppress all background processes, a system with high rate capabilities and a high granularity is needed. This applies to the detector systems, but also affects other components of the experiment such as the data acquisition chain and the services.

The resulting experimental layout is shown in Figure 3.3. At the centre of the experiment is the stopping target, a hollow double cone made from aluminised Mylar foil. Surrounding the target and made up of two layers of ultra-thin silicon pixel tracker is the *vertex detector*, followed by the first of two dedicated timing systems, the *Scintillating Fibre (SciFi) detector*. Two more silicon tracker layers, the *outer pixel tracker*, complete the *central station*. Despite minimising the material budget of the silicon tracker and the SciFi, multiple Coulomb scattering is still the dominating factor for the momentum resolution. By choosing a magnetic field in beam direction with a field strength of 1 T, the outgoing decay electrons make a half turn and are bent back into the detector region – they *recurl*. Placing two sets of double pixel tracker layers upstream and downstream of the target increases the lever arm without the need to place the tracking layers at large radii (where the acceptance for low-momentum particles would be reduced). In addition, the effect of the multiple scattering cancels to first order after half a turn, leading to the narrow tube design sketched in Figure 3.3. The last detector system through which a decay particle passes is the second dedicated timing system, the *Tile Detector*, placed radially inwards of the recurl tracking layers.




Figure 3.3: Conceptual design of the Mu3e experiment [3]. A $\mu \rightarrow eee$ decay on the target is sketched, with the flight path of the two positrons in red and the electron in blue. The particles are first detected by the two inner pixel tracking layers of the vertex detector for precise vertex determination, then pass the low-density Scintillating Fibre timing system before crossing the two outer pixel tracking layers. They are then bent back into the detector volume by the 1 T solenoidal magnetic field, passing the two recurl tracking layers before reaching the Tile Detector timing system.

In the following paragraphs, the experimental infrastructure, the detector systems, and a selection of relevant services as well as the data acquisition system are described in more detail.

3.2.1 Muon beam, stopping target and magnet

Muon beam

The Mu3e experiment not only requires a high muon intensity, but also profits in terms of background suppression from a continuous beam structure, which provides a lower instantaneous muon rate than a pulsed beam (as used for example for the Mu2e experiment). The High Intensity Proton Accelerator (HIPA) facility at PSI [39] is the best available choice for the Mu3e experiment. It generates proton beams at energies of 590 MeV with up to 1.42 MW power, which hit a graphite target to produce pions [40]. The pions subsequently decay into muons, which are fed into beam lines such as the $\pi E5$ channel to be used for particle physics and solid state physics experiments.

Figure 3.4 shows a computer-aided design (CAD) model of the $\pi E5$ beam line and the experimental area where the Mu3e experiment will be located. Mu3e shares the $\pi E5$ channel with MEG/MEG II, which is permanently installed in the rear half of the $\pi E5$ area. In order to accommodate the 3.2 m long Mu3e magnet, the compact muon beam line (CMBL) was specifically designed for Mu3e [41]. It is connected to the front beam transport elements also used by MEG/MEG II, making it possible to efficiently switch between the two experiments [3]. In order to reach the muon stopping rate of

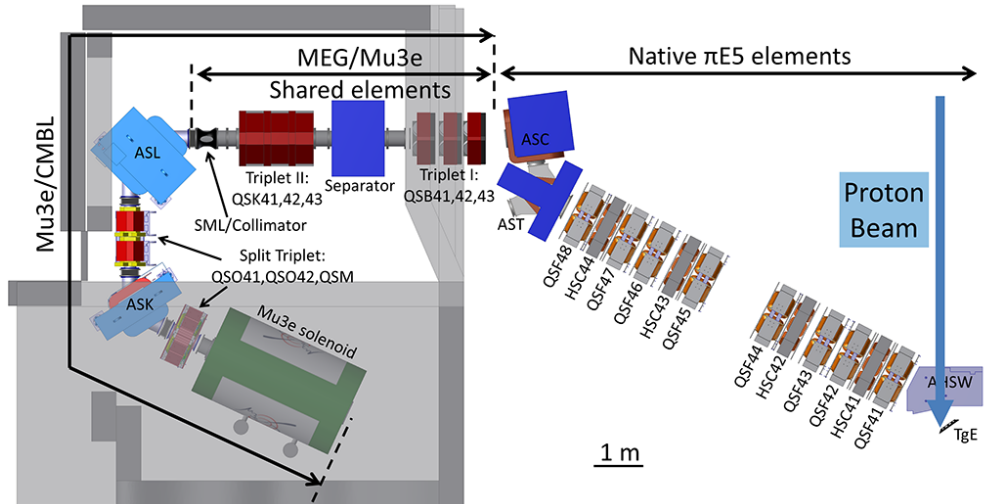


Figure 3.4: CAD model of the $\pi E5$ beam line and CMBL [3]. The $\pi E5$ channel, including the front beam transport elements, is shared by the Mu3e and MEG experiments. The CMBL is specifically designed for the Mu3e experiment in order to accommodate the Mu3e solenoid in the $\pi E5$ area.

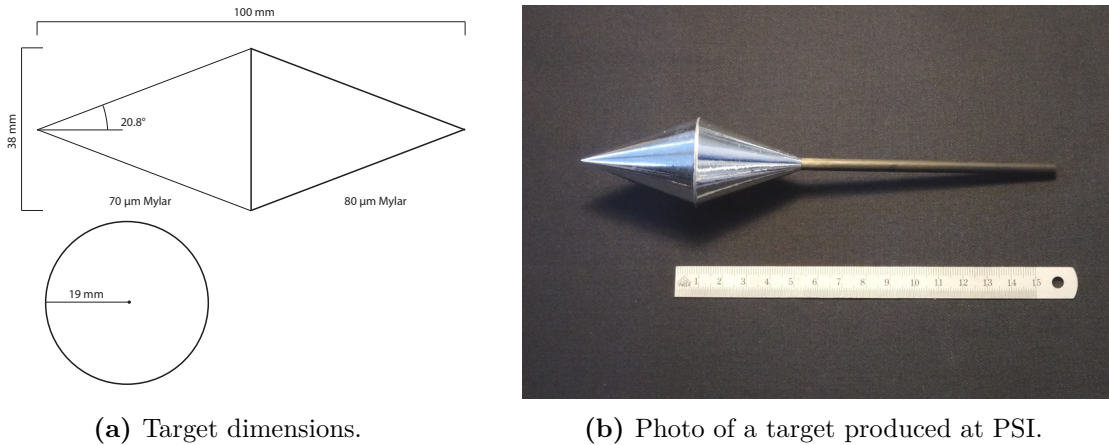


Figure 3.5: Stopping target for the Mu3e experiment [3]. The baseline design is a hollow double cone structure, similar to the target of the SINDRUM experiment. The exact dimensions are shown in (a), a finished target made from aluminised Mylar foil in (b).

$2 \cdot 10^9$ Hz needed for the Mu3e phase II sensitivity, an upgrade of the beam line is planned within the High-Intensity Muon Beams (HIMB) project [42], which is foreseen to start in 2027 [43].

Stopping target

As set out in the list of requirements for the Mu3e experiment, the Mu3e target must contain enough material to stop a large fraction of the muons. At the same time, the material amount and density must be minimised to reduce multiple scattering of the decay particles as well as background-inducing Bhabha scattering and pair production. Simulation studies of several different shapes yielded a hollow double cone structure as the baseline design, similar to the one used by the SINDRUM experiment, but made from 70 μm (80 μm) Mylar foil in the front (back) [3]. The target dimensions are sketched in Figure 3.5a, and a finished target produced at PSI is shown in Figure 3.5b. With this design, simulations show that about 95.5% of the muons that reach the target are stopped [3].

Magnet

The Mu3e magnet is a superconducting solenoid with a nominal field of 1 T, operating at a nominal temperature of 4 K [3]. The constraints on the inhomogeneities of the magnetic field along the beam direction are $\Delta B/B \leq 10^{-3}$ within ± 60 cm around the magnet center, with a long term stability of $\Delta B/B \leq 10^{-4}$ over 100 days. The magnet is about 3.2 m long and 2.0 m in width, with a warm bore of 2.7 m length and 1.0 m diameter. Its full weight is 31 tons, with 27 tons made up by the iron shielding reducing stray

fields to less than 5 mT at a distance of 1 m. As the Mu3e experiment will operate in a helium atmosphere (see Section 3.2.3), the warm bore is designed to be helium tight, with removable flanges in the magnet doors for access.

3.2.2 Detector systems

Tracking detectors

MuPix: At the basis of the Mu3e tracking detectors lies the MuPix sensor [44, 45], a high-voltage monolithic active pixel sensor (HV-MAPS) specifically developed for the Mu3e experiment. The *MuPix10* [46] is the first full-scale prototype with a size of $\approx 20 \text{ mm} \times 20 \text{ mm}$ and a pixel size of $80 \mu\text{m} \times 80 \mu\text{m}$. It provides a spatial resolution of $< 30 \mu\text{m}$ and fulfils the requirements posed by the Mu3e experiment with a thickness of $50 \mu\text{m}$, a hit detection efficiency of $\geq 99\%$ at a noise rate of less than 20 Hz/pix, and a timing resolution $\leq 20 \text{ ns}$ [45]. For the final chip, *MuPix11*, small improvements of the MuPix10 design have since been implemented and first tests have been performed, showing a promising performance [47].

Ladders: The stringent constraints on the material budget of the pixel tracking layers also affect the sensor connections and support structure. To reduce the amount of material compared to a typical printed circuit board (PCB), low-mass flex prints called high-density interconnect (HDI) circuits are used. They consist of aluminium-polyimide laminates and provide all electrical connections (power, high voltage, configuration, read-out) for the MuPix chips, which are directly glued and bonded to the HDIs. The structure of MuPix chips plus HDI is referred to as a *ladder* and forms the base element for all tracking layers. The total material budget of a ladder amounts to $X/X_0 = 0.115\%$, where X/X_0 is the detector thickness in units of the radiation length X_0 . This is sufficiently low to fulfil the Mu3e requirements.

Vertex detector: The pixel tracker of the Mu3e experiment can be divided into the vertex detector, consisting of the two inner layers around the target, and the outer tracker, which in turn can be divided into the central and recurl layers. The two layers of the vertex detector are placed at an inner radius of 23.3 mm and 29.8 mm, and consist of eight and ten ladders, respectively [3]. Six chips are assembled per ladder, yielding a length of 124.7 mm of the full vertex detector. The innermost layer of a vertex detector prototype using MuPix10 chips is shown in Figure 3.6.

Outer tracker: The outer layers of the central and the recurl stations are all identical in construction. The inner of the two layers is placed at an inner radius of 73.9 mm and consists of 24 ladders with 17 chips per ladder, while the outermost layer is positioned at 86.3 mm and consists of 28 ladders with 18 chips per ladder [3]. Adding the mounting structure for the ladders, the outer pixel tracker stations have a maximum length of 372.6 mm.





Figure 3.6: Innermost layer of the vertex detector prototype [48]. The six MuPix10 chips per ladder are visible. Due to an issue in the MuPix10 design, an adapted PCB design is used in this prototype instead of the HDIs.



(a) CAD model of the SciFi detector.



(b) Full size prototype of a SciFi ribbon.

Figure 3.7: Scintillating Fibre detector (images from [3]). (a) shows a CAD model of the SciFi, which consists of 12 fibre ribbons read out by an array of silicon photomultipliers on each end. In (b), a prototype of a full-size fibre ribbon is shown.

Timing detectors

SciFi detector: Situated in the central station between the vertex detector and the central outer tracker, the SciFi detector [37, 49] is optimised for a low material budget with $X/X_0 < 0.2\%$. At the same time, it is designed to provide a high timing

resolution of ≈ 250 ps and a spatial resolution of ≈ 100 μm at efficiencies exceeding 95% [3]. It consists of twelve fibre ribbons arranged in a cylindrical shape, as shown in Figure 3.7a. Each ribbon consists of three staggered layers of 250 μm diameter fibres, and is 300 mm long and 32.5 mm wide, matching the width of the 128-channel arrays of silicon photomultipliers (SiPMs) used for the read-out on both ends. A full-size prototype of a SciFi ribbon is shown in Figure 3.7b. The SiPM signals are read out and processed by the MuTRiG chip, an application-specific integrated circuit (ASIC) specifically designed for the two Mu3e timing detectors with a fast digital read-out for high rates.

Tile Detector: The Tile Detector is the second dedicated timing system of the Mu3e experiment and is positioned between the beam pipe and the outer tracking layers in the recurv stations. It is designed to provide a timing resolution of ≤ 100 ps and consists of scintillating tiles individually coupled to SiPMs, which are read out by the MuTRiG ASIC. Both the Tile Detector and the MuTRiG chip will be described in greater detail in Section 3.3.

3.2.3 Infrastructure

Detector cage

All detector systems, the beam pipe, and infrastructure such as cooling manifolds and power supplies for the low voltage (LV) and high voltage (HV) are installed in the *detector cage*. It is a frame structure consisting of two pairs of end wheels connected by aluminium struts to form a cylinder, which is inserted into the warm bore of the magnet via a rail system. At the centres of the end wheels, clamps are installed to mount the upstream and downstream beam pipes. The beam pipes in turn provide mounting points for the endrings needed for the installation of the detector systems. In addition, the copper bars providing the power for the central station and the ducts for the helium cooling of the central pixel detectors are assembled on the beam pipes. The end wheels also host the DC-DC converter boards for the LV power and the HV power supplies. Outside of the end wheels, two service support wheels (SSWs) are installed, which are only loosely coupled to the rest of the cage. On the SSWs, crates for the detector front-end boards as well as pipes for the cooling and patch panels for the power connections are mounted. A CAD rendering of the detector cage is shown in Figure 3.8.

Cooling systems

Helium cooling: With an estimated power density of ≈ 215 mW/cm² [48], the MuPix chips require active detector cooling. The design of the pixel tracker cooling system is again governed by the tight constraints on the material budget. Conventional water cooling or state-of-the-art bi-phase CO₂ cooling requires pipes and a method to thermally couple to the sensors, thus adding more material. A gaseous cooling approach is therefore chosen, where a stream of cooled gas is guided through the detector. The simplest solution, using air as a cooling medium, is excluded since this would add $0.33\% \cdot X_0$ to




Figure 3.8: CAD rendering of the detector cage (as of June, 2023). The detectors (in green and blue) are mounted on the upstream and downstream beam pipes at the central axis of the cage. The beam pipes are held in position by clamps at the centres of the two end wheels. The upstream beam pipe is connected to the π E5 beam line via the beam flange. The service support wheels are located outside of the end wheels and host service infrastructure such as the crates for the front-end boards and routing for the cooling pipes. A preliminary design of the helium cooling connections via the pink tubes is implemented as well.

the path of a recurling particle with a track length of approximately 1 m [48]. Given the requirement of a gas lighter than air, helium is the best choice, as it has a comparably high thermal conductivity combined with a very low material budget, and is safer to operate than hydrogen. Thus, a novel helium cooling system is designed for the Mu3e experiment, and has been tested with thermal-mechanical mock-ups in [47, 48].

Liquid cooling: The remaining subsystems of the Mu3e experiment which require cooling are the SciFi detector, the Tile Detector, and electronic components such as power supplies and read-out boards. As the cooling circuits of the two detector systems are positioned such that they do not affect the particle trajectories, and the electronics are hosted outside of the active volume in the end wheels and SSWs, liquid cooling systems with conventional pipes and metal cooling structures can be employed. In the baseline design presented in [3], water is used as a cooling medium. However, irradiation studies for the SiPMs of both the SciFi [50] and the Tile Detector [51] have shown that cooling below $-10\text{ }^\circ\text{C}$ is required in order to keep the SiPM dark count rate and power consumption at a reasonable level. A dedicated cooling system for the timing detectors is therefore currently being adapted to work with sub-zero temperatures. For

the electronics, a second liquid cooling system using a water/glycol mixture as a medium is designed to operate at room temperature.

Data acquisition system

The data acquisition (DAQ) system [3, 52] of the Mu3e experiment consists of three stages of triggerless read-out: the first stage are the detector front-end boards (FEBs), which continuously collect zero-suppressed hit information, perform time-sorting and package the data before sending it to the switching boards. The four switching boards – one each for the central pixel tracker, the SciFi, the upstream recurl station and the downstream recurl station – form the second level of the DAQ. Their purpose is to synchronously merge the time-sorted data streams from the different FEBs and subsequently forward the merged data to the third stage, the filter farm. Here, buffering and event building is performed on FPGA-level with the option to include simple clustering, sorting and selection algorithms. Afterwards, the events are transferred to GPUs, which perform a simplified fast track fitting and vertex selection. If an interesting event is identified by the event filter, the corresponding data is transferred from the FPGA buffers to the central DAQ PC and subsequently written to disk to be analysed at the central PSI computing centre. The three stages of the DAQ system reduce the data rate from about 100 Gbit/s zero-suppressed hits at $1 \cdot 10^8 \mu\text{s}$ to 100 MB/s of data for the offline analysis.

Simulation framework

The experimental parameters listed above – e.g. detector resolutions, beam intensity, and background estimations – have been implemented in the Geant4-based Mu3e software stack. It consists of a simulation which includes different muon decay generators (see [3] for a detailed description), a track reconstruction algorithm [53], a vertex fitter and a variety of analyses to estimate for example the detector performance.

A cut-based analysis, described in detail in [3], is used to estimate the signal efficiency and background suppression for Mu3e phase I. In order to reach the projected phase I sensitivity of $2 \cdot 10^{-15}$, more than $2.5 \cdot 10^{15}$ muon stops are needed at an estimated efficiency of $\approx 20\%$. After factoring in the detector acceptance and applying the selection criteria on parameters such as vertex fit quality, centre-of-mass momentum p_{cms} and reconstructed mass m_{rec} , a signal efficiency of 13% is reached. At the same time, the background events in the signal region can be suppressed to 0.57 ± 0.04 internal conversion events and 1.9 ± 1.4 Bhabha+Michel events (for $103 \text{ MeV} < m_{\text{rec}} < 110 \text{ MeV}$ and $p_{\text{cms}} < 4 \text{ MeV}$ at $2.5 \cdot 10^{15}$ muon stops) [3]. As shown in Figure 3.9a, the accidental Bhabha+Michel background can be further reduced if a more stringent lower cut on the reconstructed mass, $104 \text{ MeV} < m_{\text{rec}}$, is applied. In this case, the internal conversion events are reduced to 0.068 ± 0.013 , thus showing that the Mu3e experiment can run virtually background-free for the $2.5 \cdot 10^{15}$ muon stops of phase I. The expected single event sensitivity (SES) versus days of running time is shown in 3.9b.

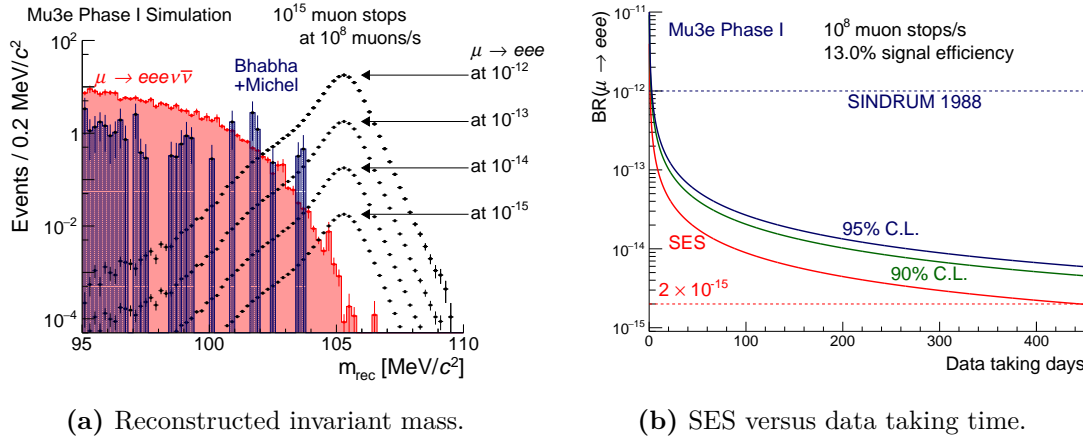



Figure 3.9: Simulated performance of the Mu3e experiment for phase I [3]. **(a)** shows the reconstructed invariant mass at various branching ratios, with a cut on the centre-of-mass momentum of $p_{\text{cms}} < 4$ MeV. The internal conversion background is shown in red, and the dominating accidental background of Bhabha scattering plus Michel decay is shown in blue. For a mass cut of $104 \text{ MeV} < m_{\text{rec}}$, the Mu3e experiment can run background-free for at least $2.5 \cdot 10^{15}$ muon stops. In **(b)**, the single event sensitivity of Mu3e phase I is shown together with the corresponding 90% and 95% CL limits versus the data taking time. After about 400 days of data taking, the planned SES of phase I is reached.

3.3 Introduction to the Tile Detector

The purpose of the Tile Detector is to provide the most precise timing measurement within the Mu3e experiment. The aim of the detector is therefore to achieve a timing resolution of ≤ 100 ps at an efficiency of close to 100%. Due to its position at the very end of the recurling particle trajectories, the stringent constraints on the material budget to avoid multiple scattering do not affect the detector design. However, space constraints require the detector to be very compact, as will be discussed in Section 4.1.1. The base design for the Tile Detector was developed in the frame of [4]. While changes and updates facilitating the production and experimental integration of the detector have been implemented since then (some in the frame of this thesis, which are discussed in Chapter 4), the original detector concept remains the same. It is introduced in the next section (Section 3.3.1), with a focus on the active detector elements, consisting of scintillating tiles and silicon photomultipliers connected to a custom read-out ASIC. In Section 3.3.2, the first technical prototype is discussed, which was assembled based on the original detector concept in [4]. The construction of the technical prototype not only marks the beginning of the author's doctoral studies, but also laid the groundwork for the main work presented in this thesis in Chapters 4, 5 and 6.



(a) Tile Detector submodule.

(b) Tile Detector station.

Figure 3.10: Original Tile Detector design, taken from [4]. In (a), the base unit of the original design, called a submodule, is shown. It consists of two 4×4 tile-SiPM arrays assembled on a common PCB with one 32-channel read-out ASIC (not visible in the figure). The flexible elements of the PCB (in dark yellow) connect the read-out ASIC to the two arrays, and allow the folding of the submodule into the depicted shape. The submodules are attached to aluminium support/cooling plates, which in turn are assembled on a pair of endrings to form (b) a cylindrical recurl station.

3.3.1 Original detector concept

The fundamental concept of the Tile Detector, as developed in [4], remains valid to date: tiles manufactured from a plastic scintillator with a short rise time are individually coupled to silicon photomultipliers, which are read out by a custom-designed ASIC. The active detector elements are assembled in a modular approach: two 4×4 tile-SiPM combinations and one 32-channel read-out ASIC on a shared PCB with flex print elements form a base unit (called a *submodule* in [4]). These base units are assembled on aluminium support plates, which also act as cooling elements. U-shaped pipes are embedded into the plates and connected to the liquid cooling system discussed in Section 3.2.3. The plates are assembled on two pairs of endrings to yield two barrels, which form the two Tile Detector recurl stations. CAD models of a submodule and a detector station, as presented in [4], are shown in Figure 3.10. A detailed description of the detector geometry and mechanics, including the updated design, will be given in Chapter 4.

Plastic scintillator

Plastic scintillators are a sub-group of organic scintillators. They are typically produced infusing a solvent base material, such as polyvinyl toluene (PVT) or polystyrene, with a primary and a secondary fluor (meaning fluorescence emitter) [54]. The base material is a scintillator and initially absorbs the energy of a traversing particle, but has only a low light yield. The primary agent is therefore a second scintillator with a high quantum efficiency, increasing the overall yield and shortening the rise time of the light pulse. As




Figure 3.11: Base ingredients of the Tile Detector. (a) The tiles are manufactured from EJ-228, a fast plastic scintillator, and individually coupled to (b) SiPMs of the type S13360-3050VE from Hamamatsu. (c) The MuTRiG, a custom 32-channel ASIC designed in-house for fast read-out and optimal timing, forms the base unit of the Tile Detector with 32 tile-SiPM elements.

Table 3.1: Properties of the EJ-228 plastic scintillator [55].

Property	Value
Scintillation efficiency (photons/1 MeV e^-)	10 200
Wavelength of maximum emission (nm)	391
Rise time (ns)	0.5
Decay time (ns)	1.4
Density (g/cm^3)	1.023
Softening point ($^\circ\text{C}$)	75
Temperature range ($^\circ\text{C}$)	-60 to 60

the fluorescence light of the primary fluor is typically in the UV range, the second fluor acts as a wavelength shifter and moves the emission spectrum into the blue/near-UV region, where photo-sensors are more sensitive.

Plastic scintillators are easy to handle and machine, are relatively cheap and can be cast in a variety of shapes, as they are liquid before polymerisation. In addition, they provide a relatively high light yield at very fast rise and decay times. This makes them very well suited for timing detectors, such as the Tile Detector. Extensive simulations of different plastic scintillators were performed in [4]. The PVT-based EJ-228 [55] from Eljen Technology was found to be the most suitable choice of all the scintillators that were studied, as it offers both a high light yield and a fast response time, and thus

provides the best time resolution ². Its properties are listed in Table 3.1. Figure 3.11a shows an EJ-228 tile which was cut to the correct Tile Detector dimensions (roughly 6 mm × 6 mm × 5 mm).

Silicon photomultiplier

A SiPM is a silicon-based solid-state photo-detector with single photon sensitivity. Its working principle is based on an avalanche photodiode (APD), a semiconductor device which allows the detection of photons through avalanches created in a high internal electric field. If an APD is operated above the breakdown voltage V_{br} of the diode, the amplification process of any photon signal is so strong ($\mathcal{O}(10^5 - 10^6)$) that a large discharge current is always flowing through the device, independent of the original amount of photons. This means that the device is not sensitive to the magnitude of the original signal anymore, but is able to detect individual photons instead. Hence, it is referred to as single photon avalanche diode (SPAD) or Geiger avalanche photodiode (G-APD), as the operation above V_{br} is called Geiger mode. In order to quench the discharge after the breakdown, a quenching resistor (*passive quenching*) or a transistor circuit (*active quenching*) can be added, which both reduce the voltage on the diode enough to stop the avalanche process. An array of G-APDs connected in parallel, each with their own quenching resistor, forms a typical SiPM.

Due to the operation in Geiger mode, each of the individual pixels of a SiPM works as a binary counter for photon detection, and can therefore not differentiate between a single and several impacting photons. However, if the photons are distributed over the sensor surface, the discharge currents of the pixels are added to a common output signal which is proportional to the number of photons. This is illustrated in Figure 3.12, which shows a typical charge spectrum of a SiPM at low light intensities. Each of the distinct peaks corresponds to a specific number of fired pixels, or photoelectrons, as marked in the figure, thus making it possible to count single photons. This spectrum is therefore also referred to as *single-photon spectrum (SPS)*. In order to ensure the detection of the individual photons, SiPMs are typically designed with small pixel sizes of a few tenths of μm distributed over areas with dimensions of a few mm. As for the scintillator material, extensive simulation studies were performed for the best choice of SiPM in [4]. A sensor size of 3 mm × 3 mm was found to be the best choice for the Tile Detector, as it offers a reasonable trade-off between a moderate signal decay time and a sufficient time resolution. Similarly, a pixel size of 50 μm was chosen as a compromise between a high photon detection efficiency, a pixel capacitance which still yields acceptable signal decay times, and a sufficient number of pixels to detect the expected number of photons. The selected SiPM is the MPPC S13360-3050VE [56] from Hamamatsu, shown in Figure 3.11b. It has a peak sensitivity wavelength of 450 nm, which sufficiently matches the EJ-228 maximum emission wavelength of 391 nm (see Table 3.1). Its photosensitive area is 3 mm × 3 mm

²Note that in [4], the scintillator BC-418 was studied, which is a commercial equivalent to EJ-228.

with a pixel size of $50\ \mu\text{m} \times 50\ \mu\text{m}$, resulting in 3 584 pixel cells per device. A selection of relevant SiPM properties is summarised in Table 3.2.




Figure 3.12: Single-photon spectrum of a SiPM [57]. The integrated charge, in units of the charge-to-digital converter (QDC) used for the measurement, is shown. Each of the peaks corresponds to a specific number of photoelectrons, i.e. fired pixels of the SiPM. The shaded peak at 0 photoelectrons corresponds to an offset of the read-out electronics.

Table 3.2: Properties of the MPPC S13360-3050VE [56]. Measurement conditions: ambient temperature $T_a = 25\ \text{°C}$, overvoltage $V_{\text{over}} = 3\ \text{V}$.

Property	Value
Effective photosensitive area (mm^2)	$3 \cdot 3$
Pixel size (μm^2)	$50 \cdot 50$
Number of pixels	3 584
Operating temperature ($^{\circ}\text{C}$)	-20 to 60
Peak sensitivity wavelength (nm)	450
Photon detection efficiency	40 %
Gain	$1.7 \cdot 10^6$
Breakdown voltage (V)	53 ± 5
Recommended operating voltage (V)	$V_{\text{br}} + 3$
Dark count (MHz)	0.5 (typ.) 1.5 (max.)

Read-out ASIC

A custom mixed-signal ASIC called MuTRiG (Muon-Timing Resolver including Gigabit link) [58] was developed at the Kirchhoff Institute for Physics in Heidelberg for the read-out of both Mu3e timing systems. The MuTRiG chip evolved from STiC (Silicon Photomultiplier Timing Chip) [59], a 64-channel ASIC with excellent timing capabilities developed for medical applications in the frame of the EndoTOFPET-US project. Both the analog front-end and the time-to-digital converter (TDC) with 50 ps bin size of STiC are inherited by MuTRiG in order to fulfill the precise timing requirements of the Mu3e timing detectors. In addition, a new digital part for fast read-out was developed specifically for MuTRiG in order to handle the high event rates of $\mathcal{O}(1 \text{ MHz/channel})$ of the SciFi detector.

In Figure 3.13, the read-out principle of MuTRiG is sketched. Two discriminator units with tunable thresholds, the *timing threshold* and the *energy threshold*, are used to digitize the analog SiPM signal. From the rising edge of the *T trigger* generated by the lower timing threshold, the time of arrival of the signal with respect to a common reset t_0 can be extracted. The higher energy threshold, which generates the *E trigger*, allows for the validation of physical signals and the suppression of SiPM dark counts. In

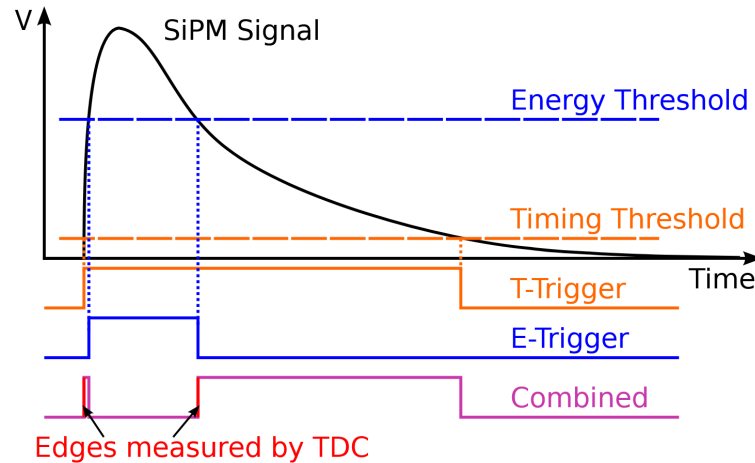


Figure 3.13: Principle of the signal discrimination in the MuTRiG chip [58]. Two discriminator units with tunable thresholds are implemented in the analog front-end. When the leading edge of an analog SiPM signal passes the lower timing threshold, a timestamp is generated by the rising edge of the T trigger. A second timestamp is generated by the falling edge of the E trigger once the SiPM signal decreases below the energy threshold. Both T and E trigger are combined to a common output processed by the TDC, which extracts the time of arrival of the event from the first timestamp, as well as the event energy in the form of a linearised ToT from the time difference between first and second timestamp.

addition, the time difference between the rising edge of the T trigger and the falling edge of the E trigger provides information on the signal energy by using a linearised time-over-threshold (ToT) method [60]. The outputs of the two discriminators are combined into one signal processed by the TDC, with the rising T trigger edge and the falling E trigger edge transformed into two rising edges in order to preserve the timing and energy information.

3.3.2 First technical prototype

Prior to the work conducted in the frame of this thesis, the Tile Detector design was developed in [4], and extensively studied in simulations as well as validated with early prototype measurements [4, 61–65]. The development, assembly, and validation of the first technical prototype of the Tile Detector marks the beginning of the work leading to this thesis. The purpose of the technical prototype was, on the one hand, to establish production and assembly tools and procedures, and on the other hand to evaluate the full detector concept, including cooling and support structures as well as read-out electronics. In contrast to the earlier prototypes, the technical prototype therefore consisted for the first time of submodules, endrings and a cooling circuit as defined in [4]. A deviation from the baseline design was needed in terms of the read-out ASIC: as the MuTRiG was not available at the time of production, the newest version of its predecessor, STiC3.1, was used instead. Given the same analog front-end and TDC of MuTRiG and STiC, this is a valid substitute for the performance evaluation of the prototype. The first version of a read-out board for the submodules was designed and produced as well.

Overall, three submodules with a total of 96 tile-SiPM channels were built and tested using an electron beam at the Deutsches Elektron-Synchrotron (DESY) testbeam facility [66] in Hamburg. Two of them were assembled on a cooling plate mounted on a pair of endrings, as foreseen in the Tile Detector design. The third submodule was assembled on a custom cooling block and installed in front of the Tile Detector-like setup in such a way that one of the 4×4 tile-SiPM arrays was aligned in parallel to the beam to serve as a time reference. A sketch of the setup is shown in Figure 3.14a. It was placed on a large moving stage, with the reference submodule additionally mounted on a smaller moving stage. The main Tile Detector-like system could also be rotated along two angles with respect to the beam axis. This setup allowed the time reference to always remain in the centre of the beam while measuring different areas of the other two submodules under various incident angles of the beam particles. The two submodules on the Tile Detector setup were read out by the newly developed read-out board prototype, whose data stream was merged with that of the reference submodule on a custom FPGA board.

Over the course of two testbeam campaigns at DESY in 2018, the technical prototype was extensively calibrated, evaluated and optimised for the best timing performance, as described in [3] and [5]. In order to evaluate the time resolution of individual channels, the coincidence time resolution (CTR) between at least three channels at a time was measured. The single-channel time resolution σ_i of a channel i can then be extracted




Figure 3.14: Testbeam measurements with the first technical prototype of the Tile Detector [5]. In (a), the testbeam setup is sketched, consisting of two submodules assembled on a cooling plate and endrings in a Tile Detector-like structure and a third submodule as a timing reference. The full setup, and the two sub-setups, were movable and rotatable with respect to each and other and to the incident electron beam. (b) shows the single-channel time resolution, which was extracted from the testbeam data following Equation 3.4. Only channels from the same submodule contribute to the red distribution, while channels from the timing reference are taken into account in the calculation of the blue distribution. The additional jitter of approximately 46 ps between the two curves is caused by the non-optimal design of a part in the DAQ chain.

from the CTR between the selected channel and two further channels j and k as

$$\sigma_i = \sqrt{\frac{\sigma_{i,j}^2 + \sigma_{i,k}^2 - \sigma_{j,k}^2}{2}}, \quad (3.4)$$

$$\sigma_{i,j}^2 = \sigma_i^2 + \sigma_j^2 \quad (\text{and } \{i, k\}, \{j, k\} \text{ equivalently}),$$

where $\sigma_{i,j}$, $\sigma_{i,k}$ and $\sigma_{j,k}$ represent the CTR of the respective channel combinations. By aligning the detector along the beam direction, it is possible to calculate the single-channel time resolution of all channels of one submodule by using only channels read out by the same chip. The distribution of this *internal* timing performance is shown in red in Figure 3.14b. It yields an average single-channel time resolution of (46.8 ± 7.6) ps. Adding *external* channels of other submodules (such as the reference system) to the calculation in Equation 3.4 results in an additional average time jitter of (45.5 ± 3.2) ps, shown in the blue curve in Figure 3.14b. This is caused by the non-optimal design of the FPGA board used to merge the data streams. However, even when including this extra jitter, the timing performance of the technical prototype is still below the required resolution of 100 ps.

In addition to the performance evaluation, the cooling system was tested and validated, and the functionality of the read-out electronics and the DAQ chain was verified. This successful evaluation of the full detector concept paved the way for the next step in the Tile Detector development: advancing from the production of individual prototype components to a manufacturing line for mass production, towards the finalisation of full-scale detector construction. The manufacturing line, which was developed based on the technical prototype during this thesis and which was used in the production of two full Tile Detector modules, will be presented in Chapter 4. The development and evaluation of quality control procedures accompanying the detector production will be discussed in Chapter 5; and the cooling simulation developed and optimised using the first technical prototype will be the topic of Chapter 6.

Chapter 4

Tile Detector construction

The construction and performance evaluation of the first technical prototype, as described in Section 3.3.2, validated the full Tile Detector concept developed in [4]. Thus, the design of the prototype, as well as the manufacturing and assembly processes developed during its production, serve as the basis for the construction of the detector as presented in this chapter. As the detector design was adapted and modified in several iterations with the first technical prototype as the baseline, Section 4.1 presents the updated Tile Detector concept and discusses the motivation for these changes. Sections 4.2–4.4 then elaborate on the construction and assembly of the updated detector design: the production of the scintillating tiles is discussed in Section 4.2, while the subject of Section 4.3 is the development of a production line for the tile matrices, the smallest modular base elements of the Tile Detector. Lastly, the assembly and performance evaluation of two preproduction modules are described in Section 4.4.

4.1 Updated detector design

The modular composition of the Tile Detector, which has already briefly been mentioned during the introduction of the original detector concept (Section 3.3.1), is described in more detail in the following paragraphs, starting with the full detector and subsequently increasing the granularity of the detector elements down to the smallest building blocks. In Section 4.1.1, considerations relevant for the design choices and/or changes of the different detector elements are discussed.

Tile station

The Tile Detector consists of two independent *stations*, placed upstream and downstream of the target as shown in Figure 3.3. One station has 2 912 read-out channels, resulting in 5 824 channels for the full Tile Detector. Each station is approximately 342 mm in length with a diameter of 128 mm, and can be split into seven *modules* mounted on two *endrings* (see Figure 4.1a). The endrings in turn are mounted directly on the beam pipe, as depicted in Figure 4.1b, and are manufactured from carbon fibre reinforced polyether ether ketone (PEEK) with a thickness of 3.5 mm. At the given thickness and endring design, this choice of material provides a suitable rigidity as opposed to e.g. aluminium, which was used for the endrings of the first technical prototype. In addition, choosing a




Figure 4.1: CAD model of a Tile Detector station. In (a), an exploded view of a station is shown with approximate outer dimensions. One station is made up of seven modules, with one of them highlighted in the picture. The modules are mounted on a pair of endrings, one of which is visible in (b) the front view of the station. The mounting structures on the beam pipe, on which the endrings are assembled, are pointed out as well.

plastic instead of a metal as the endring material avoids a closed conducting loop, and thus prevents the flow of eddy currents in the magnetic field of the Mu3e experiment.

Tile module

A module, shown in Figure 4.2, is an assembly of a cooling/support plate, 26 *tile matrices* (in two rows at 13 matrices each) and a *Tile Module Board (TMB)* hosting the electronics to read out 416 channels per module. The matrices are mounted on the radially outward side of the cooling plate. In this way, the active material of the Tile Detector (i.e. scintillators and SiPMs) is hit first by the recurling electrons before they traverse the support plate and services further inwards, ultimately stopping in the beam pipe. The TMB is attached to the face of the cooling plate pointing radially inwards, where the cooling pipe is embedded. In this way, the heat produced by the read-out electronics on the TMB, which makes up the bulk of the Tile Detector heat load, can be most effectively cooled. Given the sevenfold geometry of the Tile Detector, the angle between two modules is designed to be $360^\circ/7 \approx 51.4^\circ$; the angle between two matrices, both within a module and between neighbouring modules, is consequently $360^\circ/14 \approx 25.7^\circ$.

Read-out electronics

The read-out board of the Tile Detector, the TMB, is shown in Figure 4.3. In total, 13 MuTRiG ASICs in a commercial BGA package are assembled on one TMB. Each MuTRiG processes the signals of 32 tile-SiPM channels, resulting in the above mentioned




Figure 4.2: CAD model of a Tile Detector module. (a) shows an isometric view of a module, which consists of 416 tile-SiPM channels assembled on an aluminium support/cooling structure with embedded stainless steel pipes. In (b), a front view of a module is shown. The active detector material (tiles and SiPMs) is placed on the radially outward side of the support plate, while the read-out electronics, including the MuTRiG ASICs, are hosted on the TMB mounted on the radially inward side. The connection between active detector part and read-out electronics is provided by flex print connectors.

416 channels per TMB/module. The TMB provides both the HV and LV power needed for the SiPMs and the ASICs, respectively. It also hosts a number of components such as distribution chips for the clock and reset signals, temperature sensors, power monitors, and furthermore transfers the digitised data of the MuTRiGs to the Mu3e DAQ system. The 26 matrices that can be read out by one TMB are connected using low-height flexible printed circuit (FPC) connectors, as shown for one matrix in Figure 4.3.

Cooling and support

The cooling circuit of a Tile Detector module, shown in Figure 4.4, consists of a support plate with a stainless-steel pipe embedded in a U-shaped channel on the bottom side. The plate is machined from aluminium, which provides a good heat conductance at a relatively low weight. Counterbore holes and precision holes for alignment pins are milled into the top side of the plate for the tile matrix assembly, shown in Figure 4.4a. The pipe is glued into the channel on the bottom side, shown in Figure 4.4b, with copper shavings mixed into the glue to increase the thermal conductivity between pipe and plate. The pipe inlet and outlet are connected to the liquid cooling system used for both the SciFi detector and the Tile Detector, which is foreseen to operate at approximately -20°C . The liquid cooling system, together with design considerations such as pipe dimensions and thermodynamic parameters, is discussed in detail in Chapter 6.






Figure 4.3: Photo of a TMB. The 13 MuTRiGs, assembled in BGA packages, are connected to the tile matrices via low-height FPC connectors, here shown for one exemplary matrix. On the edge of the photo, the signal connector (on the bottom-left) and the multi-coloured LV connections are visible. In addition, the TMB also provides the HV for the SiPMs via four thin coaxial cables, of which one exemplary cable is shown at the top (in black).



(a) Front/top view.




(b) Bottom view.


Figure 4.4: Photo of a Tile Detector cooling circuit, consisting of an aluminium support plate and a cooling pipe. (a) The regularly spaced counterbore holes on the top side of the plate are used for the tile matrix assembly. (b) The stainless steel pipe is glued into a U-shaped groove on the bottom side of the plate.

Tile matrix

The smallest modular building block of the Tile Detector is the *tile matrix*. It consists of 16 plastic scintillator cubes attached to 16 SiPMs, arranged in a 4×4 grid on a custom-designed PCB with a flex print connection, as shown in Figure 4.5. A pair of matrices is connected to the 32 channels of one MuTRiG, resulting in 26 matrices read out by 13 MuTRiGs per module. Attached to the matrix PCB are two cylindrical nuts with an internal thread, which are used to screw the tile matrix to the cooling plate. As shown in Figure 4.5b, the scintillator cubes are wrapped in reflective foil before they are glued to the SiPMs, in order to reduce optical cross-talk between neighbouring tiles. Both the wrapping and the gluing procedures will be explained in greater detail in Section 4.3.2.



(a) CAD model of a tile matrix.



(b) Photo of a tile matrix.

Figure 4.5: The smallest modular building block of the Tile Detector, the tile matrix, shown as (a) a CAD model and (b) as a photo. The scintillator tiles are wrapped in reflective foil and glued in 4×4 pattern onto the SiPMs. The custom-designed PCB provides a flex print connection for the signal transfer to the MuTRiG ASIC.

4.1.1 Design considerations

Over the course of this thesis, the Tile Detector design has been subject to various alterations and revisions, changing from the technical prototype shown in Section 3.3.2 to the design presented above. Some of the modifications were driven by insights gained during the mechanical assembly, while others were necessary due to conflicts with other detector systems or the Mu3e infrastructure. Furthermore, constraints by the read-out electronics, which were developed during that time, had to be taken into account. The following paragraphs will point out the most prominent changes and give a short overview of the considerations concerning the respective new design.




Figure 4.6: Cross section of the Mu3e CAD model (downstream side). The outline of the Tile Detector is shown in blue. Radially inwards lie the ducts for the helium cooling, the copper bars carrying the power and ground for the central pixel layers and the SciFi detector, and ultimately the beam pipe. Radially outwards, the Tile Detector is restricted by the two layers of the recurl outer pixel tracker.

Overall detector size and dimensions

The overall Tile Detector size is driven by the detector and infrastructure systems surrounding it. This is illustrated, for the case of the Tile Detector diameter, in Figure 4.6: radially inwards, the Tile Detector is limited by the beam pipe and the services for the central detectors, namely the ground and power lines (in the form of copper bus bars glued to the beam pipe), the helium cooling (in the form of ducts glued on top of the bus bars), and signal lines as well as the SciFi cooling pipes (routed through three channels between the helium ducts). Radially outwards, the confinement is given by the inner radius of the recurl outer tracker, with some safety margin. In terms of length, the Tile Detector design is limited by the position of the endrings of the recurl outer tracker, and thus by the length of the outer tracker ladders, which is approximately 372 mm. Taking the space for the mounting structures of the ladders on the inside of the outer tracker endrings into account results in a length of approximately 350 mm allocated for the Tile Detector.

The constraints for both the length and the radial position/expansion of the Tile Detector have changed with respect to the baseline design in [4], as infrastructure and detector components such as helium ducts or the outer pixel endrings have been newly

introduced since or undergone major modifications. Therefore, length and thickness of the Tile Detector have been repeatedly adjusted and reduced according to the experimental situation, all the time leaving the scintillator size unchanged as to not compromise the timing resolution. To facilitate this process, and in order to spot conflicts early on in the design, an envelope of the Tile Detector was defined and included in the Mu3e CAD model (shown in Figure 4.6). The volume within this envelope is allocated to the detector and allows a first estimate of whether design changes are compatible with the full CAD model, especially at the interfaces to the other systems or the infrastructure. However, the use of the CAD model for discovering conflicts between the different subsystems is limited. Therefore, dedicated cabling and routing workshops within the Mu3e collaboration are regularly conducted using mock-ups of detectors and infrastructure in various stages, in order to identify potential showstoppers in the assembly process early on.

TMB and matrix PCB design


Given the limitations on the radius elaborated above, one of the most crucial design changes with regard to the technical prototype is the merging of the MuTRiGs and the prototype read-out board, resulting in the TMB in its current form. It reduced the thickness of the Tile Detector read-out electronics by 7.2 mm, from approximately 9.5 mm (MuTRiG on submodule PCB plus prototype read-out board) down to 2.3 mm (TMB with assembled MuTRiGs), thus freeing the space now occupied by the helium ducts.

Another change which goes hand in hand with the merging of the MuTRiGs and the TMB is the splitting of a submodule into separate tile-SiPM elements and the MuTRiG, thus creating the current design of the tile matrices. It allows the exchange of one or more matrices without having to exchange the respective MuTRiG (or the potentially still working partner matrix). As studies concerning irradiation damage of the SiPMs suggest that the matrices closest to the target may need to be exchanged before the end of the phase I run time [51], mechanically decoupling the MuTRiG from the matrices will reduce the amount of exchanged detector material by replacing only the problematic matrix (or matrices).


Support plate

Similarly to the read-out electronics, the support plate was slimmed down by around 2 mm to create more space for the services close to the beam pipe. As a result, the pipe size was reduced as well; the impact of this design choice on the cooling is discussed in Chapter 6.


Another significant change in the support plate layout is driven by an improvement of the matrix assembly method compared to the technical prototype. In the technical prototype design, the assembly points of one submodule are covered by the neighbouring submodule, as sketched in Figure 4.7a. This allows only one fixed sequence of assembly (and, consequently, disassembly), meaning that if the first submodule in the sequence needed replacement, all other submodules would have to be disassembled first. Another



(a) Sketch demonstrating the sequential (dis-)assembly of the submodules.



(b) Sketch of the submodule assembly.



(c) Sketch of the tile matrix assembly.

Figure 4.7: CAD based sketches showing the assembly differences of the technical prototype submodule and the new tile matrix design. **(a)** In the technical prototype design, the mounting points of a submodule are covered by the neighbouring submodule. As a result, only a sequential (dis-)assembly of the submodules is possible. **(b)** The submodule must be assembled in a fixed sequence: first, one of the PCBs with a tile-SiPM array, then the ASIC PCB, lastly the other tile-SiPM PCB. As the length of the flex print connections between tile-SiPM arrays and ASIC PCB is limited, the second tile-SiPM PCB must be assembled at an angle, making it difficult to position the screws and increasing the risk of damaging the submodule. The issues in (a) and (b) are avoided by the new matrix design in (c), which allows the individual assembly of the matrices due to the separation of the tile-SiPM subsystems and the ASIC.

challenge of the technical prototype assembly arises from the submodule PCB design: As the submodules must be mounted as a whole, the best sequence of assembly is first to fix one 4×4 tile-SiPM PCB to the support plate, then to fold the MuTRiG PCB around the plate and attach it, and lastly to fix the other 4×4 tile-SiPM PCB. However, due to the limited length of the flex print connections, the second tile-SiPM PCB cannot be positioned in parallel to the support plate surface during the assembly, as sketched in Figure 4.7b. This can be mitigated by selecting very short screws, such that mounting the submodules is in principle possible; still, the procedure is difficult and increases the

risk of damaging the tile-SiPM connection during the assembly, in the worst case even breaking tiles off.

With the change from submodule and prototype read-out board to tile matrices and TMB, through-holes were added to the support plate, which allow the assembly of individual matrices from the back side (that is, the TMB side) of the plate. Counterbore holes on the matrix side of the plate (as shown in Figure 4.4) allow the insertion of the cylindrical nuts, which are attached to the tile matrix. Due to the internal thread of the nuts, the matrix can be fastened to the support plate and correctly aligned using tapered screws, which are countersunk on the back side of the plate. The general assembly idea is sketched in Figure 4.7c. This new approach is not without drawbacks: the TMB now covers the matrix screws, which thus adds the disadvantage of having to remove the TMB first before disassembling any matrices. Still, it solves the assembly challenges discussed above and allows for more flexibility than the original technical prototype.

Matrix and scintillating tile sizes

The structure of the tile matrix is largely inherited from the technical prototype, which in turn is derived from the baseline design in [4]. This also includes the dimensions of the 4×4 tile-SiPM array. The pitch of the SiPMs on the PCB is 6.41 mm in z direction (along the beam) and 6.5 mm in φ direction (i.e. in the transverse plane). Keeping in mind that the thickness of the reflective wrapping foil – approximately 80 μm – must be taken into account, the baseline for the tile width and length is thus defined as $l_{\varphi,\text{tile}} = 6.3$ mm and $l_{z,\text{tile}} = 6.21$ mm. This leaves a safety margin of 200 μm for the $\approx 2 \cdot 80$ μm wrapping thickness between two tiles, provided they are correctly centred on the SiPMs. Simulation studies in [4] found a tile height of $h_{\text{tile}} = 5$ mm to be the best compromise between a high timing resolution and a low hit rate and pile-up. Lastly, the angle of the tiles on the edges of a matrix is defined as $\alpha_{\text{tile}} \approx 77^\circ$, deriving from the angle between two matrices, which must be $360^\circ/14 \approx 25.7^\circ$ to form a cylindrical shape (as discussed above in the updated detector design). A sketch of the two tile types (with and without angle) with the respective dimensions is shown in Figure 4.8a. The width and length of a tile matrix can then be derived as $l_{\varphi,\text{matrix}} = 26.0$ mm and $l_{z,\text{matrix}} = 25.64$ mm, which are sketched in Figure 4.8b and 4.8c, respectively.

Ideally, the matrices should be put as close together as possible in order to cover a maximised area, with as little dead space in between as possible. In terms of mechanical assembly, however, experience has shown that additional space between two neighbouring matrices is needed; especially during the construction of the first technical prototype, the tiles were prone to breaking off as placement tolerances had not been adequately taken into account. As simulation studies in [4] and [38] have shown that the tiles are most likely hit under impact angles $\neq 0^\circ$, and that the influence of the inactive area between tiles is thus negligible if it is small with respect to the tile size, an additional gap of 150 μm was added between matrices in both φ and z directions to facilitate the assembly. To accommodate for this extra space in φ direction, the radius at which the Tile Detector is placed has to be increased by 0.34 mm. This increase was cross-checked




Figure 4.8: Definition of the dimensions of the scintillating tiles and the tile matrices. (a) shows the two different tile types, which share the same base area, given by the width $l_{\varphi,\text{tile}} = 6.3$ mm and the length $l_{z,\text{tile}} = 6.21$ mm. The height $h_{\text{tile}} = 5$ mm is also the same for both types. The angle of the edge tile, which is needed for the detector to form a barrel, is $\alpha_{\text{tile}} \approx 77^\circ$. Derived from the tile sizes are (b) the matrix width $l_{\varphi,\text{matrix}} = 26.0$ mm and (c) the matrix length $l_{z,\text{matrix}} = 25.64$ mm.

in the Mu3e CAD model and deemed feasible, as the constraints posed by the recurv pixel layers on the outer Tile Detector radius are not as tight as those on the inner radius. In z direction, the additional gap results in a tile matrix pitch of 25.79 mm, and thus in an active area with the length of ≈ 335 mm. Adding the thickness of the endrings yields the full Tile Detector length of ≈ 342 mm, which is well inside the available space of 350 mm discussed above. An extensive summary of the Tile Detector dimensions can be found in Table A.1 and Figure A.1 in Appendix A. The finalisation of the relevant dimensions of the Tile Detector allows the start of the next step, which is the actual production of the detector.

4.2 Scintillating tile production

4.2.1 Development of a new production method

One of the initial steps to building the Tile Detector is the production of the scintillator tiles. In this regard, several points must be kept in mind:

- **Dimensions:**

The size of the scintillating tiles must be precise enough to allow the assembly of the tiles in the pitch and with the gap size given in Section 4.1, while at the same time covering as much of the Tile Detector surface as possible.

- **Surface quality:**

The scintillator surfaces should be polished well enough to ensure a good reflectivity, and thus a good light collection efficiency of the SiPM attached to the tile.

- **Machining process:**

As the softening point of the scintillator is located around 75 °C [55], a coolant is needed during machining. However, typical metal cutting oils must be avoided, as this would damage the material. Instead, fine (soapy) water sprays are recommended.

In order to meet the requirements listed above, the scintillating tiles produced for the first technical prototype were cut and polished by hand in the mechanical workshop of the Kirchhoff Institute for Physics (KIP), resulting in a production time of approximately 70 working hours for 100 tiles. Extrapolated to the almost 6000 tiles needed for the full detector, this would result in a working time of around two to three years (for a single person), and was thus deemed not feasible. Instead, a new method was developed in collaboration with the KIP mechanical workshop to greatly speed up tile production, while keeping in line with the requirements listed above. This new method is illustrated in Figure 4.9 (for the example of one matrix) and makes use of two clamping mechanisms in order to produce a large amount of tiles in one machining cycle.

The starting point for the production is a scintillator plate which was saw-cut to a size of 262 mm × 36 mm × 5.5 mm by the manufacturer. The top and bottom faces are diamond-milled with a thickness tolerance of ±0.18 mm. The width of the plate is chosen such that one tile matrix can be cut out, while the length corresponds to that of eight matrices plus a margin to account for the width of the milling tool. This plate is clamped down to a vacuum plate and 5 mm deep trenches are cut into it, as shown in Fig 4.9a. The angled faces of the edge tiles are milled by first cutting small steps to define the correct angle, and then using a conical milling tool to smoothen the surface. As the depth of the trenches is only 5 mm, the tiles are not fully cut out of the plate, but instead remain connected by a sheet of scintillator of 0.5 mm thickness.

In the next step, the plate is flipped, with the tiles now facing downwards and the thin connecting sheet of scintillator at the top, and mounted on the second clamping tool, called *ice-vice* [67]. The ice-vice is a clamping plate that is cooled down to −10 °C using pressurised air, thus making it possible to freeze delicate parts into place for milling, instead of conventionally clamping them by force. In the case of the Tile Detector scintillator production, the trenches between the tiles are filled with water, freezing them to the ice-vice. Once the full scintillator plate is fixed, the thin 0.5 mm sheet still connecting all the tiles together is milled away from the top, as shown in Figure 4.9b. The setup is then brought back to room temperature and only the finished tiles remain, as shown for a single matrix in Figure 4.9c. In the final production process, one plate - yielding eight matrices - was machined per day. The production time for the full Tile Detector, i.e. 364 matrices (plus spares), was thereby reduced from the original estimate of two to three years to approximately four months.



(a) Scintillator milling on the vacuum plate.

(b) Scintillator milling on the ice-vice.

(c) Final matrix.

Figure 4.9: Scintillating tile production of one matrix. **(a)** In the first step, the tiles needed for one tile matrix are cut from a scintillator plate, which is clamped on a vacuum plate. A small amount of material (thickness 0.5 mm) is left at the bottom to hold the tiles together. **(b)** The tiles are flipped and clamped on the ice-vice by filling the gaps with water and then lowering the temperature. Once everything is frozen into place, the 0.5 mm sheet of scintillator (now at the top) is milled away until **(c)** only the tiles of a full matrix remain.

4.2.2 Performance studies of the new tiles

Before the production of the full amount of tiles was approved, a small sample of tiles was manufactured using the new ice-vice method in order to compare it to the old hand-polished tiles used in the technical prototype. Figure 4.10 shows the two tile types side-by-side: the left tile was manually cut and diamond-polished, while the right tile was machined using the ice-vice. It is notable that the surfaces of the ice-vice tile are more opaque than those of the hand-polished prototype tile. An exception is the bottom surface of the ice-vice tile (marked in Figure 4.10), which is slightly clearer than the tile's sides. This surface is the only part of the tile which is not machined during the new production method, and thus remains in the state in which the scintillator plates are delivered by the manufacturer. It is therefore selected as the face which is glued to the SiPM to ensure a good light collection efficiency.

In order to investigate whether the properties of the new ice-vice tiles are comparable to those of the tiles used in the technical prototype, tile matrix prototypes were assembled with the different tile types and taken to the DESY testbeam facility [66] to study

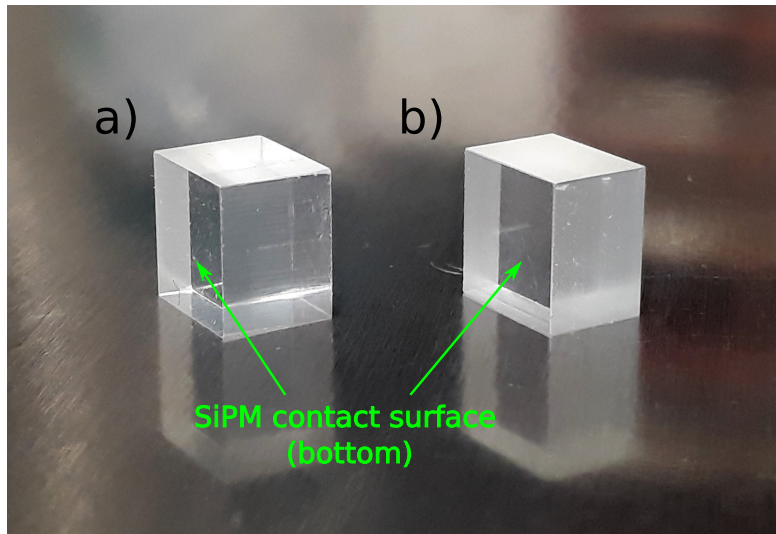


Figure 4.10: Comparison of the two tile production methods. The tile in a) was cut and diamond-polished by hand, while the tile in b) was machined with the new production method using the ice-vice. The surfaces of the old, handmade tile are distinctly clearer than those of the new ice-vice tile with the exception of the bottom surface, which is left in the state as delivered by the manufacturer of the raw scintillator plates. This surface is therefore chosen as the contact surface for the SiPMs.

the timing performance. In Figure 4.11, a sketch of the setup used for the testbeam measurement is shown. It is based on a MuTRiG stand-alone system, consisting of a carrier board with a wire-bonded MuTRiG ASIC, a motherboard with the necessary connections for the MuTRiG operation (supply power, clock and reset signals, etc.), and an FPGA with a custom DAQ software which was used prior to the Mu3e DAQ system. This stand-alone setup was originally created for the STiC ASIC, and was used for the first tests of the MuTRiG before the TMB was developed. For this testbeam, it was connected to two tile matrices via a custom-made adapter board, with the low-height FPC connectors for the matrices on one end, and connections for the carrier board on the other side. Furthermore, the adapter board was equipped with two coaxial connectors for the SiPM HV, and pin connectors for the I2C read-out of the matrix temperature sensors.

During the testbeam campaign, the setup was arranged such that the particle beam was passing through both matrices assembled on the adapter board, as sketched in Figure 4.11. In this way, eight tiles in a row were hit, defining a *tower*. The first matrix reached by the beam served as a reference, while the second matrix was the one investigated; during the testbeam, this position was first held by the matrix assembled with the old prototype tiles, and later by the one equipped with the new ice-vice tiles. In this way, a direct comparison of the two matrices was possible, without the need to exclude

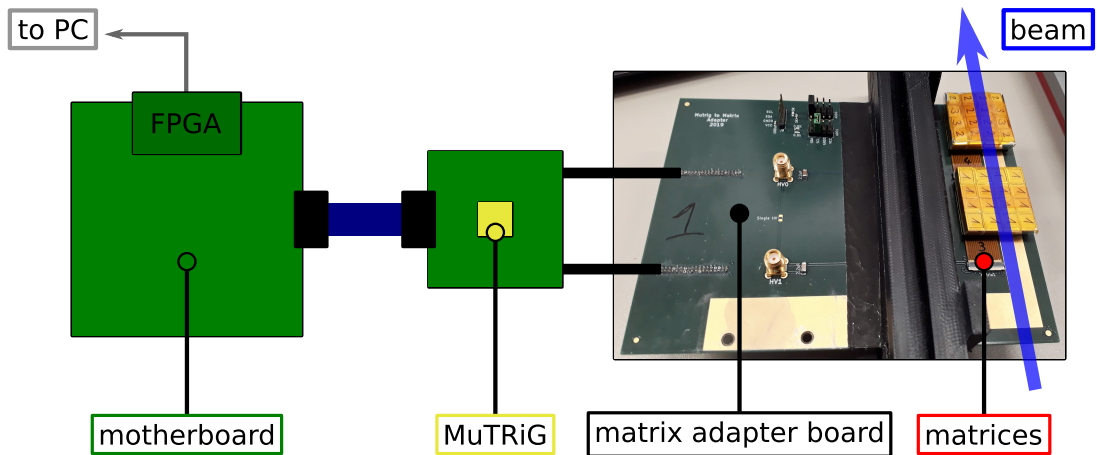


Figure 4.11: Sketch of the setup used to investigate the different tile types at the DESY testbeam facility. Two matrices are connected to one MuTRiG using an adapter board, which provides the temperature read-out as well as the HV connections for the SiPM operation. The MuTRiG is wire-bonded on a carrier PCB connected via a high-speed twinax cable to a motherboard, which delivers the ASIC supply voltages and clock and sends the ASIC data via an FPGA board to a PC for data acquisition.

MuTRiG-specific channel-by-channel variations. In order to compare the two tile types in terms of timing performance, the matrices were calibrated using the following method: first, one matrix was set to a rough working point, while a parameter scan was performed for the other matrix; then, the process was repeated vice versa. The parameters chosen for this investigation were the SiPM HV, and the time and energy threshold of the MuTRiG, which were all selected for optimum timing resolution. The timing performance of a given channel i was then determined by calculating the single-channel time resolution defined in Equation 3.4 for all 21 permutations in the same tower as the channel in question, and subsequently averaging the results to arrive at the final average single-channel time resolution of channel i .

The method described above was applied to all channels of the two tile matrices under investigation. The resulting single-channel time resolutions of the $2 \cdot 16$ channels are shown in Fig. 4.12; no energy cuts or time-walk corrections are applied. With the new ice-vice tile matrix, an average resolution of 35.4 ps with a standard deviation of 3.0 ps can be achieved, while the matrix with the old prototype tiles reaches an average resolution of 37.0 ps with a standard deviation of 3.0 ps. In conclusion, the tiles manufactured with the new production method perform equally well as the old hand-polished tiles. The new ice-vice method is therefore deemed to be suitable for the full tile production.

Lastly, based on these testbeam results, a conclusion can be drawn concerning the general performance of the tile matrix design: With the former submodule design of




Figure 4.12: Comparison of the timing resolution achieved with the old prototype tiles and new ice-vice tiles at the DESY testbeam. The 16 channels of the matrix equipped with the prototype tiles have an average single-channel time resolution of (37.0 ± 3.0) ps, while for the matrix with the ice-vice tiles, an average resolution of (35.4 ± 3.0) ps is measured.

the first technical prototype, an average single-channel time resolution of (46.8 ± 7.6) ps within one submodule (i.e. 32 channels read out by one MuTRiG) could be achieved at a similar DESY testbeam campaign (Section 3.3.2). Comparing this average value, as well as Figure 3.14b, to the results achieved with the tile matrix measurements, a clear improvement both in the absolute value of the time resolution and in the spread of the individual channels can be observed, thus validating the new matrix design.

4.3 Development of a tile matrix production line

The production of the smallest building block of the Tile Detector, the tile matrix, can be separated into two steps: first, the electronic and mechanical components are assembled on the PCB of the matrix, the *matrix-flex*. This includes the SiPMs and a temperature sensor, as well as the two cylindrical nuts needed to attach the final tile matrix to the cooling plate, also referred to as *spacers*. Following the assembly of the matrix-flex, the scintillating tiles are glued to the SiPMs. Prior to the gluing step, the tiles are individually wrapped in reflective foil in order to reduce optical cross-talk between neighbouring channels. Once the curing time of the glue has passed, the matrix is ready for module assembly.

The two production steps described above were developed during the construction of the first technical prototype of the Tile Detector. With the change of the tile matrix design, they were accordingly adapted, and quality control (QC) procedures were developed and integrated in the production process. The finalised production line will be used for the assembly and QC of the 364 matrices (plus spares) needed for the full detector, and will be done fully in-house at KIP in Heidelberg. In the following sections, the two production steps are described in more detail; the QC procedures will be explained in Chapter 5. All assembly steps, tools, and methods were thoroughly tested and finalised during the construction of two *preproduction modules*; they therefore constitute the first mass production test for the Tile Detector, thus providing valuable input to establish the full production line for the final detector.

4.3.1 SiPM matrix

In the first production step, the matrix-flex PCB is equipped with all electronic components, including the 16 SiPMs per matrix; this intermediate production state is therefore referred to as *SiPM matrix*. Before the start of the assembly, the SiPMs are sorted into groups of 16 based on their operating voltage V_{op} , according to a data sheet provided by Hamamatsu. Figure 4.13 shows the distribution of V_{op} of the 1000 SiPMs used in the




Figure 4.13: Distribution of the operating voltage V_{op} of the 1000 SiPMs used for the production of the two preproduction modules (data provided by Hamamatsu). The SiPMs can be sorted into two groups, shown in blue and yellow, with a difference of ≈ 0.5 V between the groups. One preproduction module was assembled with sensors from the blue group, while for the second module, sensors from the yellow group were used.




Figure 4.14: Photo of an assembled SiPM matrix, with 16 SiPMs arranged in a 4×4 grid. The small black component in the centre of the matrix is the temperature sensor; the metal cylinders to the left and right of the temperature sensor are the spacers used for the mechanical assembly of the matrix on the support plate.

production of the two preproduction modules. Two groups are visible, with a difference in V_{op} of approximately 0.5 V. This is attributed to variations in the production, in particular between different wafers. While the first group contains 580 SiPMs, the second group consists of only 420 SiPMs - just enough to produce the 26 matrices needed for one full module. It was therefore decided to group the matrices for the two preproduction modules according to the spread of operating voltages of the SiPMs in order to have a more consistent distribution of V_{op} within the modules.

Once the SiPMs have been selected, the matrices are assembled with all electronic components plus the mechanical spacers in batches of 10 matrices in the electronics workshop of KIP. Figure 4.14 shows a photo of an assembled SiPM matrix. The matrices are numbered (first by hand, as seen in Figure 4.14, later using a QR code) in accordance with their respective entries in the Mu3e internal database, which is used to keep track of parts and production progress, as well as to store results from QC tests of the individual matrices in their various production states.

In order to proceed to the next production step, the SiPM matrices need to pass a first QC test. For the two preproduction modules, a simplified QC procedure – described in Section 5.2 – is used to test the general functionality of the matrix by measuring the SiPM noise rate and by reading out the temperature sensor. If the temperature and the noise rates are as expected, the matrix passes the test; if not, the matrix is investigated and reworked.

4.3.2 Tile matrix

The second production step consists of the wrapping and gluing of the scintillating tiles. For both processes, custom tools were developed using CAD and 3D-printing for prototyping. After the prototyping phase, the tools were finalised and produced with the help of the KIP mechanical and electronic workshops for increased precision, and semi-automation in case of the wrapping tools. In the next paragraphs, the wrapping and gluing tools and procedures will be discussed, following the actual order of production as it was performed for the two preproduction modules.

Tile wrapping

The scintillating tiles are wrapped with enhanced specular reflector (ESR) foil to reduce optical cross-talk between neighbouring tiles. The wrapping step is mostly detached from the gluing process and can therefore be done in advance for a large amount of tiles. First, the ESR foil must be cut into a shape which allows the folding around a tile; the chosen shape is shown in Figure 4.15a. A window of $3.2\text{ mm} \times 3.2\text{ mm}$ is cut in the centre of the foil for the SiPM. To obtain a foil in this shape, a laser cutter is used which was tuned to fully detach the outlines of the foils, and to remove only a small amount of material along the folding edges without cutting through to allow for easier folding. Per sheet of ESR foil ¹, around ≈ 150 individual foils can be cut in approximately 45 minutes using the laser cutter, resulting in a working time of only a few hours for the two preproduction modules and a few days for the full Tile Detector. Before the foils are removed from their sheet to be used for wrapping, a piece of adhesive tape is applied to the square end of each foil (seen in Figure 4.15a as a red square). When the tape has been applied to all foils in a sheet, all preparations for the wrapping are finished and the sheet can be either stored away for later use, or immediately used in the wrapping process using the custom-developed wrapping tools.

The wrapping tools exist in two versions: one for the centre tiles and one for the edge tiles. Both tools are shown in Figure 4.16a, with the white centre tile tool on the left and the red edge tile tool on the right. The key difference of the two versions lies in the central aluminium block, which is the core of both tools; a zoomed-in picture of one of the blocks is shown in Figure 4.16b. In this aluminium block, a funnel is cut in the shape of the respective (centre or edge) tile. When an unfolded foil is placed above the funnel with a tile on top, the top lever of the tool pushes both tile and foil into the funnel, with the side walls of the foil folding upwards in the process. Through slits on the four sides of the aluminium block, small side levers push the top covers of the foil down, closing them on top of the tile like an envelope. The order of the side levers' movement matters here: First, the square end of the foil with the double-sided tape on the back is pushed down; second, the three triangular ends are folded over; last, the side levers are retracted and the top lever pushes down once more, pressing the top covers to the tape and holding

¹size: 11 inch \times 11 inch




Figure 4.15: Examples of tile wrapping. (a) shows a piece of ESR foil cut into the shape needed to wrap a centre tile. In the middle of the foil, a square window is cut out for the SiPM. The red square on the right end of the foil is a piece of adhesive tape used to hold the foil together once it is folded around the tile. In (b), two wrapped tiles are shown, a centre tile (on the left) and an edge tile (on the right). The bottom side with the SiPM window is facing up/front in this picture for better visibility.

them in place. Lastly, the wrapped tile is ejected by a bottom lever. Two wrapped tiles – one centre tile, one edge tile – are shown in Figure 4.15b.

The core part of the tools, the central aluminium block, was first developed using CAD and a prototype was 3D-printed to test the feasibility of the design. In this prototype, all mechanical processes, like pushing down the tiles and moving the levers to fold the foils, were purely manual. This, while proving that the concept was working, was considered not feasible for mass production in terms of reproducibility and working time per tile; therefore, the semi-automatic version of the tools described above and shown in Figure 4.16 was developed in cooperation with the mechanical and electronics workshops of KIP. Small servo motors drive the movement of the top, bottom, and side levers. The motion-sequence is controlled by an Arduino and can be easily adapted and tuned, e.g. in terms of motor speed, time of remaining in a certain position, or breaks in the sequence which demand a physical input to continue. Different sequences can be stored on the Arduino and selected directly from a menu on the tool's display. The selected sequence can be started by pressing the yellow button on top of the tool (see Figure 4.16a).

Using the semi-automatic tools, only two manual steps remain in the wrapping procedure: placing the foil and tile on top of the aluminium block at the start of the sequence, and taking the wrapped tile out at the end. The full wrapping process is thus greatly streamlined, allowing the user to wrap tiles for several hours in one sitting. The time needed per tile is reduced by a factor ≈ 10 – 20 compared to the purely manual prototype, resulting in an expected production time for the full amount of tiles of ≈ 2 – 3 months, i.e. on a similar scale as the scintillator production time. While the tile wrapping

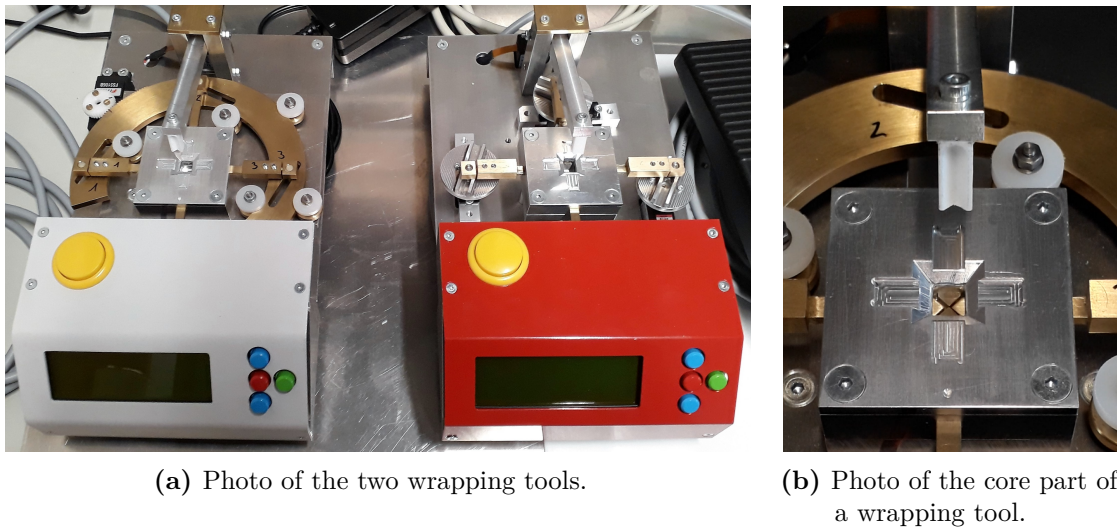



Figure 4.16: Tile wrapping tools. In (a), the two semi-automatic tools are shown: the white tool on the left for the centre tiles, the red tool on the right for the edge tiles. The key difference between the two lies in the core part of the tool, which is the aluminium block in the centre. (b) shows a zoom of such an aluminium block, here for the centre tiles. The shapes of the funnel and of the channels on the top match the shapes of the centre tile and its foil, respectively. Correspondingly, the aluminium block of the edge tile tool is cut slightly differently to exactly match the shape of the edge tile and its foil.

using the semi-automatic tools is thus still a time-consuming procedure, this estimate demonstrates that a purely manual process would not have been feasible for the full Tile Detector production. Furthermore, the automation of the most relevant steps of the wrapping procedure yields more consistent results than the previous process, ensuring a high quality standard of the wrapped tiles.

Tile gluing

The gluing of the tiles to the SiPMs is done matrix-wise, i.e. 16 tiles in a 4×4 grid are glued at the same time. The glue ² is transparent (in order to let the scintillating light travel from tile to SiPM) and highly viscous, with a curing time of several hours to a day (depending on temperature and humidity). This means that the tiles need to be held in place until the glue has cured. However, as the glue is viscous and the foils are somewhat stiff, the tiles cannot simply be placed on top of the SiPMs and be left there until the glue has dried, but need to be actively pushed together and down onto the SiPM surface to stay in position. Otherwise, the SiPM windows in the reflective

²DOWSIL 3145 RTV MIL-A-46146



(a) Photo of a disassembled gluing tool.

(b) Photo of an assembled gluing tool.

Figure 4.17: Photo of a gluing tool, in (a) disassembled and (b) assembled state. The core brass block holds the tiles in place, while the pedestal is used to mount the SiPM matrix. Both brass block and pedestal are assembled on the base plate using alignment pins and screws.

foils might be misaligned with respect to the SiPM surfaces, leading to a loss of light. Furthermore, a strong misalignment of the tiles could cause mechanical problems and the risk of breaking tiles off, as the gap sizes discussed in Section 4.1.1 would then not offer enough margin for a safe matrix assembly.

These requirements have led to the development of tools which allow the gluing of the tiles in the final 4×4 matrix form as a whole to the SiPM matrix. Similar to the wrapping tools, the development was done in CAD software and started with 3D-printed prototypes. Once the proof of concept was established, the final tools were produced at the KIP mechanical workshop for high precision and reproducibility. One of the final tools is shown in Figure 4.17, both in disassembled and assembled state. Broadly speaking, the tool can be separated into three main parts: the core brass block which holds the tiles in place, seen on the right side of Figure 4.17a; the pedestal to which the SiPM matrix is screwed, in the front in Figure 4.17a; and the base plate, on which the core block and the pedestal are assembled using both screws and alignment pins. The core block in turn is made up of several components: a funnel cut into the brass at an angle matching that of the edge tiles; a removable front wall that can be slid into place along two grooves; a movable back wall that can be tightened using a screw, thus allowing the fixation of the 16 tiles in the funnel; and a stamp shaped to fit into the funnel in order to push the tiles down onto the SiPMs, which can be tightened with screws as well.

In order to ensure the correct positioning of the tiles with respect to the SiPMs during the gluing, an alignment procedure was developed using acrylic glass mock-ups of the SiPM matrix and a 4×4 tile array. Circles etched into the mock-ups mark the centres of the SiPMs and of the tile bottom surfaces. By mounting the mock-ups on the pedestal and the core brass block, and subsequently aligning the centre marks under a microscope, the pedestal and core block can be correctly positioned relative to each other. They are then fastened on the base plate using the screws, and the precision holes for the alignment pins are drilled through the pedestal and brass block into the plate to maintain

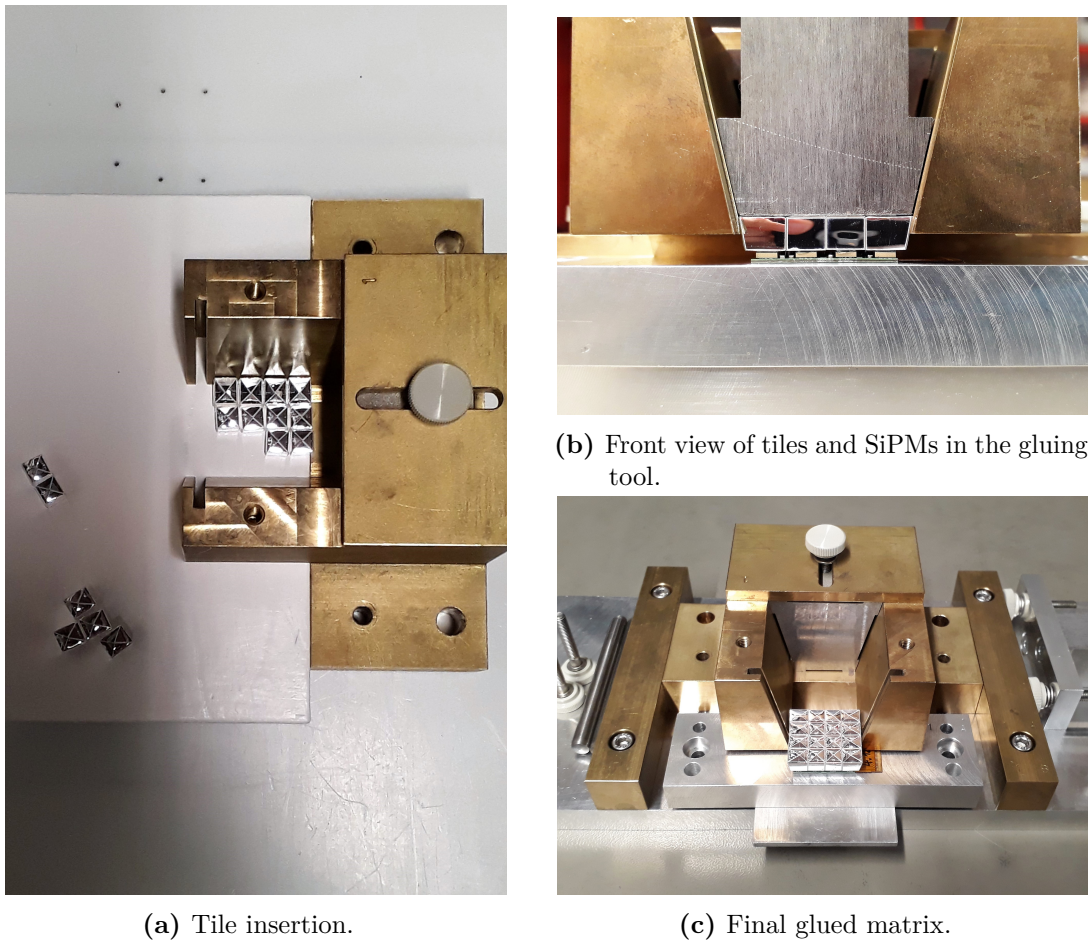


Figure 4.18: Gluing procedure using the custom-made gluing tool. **(a)** In the first step, 16 tiles are arranged in the 4×4 matrix layout in the core block of the gluing tool. **(b)** After dispensing the glue onto the SiPMs, which are mounted on the pedestal, the tiles are pushed down onto the SiPM surface using the gluing tool stamp. **(c)** After a curing time of ≈ 6 h, the tool is disassembled again and the finished matrix is dismounted.

the correct positioning. Once the alignment process is completed, the tool is ready to be used for the tile gluing. The full gluing procedure can be summarised as follows:

- 16 tiles (8 edge tiles, 8 centre tiles) are arranged in the 4×4 matrix layout in the core block of the gluing tool, as shown in Figure 4.18a.
- The front wall is inserted and the back wall is moved to push the tiles together. The screw of the back wall is tightened such that the tiles are held in place, but can still be pushed down by the stamp in a later step.

- The SiPM matrix is mounted by inserting long screws through the pedestal from the backside to fasten the spacers.
- Both the pedestal with the SiPM matrix and the core block with the tiles are mounted on the base plate using the alignment pins.
- The tiles are pushed on top of the SiPMs using the stamp, as shown in in Figure 4.18b.
- The screws of the stamp and back wall are tightened, the positioning is finished.
- The core block with the tiles is removed again.
- The glue is applied to the SiPMs using a pressured air glue dispenser. The optimal amount of glue is found empirically: too little glue increases the risk of a detachment of the tiles from the SiPMs, while too much glue will ooze out and slightly displace the tiles. A timer in the glue dispenser is therefore used to ensure a consistent amount of glue.
- The core block is slid into position from the top down, using the alignment pins as guidance.
- Once the tiles are in place, the core block is screwed tight onto the base plate and left there for approximately six hours to cure.
- After the curing time is over, the tool is dismantled again and the finished matrix is dismantled, shown in Figure 4.18c.

The time needed for the gluing of one matrix is approximately 30 minutes. In total, four tools were produced by the KIP mechanical workshop, yielding a workload of about two hours for four matrices per day (not including the QC measurements, which are described in the next chapter). Given the curing time of ≈ 6 h, it is therefore possible to glue up to eight matrices per day, resulting in ≈ 2 –3 months of matrix production for the full detector.

4.4 Module assembly and evaluation

Beside the matrices, several components are still needed in order to produce a full Tile Detector module: 13 MuTRiGs, which need to be mounted on a fully assembled TMB, and an aluminium support plate with a cooling pipe embedded.

The MuTRiG design is maintained in the Tile Detector group at KIP (see also Section 3.3.1 and [58]); the production and packaging into a standard ball grid array (BGA) package are then performed by external companies. In the case of the MuTRiGs of the two preproduction modules, the packaged ASICs were directly mounted on the TMB and tested there. For the full detector, a test setup is foreseen which allows the qualification of the packaged chips before assembly, by using a spring pin socket which




Figure 4.19: Fully assembled preproduction module. **(a)** shows the module lying on the TMB side with the tile matrices up. Each matrix is covered with a QR code sticker, linking it to its individual database entry. The multi-coloured cables in the front of the picture are the LV connections of the TMB. **(b)** shows the same module flipped onto the tile matrix side with the TMB up. The 13 MuTRiGs in their BGA packages are assembled in the centre of the TMB, with the flex-prints of the matrices plugged into the connectors on the TMB edges.

establishes electrical connection to the chip without the need of soldering. With the test setup in place, only those MuTRiGs that pass this QC step will be assembled to the TMBs for the final Tile Detector. The TMB assembly is done in-house at the KIP electronics workshop. For the two TMBs of the preproduction modules, basic functionality tests using the full DAQ chain with a Tile Detector specific detector adapter board (DAB), a FEB and a read-out PC were performed. For the full detector, more elaborate QC tests are foreseen, which are based on the insights gained from the two preproduction modules and are currently under development.

The design of the cooling circuit was continuously adapted in CAD over the course of this thesis, in order to comply with the space requirements of the surrounding detector systems and services (as discussed in Section 4.1.1). During this time, several prototypes were both 3D-printed and machined in-house. The experience with the machining of the support plates at the KIP mechanical workshop led to the final plates also being produced in-house. The bending and gluing of the stainless steel pipes into the cooling

plate is done in-house as well. Once all the individual components are ready, the full module assembly is done by the Mu3e group at KIP in the following order:


1. The matrices are mounted on the top side of the support plate and secured by screws from the backside, as described in Section 4.1.1.
2. The TMB is screwed to the bottom side of the support plate, thus covering the cooling pipe and the access to the matrices.
3. The matrix flex-prints are inserted and locked into the connectors on the TMB.

A fully assembled module is shown in Figure 4.19. The modules will undergo QC tests as well, which are also still under development as discussed in Chapter 5. The cooling circuit is discussed in more detail in Chapter 6.

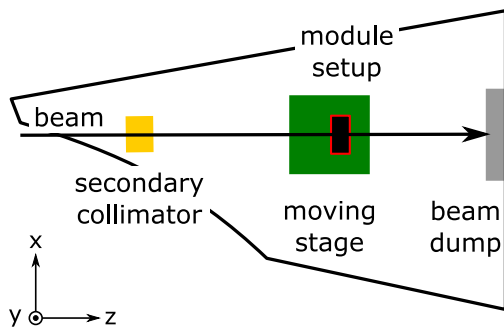
The assembly of the two preproduction modules was the first full module production of the Tile Detector, and thus forms an important milestone in the Tile Detector development. The modules were commissioned at the DESY testbeam facility [66]. The focus of the commissioning lay mainly on the investigation of meaningful QC procedures and in-situ calibrations suited for the Tile Detector operation in the full experiment. Module-wide synchronisation of the MuTRiG ASICs was investigated in hardware tests of the clock and reset signals in [68], while a first calibration scheme including the settings of the MuTRiG time and energy thresholds was developed in [69]. In addition, a first online monitoring interface with an integrated alarm system was developed and tested during this testbeam, as described in [70].

The layout of the testbeam facility at the DESY II synchrotron is shown in Figure 4.20a, with a more detailed sketch of the testbeam area (area TB22 in case of this testbeam) in Figure 4.20b. In the DESY II ring, a bunch of $\approx 10^{10}$ electrons or positrons is accelerated from an injection energy of $E_{\min} = 0.45$ GeV to typically $E_{\max} = 6.3$ GeV [66]. By placing a carbon fibre target in the circulating beam, bremsstrahlung photons are generated, which are converted to electrons and positrons via pair production in a secondary target made out of a metal plate. A dipole magnet is used for the selection of particle charge and momentum, before the beam is collimated and guided into the areas accessible by the testbeam users. Both the dipole magnet settings and the secondary target material and thickness are remote-controllable and can be selected by the users, depending on the need of particle type, momentum and rate. For the testbeam campaign discussed here, a 2.2 GeV e^- beam with a peak electron rate of ≈ 30 kHz was used. In the testbeam areas, secondary collimators with interchangeable lead insets are installed. The lead insets are available in various collimator sizes and shapes and can be manually exchanged by the testbeam users. In the testbeam discussed here, a collimator size of $18 \text{ mm} \times 18 \text{ mm}$ is selected to cover most of the surface area of a matrix without losing too much beam intensity.


The setup used in the testbeam campaign is placed on a remote-controllable moving stage in the centre of the area between collimator and beam dump, as sketched in Figure 4.20b. It is based on the geometry of the Mu3e experiment: the preproduction



(a) Layout of the DESY testbeam facility [71].



(b) Sketch of the testbeam area used for the measurement campaign.



(c) Photo of the setup used for the measurement campaign.

Figure 4.20: Commissioning setup of the preproduction modules at the DESY testbeam facility. (a) gives an overview of the testbeam facility at the DESY II synchrotron, with a sketch of the beam generation via bremsstrahlung and electron/positron conversion. In (b), the layout of the user area of the testbeam campaign (TB22) is sketched, with a corresponding photograph of the module setup shown in (c).

modules are assembled on two 3D-printed Tile Detector endrings, which in turn are mounted on a mock-up beam pipe with the correct dimensions of the real Mu3e beam pipe, thus allowing the validation of the installation procedure. In order to ensure the light reduction needed for the SiPM operation, the setup is covered by a 9.6 mm thick PVC pipe with an outer diameter of 250 mm. The pipe is assembled on two 3D-printed flange plates, which are mounted on the mock-up beam pipe as well. In addition, the flanges provide feedthroughs for all services, including liquid cooling, signal and power lines, and environmental sensor connections. The whole system is fixed on an optical




Figure 4.21: Commissioning setup of the preproduction modules at the DESY testbeam facility, with the PVC pipe removed. One of the preproduction modules is assembled on two 3D-printed Tile Detector endrings, which in turn are mounted on a mock-up beam. 3D-printed mock-up support plates (in white, opposite of the module) are used to stabilise the setup. Two large flange plates, one shown on the right side of the picture, provide feedthroughs for all services such as liquid cooling (at the top of the flange plate) or LV lines (multi-coloured cables at the bottom of the flange). In the shown configuration, the module is hit perpendicularly by a 2.2 GeV electron beam.

breadboard via 3D-printed mounts for the mock-up beam pipe, which is then placed on the moving stage in the testbeam area as shown in Figure 4.20c. The supplies for the LV and HV power, as well as the read-out PCs, are hosted in a service rack and are connected via patch panels to the user huts for remote control and monitoring.

Figure 4.21 shows the setup with the PVC pipe removed so that the module and the flange with the service feedthroughs are visible. In the configuration shown in this picture, the tile matrices are positioned perpendicularly to the beam, as indicated by the beam arrow. This is defined as $\theta = 90^\circ$, where θ is the angle between the beam axis and the matrix surface on the breadboard plane. The breadboard can also be manually rotated so that the matrices are positioned along the beam axis and the beam traverses the full module ($\theta = 0^\circ$). Furthermore, the mounts of the mock-up beam pipe are designed to allow manual rotations around the pipe/ φ axis. Lastly, the remote-controlled moving stage allows the positioning in x and y even while the beam is on, which makes time-effective scanning of the modules for calibration possible.

A first result of the commissioning can be seen in Figure 4.22. A hit map of one module is shown as a projection in the local z - φ plane after applying the new calibration




Figure 4.22: Hit map of a preproduction module after calibration during the testbeam campaign at DESY. The sketch on the right side shows the beam orientation with respect to the module: the beam is perpendicular to the module, $\theta = 90^\circ$, and the two matrix rows are hit under the same φ angle. The beam centered on MuTRiG #6 at $(n_{\text{tile},z}, n_{\text{tile},\varphi}) = (26, 4)$ is visible. Several errors, causing for example the data loss of MuTRiG #11 at $n_{\text{tile},z} = [4, 8]$ and the repeating pattern of missing/reduced data at $n_{\text{tile},\varphi} = 3$ (channel 0 of each MuTRiG), have been identified and corrected.

scheme. Each bin in the hit map corresponds to one channel, i.e. to one tile-SiPM detector element. The orientation of the module corresponds to the one shown in Figure 4.21, that is perpendicular to the beam ($\theta = 90^\circ$) with both rows of matrices hit under the same angle (illustrated in the sketch on the right side of Figure 4.22). Several observations can be noted:

- The vast majority of the 416 channels could be calibrated to a meaningful working point, so that the beam, which was centered around MuTRiG #6 at $(n_{\text{tile},z}, n_{\text{tile},\varphi}) = (26, 4)$, is clearly visible.
- The missing data marked by the red rectangle (corresponding to MuTRiG #11) and the repeated pattern of missing/reduced data highlighted by the black rectangles at $n_{\text{tile},\varphi} = 3$ (corresponding to channel 0 of each MuTRiG) are both caused by connection problems, which have been corrected in the meantime.
- The increased amount of entries at high $n_{\text{tile},z}$ was visible during the whole testbeam, independent of the beam position. It was most likely caused by limitations in the read-out chain, as there are no sorting and event building algorithms implemented yet in the first version of the FPGA firmware.

It is notable that several channels take no or only little data in Figure 4.22 without a known cause (i.e. which are not marked either in red or in black). In order to identify whether this is due to an insufficient exposure by the beam or whether these channels indeed exhibit a problematic behaviour, Figure 4.23 shows the same module rotated to $\theta = 0^\circ$, such that the beam is now directed *along* the module. In this way, the uniformity of the beam exposure is greatly increased (disregarding the first two ASICs, which show an increased data income due to the above mentioned issues in the read-out chain). All




Figure 4.23: Hit map of a preproduction module after calibration during the testbeam campaign at DESY. The sketch on the right side shows the beam orientation with respect to the module, which is rotated to $\theta = 0^\circ$ so that the beam is directed along the entire length of the module. Again, the repeating pattern of the missing/reduced data of the channels 0 (at $n_{\text{tile},\varphi} = 3$) has been marked by black rectangles. In addition, all anomalous channels which could not be attributed to any known error or bug are highlighted in red. Counting only these channels, and disregarding the channels 0 and the missing MuTRiG #11 at $n_{\text{tile},z} = [4, 8]$, results in a fraction of 3.6% of channels which could not be tuned to a meaningful working point.

channels that show an anomalous behaviour, and are not covered by the faulty channel 0 connection or MuTRiG #11, are marked using red rectangles; for easier distinction, the faulty channels 0 are again highlighted in black.

Counting all anomalous channels in Figure 4.23 results in 59 out of 416 channels not working as expected, resulting in a yield of 85.8% of fully functioning channels. However, when detracting the channels with a known (and fixable) error or bug, only 15 out of 416 channels, or 3.6%, remain which could not be calibrated to a meaningful working point. In conclusion, the first version of the calibration scheme can already be expected to successfully tune more than 90% of the channels. Besides the further advancement and optimisation of the calibration algorithms, an improvement of the yield is also expected through a new iteration of the MuTRiG ASIC. This version of the chip features an increased range of the threshold settings and improved threshold control, allowing for a better tuning of the individual channels. At the writing of this thesis, the new ASICs are being tested in a dedicated QC system (see Section 5.4), which will be used to evaluate the chips before the TMB assembly. Thus, considering the error corrections already implemented to fix the channel 0 and MuTRiG #11 issues, the new ASIC version, and the QC schemes for both the MuTRiGs and the TMBs, which will be in place for the full production, a yield of close to 100% can be expected for the final modules of the Tile Detector.

Chapter 5

Development and evaluation of quality control procedures

62 matrices were assembled in total for the two preproduction modules. Based on these matrices, both the full production line described in Chapter 4 and the quality control (QC) concepts, which will be discussed in this chapter, were developed. The purpose of the QC is to evaluate all detector components after their production or assembly, in order to ensure a high detector performance. The workflow of the detector production and QC procedures is sketched in Figure 5.1. Following the order of production in Chapter 4, first the QC of the scintillating tiles is described in Section 5.1, followed by the development of the SiPM matrix QC in Section 5.2. In Section 5.3, the design of a QC system for the tile matrices is discussed, with an outlook given in Section 5.3.3 on a common matrix QC system covering both the SiPM and tile QC steps. Lastly, QC concepts for the MuTRiG, TMB, cooling circuit and module are introduced in Section 5.4.

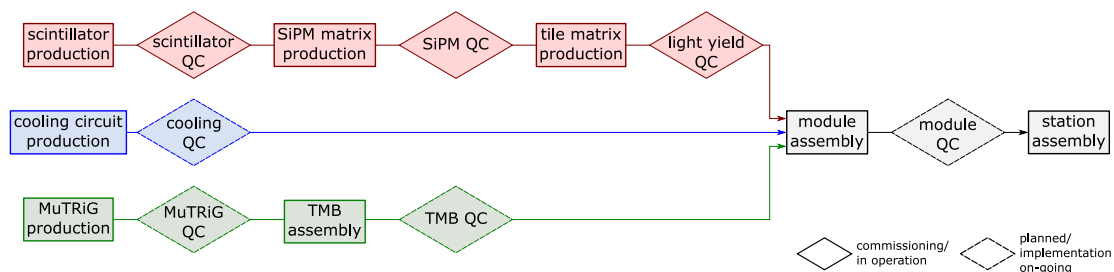


Figure 5.1: Sketch of the production and QC workflow. The manufacturing phase of the Tile Detector can be separated into three independent chains: the tile matrix production in red, the cooling circuit production in blue, and the read-out electronics production in green. Once the production in all chains is completed, the modules, and ultimately the two stations, can be assembled. Each production/assembly step (in a rectangular frame) is followed by a dedicated QC procedure (in a diamond frame). The QC systems in the frames with continuous lines are either currently being commissioned or already in operation, while the QC procedures in the dashed frames are still in the planning or implementation phase.

5.1 Quality control of the scintillating tiles

In the first step of the tile matrix production line in Figure 5.1, the scintillating tiles are manufactured. The QC concept for the tiles foresees two basic steps: first, a visual inspection of the scintillator cubes, and second, the measurement of height, length and (in the case of the centre tiles) width of the tiles. Here, the visual examination is limited to obvious visible defects, such as scratches, dents or cracks; a more detailed study of the scintillator quality is later performed during the tile matrix QC (Section 5.3). The measurement of the tile size is used to verify whether the scintillator cubes are within the defined dimensions, given by the following specifications:

$$h_{\text{tile}} = 5.00_{-0.05}^{+0.00} \text{ mm}, \quad l_{z,\text{tile}} = 6.21_{-0.02}^{+0.00} \text{ mm}, \quad l_{\varphi,\text{tile}} = 6.30_{-0.02}^{+0.00} \text{ mm} \quad (5.1)$$

Given the space constraints discussed in Section 4.1.1, the scintillator cubes should not be larger than the specified value in any dimension; therefore, the sizes in the specifications are given as an upper limit. On the other end, if the cubes are too small in length and width ($l_{z,\text{tile}}$ and $l_{\varphi,\text{tile}}$), the 4×4 tile matrix might be pushed together too much during the gluing step, thus misaligning the tiles relative to the SiPMs; therefore, a maximum deviation of $-20 \mu\text{m}$ is required for length and width. The specifications are more relaxed for the tile height, allowing a maximum deviation of $-50 \mu\text{m}$, since this has no direct mechanical implications on the tile matrix. Detailed dimension checks were performed for the first production tests, using extensive measurements provided by the KIP mechanical workshop. Based on these checks, the machining process was fine-tuned and optimised for mass production, as will be described in the following paragraphs.

Once the new production method using the ice-vice concept had been verified as discussed in Section 4.2.2, the machining of the tiles was expanded to a full scintillator plate yielding 4×32 tiles (equaling eight matrices). While it was no problem to machine the full scintillator strip on the vacuum plate, the subsequent ice-vice step proved to be challenging: the scintillator strip was bending due to the water expanding during the freezing process, with the detachment of the strip ends from the ice-vice plate as a consequence. Adding weights on top of the strip was tried as a means of keeping the strip attached, but did not significantly improve the situation. As this made any further machining impossible, an alternative procedure was tested: the full strip was machined on the vacuum plate, but divided into “matrix pieces”, i.e. pieces of 4×4 tiles, for the ice-vice step. Figures 5.2 and 5.3 show the results of the size measurements of one exemplary 4×4 scintillator matrix. In both figures, the deviation from the norm defined in Equation 5.1 is illustrated. The length $l_{z,\text{tile}}$ and width $l_{\varphi,\text{tile}}$ (only for the centre tiles), shown in Figure 5.2, are machined very precisely and deviate only in the range of $[-4 \mu\text{m}, 3 \mu\text{m}]$ from the specified values. Since it was not expected that the machining would be this precise, this was considered to be sufficiently accurate for the matrix production.

The deviation in the height of the scintillator tiles, on the other hand, is not as uniform as the length and width deviations, as illustrated in Figure 5.3. For each cube, the





Figure 5.2: Measured deviations of the specified length $l_{z,tile}$ and width $l_{\phi,tile}$ of 16 scintillator tiles, produced as a 4×4 matrix. Due to the angle of the edge tiles, the width can only be measured for the centre tiles. The deviations from the specifications defined in Equation 5.1 are of the order of 1‰, ranging from $-4 \mu\text{m}$ to $3 \mu\text{m}$.

height was measured in four quadrants. This was necessary because, as Figure 5.3a demonstrates, there are large height differences within the tiles at the border of the 4×4 matrix. In particular the edge tiles at the four corners, corresponding to the tile numbers 1, 4, 13 and 16 in Figure 5.3b, show a large non-uniformity of the height. The tile with the largest internal difference is the edge tile in the bottom-right corner (tile number 4 in Figure 5.3b), with the deviations of ranging from $-31 \mu\text{m}$ to $8 \mu\text{m}$. Only the four tiles in the centre of the matrix (tile numbers 6, 7, 10 and 11 in Figure 5.3b) are quite uniform in height. This behaviour can be explained by the displacement of the tiles by the freezing water: despite having weights on top, the edges of the matrix piece are pushed out- and upwards by the ice. When they are milled flat in the next step, this results in a top surface with a slope. Overall, deviations ranging from $-33 \mu\text{m}$ to $11 \mu\text{m}$ are measured over the full 4×4 matrix.

In order to investigate the extent of this non-uniformity in the tile height, a strip of 4×28 tiles was investigated in the next step. However, for this larger data set, the height is measured once per tile, instead of four times for each quadrant of the tile. The resulting deviations from the specifications are shown as a colour map in Figure 5.4a and projected to a one-dimensional graph in Figure 5.4b. The dashed red lines in both figures indicate where the strip was separated into smaller pieces for the ice-vice machining. In general, the uniformity of most 4×4 pieces is better than in the example shown in Figure 5.3. Over the full strip, the deviations range from $-34 \mu\text{m}$ to $14 \mu\text{m}$, which is similar to the range quoted for the single 4×4 matrix piece above. However, it should



(a) Height deviations, colour map.


(b) Height deviations, 1D graph.

Figure 5.3: Measured deviations of the tile height of 16 scintillator tiles, produced as a 4×4 matrix. The height is measured in four quadrants for each tile. (a) shows the resulting deviations from the specifications defined in Equation 5.1 as a 2D colour map. It is clearly visible that the border tiles have large height differences within the individual tile, while the four tiles in the centre are quite uniform. For better visibility of the differences, the same data points are shown in (b) as a 1D graph, with the average deviation and the spread shown in grey as a guide to the eye. The height deviations range from $-33 \mu\text{m}$ to $11 \mu\text{m}$ over the full 4×4 matrix.


be noted that for the strip, the height of the tiles is measured only in the centre, and that the deviations would be expected to be larger if the corners of the tiles had been explicitly measured.

In general, the height deviations are still in a range considered to be acceptable for the Tile Detector production. However, as a more uniform height of the tiles would be better both in terms of mechanical assembly and in terms of expected light yield, another step of improvement was tried: the strip is divided into pieces of 2×2 tiles instead of 4×4 , as the weight put on top during the freezing step can be distributed better this way. The results for 16 tiles, produced as four 2×2 pieces, are shown in Figures 5.5 and 5.6.

The deviations from the specified length and width of the 2×2 pieces, shown in Figure 5.5 are very similar to those of the single 4×4 piece shown in Figure 5.2. Likewise, they are considered satisfactory for the tile production, and in addition provide a cross-check that the machinery is still correctly set up. A significant improvement of the height uniformity is achieved for the 2×2 pieces compared to the 4×4 matrix, as Figure 5.6 demonstrates. The displacement of the edges by the ice is still partially visible in Figure 5.6a, but not nearly as prominent as in Figure 5.3a. The projection of the data in Figure 5.6b also shows that the height of the tiles is now much more uniform compared to Figure 5.3b, with the range of the deviation significantly reduced from $[-33 \mu\text{m}, 11 \mu\text{m}]$ to $[-14 \mu\text{m}, 5 \mu\text{m}]$. Given this clear improvement in the tile production and the fact that separating the strip into 2×2 pieces does not significantly increase the



(a) Height deviations, colour map.



(b) Height deviations, 1D graph.

Figure 5.4: Measured deviations of the tile height of 4×28 scintillator tiles, produced as 4×4 matrix pieces. The height is measured once in the centre of each tile. In (a), the resulting deviations from the norm (see Equation 5.1) is shown as a colour map. The horizontal red lines mark the edge and centre tile rows, while the vertical dashed lines indicate the separation into 4×4 pieces. The height uniformity is generally better than in Figure 5.3, except for the third matrix from the left. This is also clearly visualised by the dip at column number 9 in (b), which shows the same data projected to a 1D graph. Overall, the deviations range from $-34 \mu\text{m}$ to $14 \mu\text{m}$.

working time, this method was chosen for the full scintillator production. Following this extensive fine-tuning of the manufacturing process, the monitoring of the tile dimensions was only performed on random samples during the tile production to ensure the correct settings of the machinery.





Figure 5.5: Measured deviations of the specified length $l_{z,\text{tile}}$ and width $l_{\phi,\text{tile}}$ of 16 scintillator tiles, produced as four 2×2 matrix pieces. The results are very similar to those presented in Figure 5.2, and are equally considered to be sufficiently accurate for the tile production.



(a) Height deviations, colour map.

(b) Height deviations, 1D graph.

Figure 5.6: Measured deviations of the tile height of 16 scintillator tiles, produced as four 2×2 matrix pieces (illustrated by the black dotted lines). The height is measured in four quadrants for each tile. A clear improvement in the uniformity is seen compared to Figure 5.3, both in the 2D colour map in (a) and in the projection of the data in (b). The height deviations now only range from $-14 \mu\text{m}$ to $5 \mu\text{m}$ over the four 2×2 pieces, as opposed to the range of $-33 \mu\text{m}$ to $11 \mu\text{m}$ of the full 4×4 matrix piece.

5.2 The SiPM matrix quality control system

Following the workflow in Figure 5.1, the next step in the matrix production line is the assembly and QC of the SiPM matrix. The purpose of the SiPM QC system is to validate the functionality of the SiPM matrices and to approve them for the gluing step. The setup presented in this section is the first version of the SiPM QC, and is designed to verify the basic functionality of the SiPM matrices by testing

- electrical connections,
- SiPM powering and data taking, in the form of noise measurements,
- and temperature sensor read-out.

To this end, the same components/PCBs of the MuTRiG stand-alone system, as previously shown in Figure 4.11, are used. However, the setup for the SiPM QC is more compact, as illustrated in Figure 5.7. This makes it possible to put the system in a climate chamber for both a controllable temperature environment and light shielding needed for the SiPM noise rate measurement. An improved version of the setup, which provides a broader range of SiPM characterisation features, will be presented in Section 5.3.3.

Since the noise rate is dependent on temperature, all SiPM measurements presented in the next paragraphs have been executed at a fixed temperature of 10 °C. Furthermore, the SiPMs are operated at a voltage of $V_{op} = 55$ V, corresponding to approximately 3 V





Figure 5.7: Sketch and photo of the first version of the SiPM QC measurement setup. The same components of the MuTRiG stand-alone setup as described in Section 4.2.2 are used, but assembled in a more compact way in order to fit in a climate chamber for temperature control and light shielding.

overvoltage as recommended in the data sheet of the manufacturer (see Table 3.2 and [56]). As a first step, the matrix is connected to the QC setup, with only the temperature sensor powered and read out. If the temperature output is as expected, the next measurements are performed; if not, the matrix is sent back to the electronics workshop for repairs. In the next step, the QC setup itself, including the MuTRiG, is powered. Any obvious problems with the electrical connections of the SiPMs such as shorts can be spotted in the MuTRiG current, which is protected by the current limit of the power supply.

If this check is successful as well, the last step of the SiPM QC is the measurement of the SiPM noise rate. The data are taken for each channel in a fixed sequence: First, the time threshold is identified where the electronic noise floor of the ASIC is reached. Then, the time threshold is set to two MuTRiG DAC values higher, to be safely out of the electronic noise. This timing threshold can be considered to be the lowest possible threshold at which the MuTRiG can be operated without the possibility of hitting the electronic noise floor. Lastly, the noise rate is determined for this threshold value by measuring the time difference Δt between consecutive hits in the SiPM. The resulting exponential distribution, shown for one exemplary channel in Figure 5.8a, can then be fitted with an exponential function of the form $f(t) = A \cdot e^{a-b \cdot t}$, which allows the



(a) Noise rate of channel 2 of matrix 30. (b) Noise rate distribution of all channels of matrix 30.

Figure 5.8: SiPM QC results from matrix ID = 30. (a) shows the distribution of the time difference between consecutive hits for one channel of matrix number 30. The distribution is fitted using an exponential function to extract the noise rate. The fit range is selected such that the starting point at 8 μ s is safely away from the electronic noise floor, and that the end point is the first bin with less than 100 entries to ensure enough data is still available for the fit to converge. This conservative limit is chosen as it works for all channels of all 62 matrices. In (b), the noise rate distribution of all 16 channels of matrix number 30 is shown. It is rather narrow and uniform, with an average value of 158.69 kHz and a standard deviation of 13.84 kHz, and without any large outliers.

extraction of the noise rate b . The SiPM is defined to pass the test if the measured noise rate is below 500 kHz. The reasoning of this sequence for the noise rate measurement is that the ASICs are expected to be operated at much higher timing thresholds due to a large number of photons in the tiles. Consequently, the SiPMs will have a much better noise rate than they have at this “worst possible” DAC setting. Therefore, they are expected to perform well during the Mu3e data taking if they already function normally at such a low timing threshold.

This process is repeated for all 16 channels of a matrix. As an example, the distribution of the noise of SiPM matrix number 30 is shown in Figure 5.8b. The average noise rate of the 16 channels is 158.7 kHz with a standard deviation of 13.8 kHz, a minimum value of 138.0 kHz and a maximum value of 182.8 kHz. It is evident that all channels are clearly below the 500 kHz limit set for the QC condition, and furthermore, that the spread of the distribution is low. Therefore, this matrix is considered to be well-performing and passes the QC test.

In order to compare the full amount of preproduction matrices, the measurements described above are performed for all 62 matrices. Figure 5.9 shows the resulting noise rate data as the matrix average value with a box representing the minimum and maximum values, in order to judge the matrix-internal spread of the channels. For reference, the matrix detailed in Figure 5.8 has the matrix ID 30 in Figure 5.9. From the overview


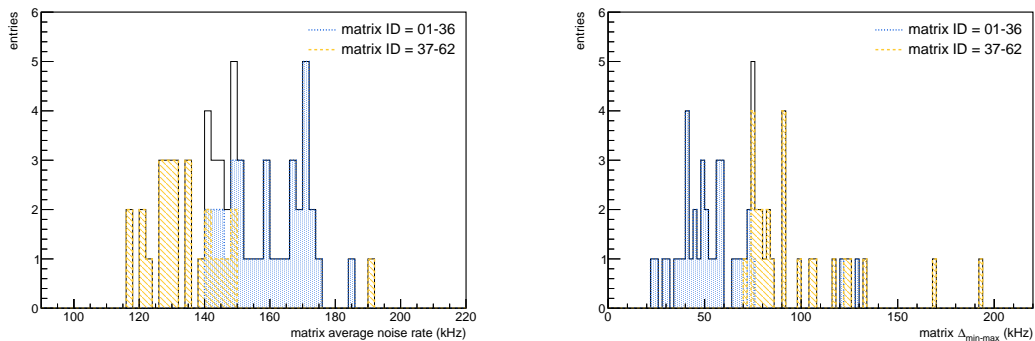


Figure 5.9: Summary of the QC results from all 62 SiPM matrices of the preproduction modules. The average noise rate per matrix (16 channels) is shown, with a box representing the minimum and maximum values of each matrix to visualize the spread. The dotted black line separates the matrices belonging to either of the two groups of SiPMs, as shown in Figure 4.13.

provided by this figure, several conclusions can be drawn:

1. All 62 matrices are clearly below the 500 kHz specification, meaning all have successfully passed this basic QC functionality test.
2. There seems to be a difference between the two groups of SiPMs (see Figure 4.13), which are separated by a black dotted line: the average noise rate of the first group is slightly higher than that of the second group, and the spread is smaller in the first group than in the second. This is further illustrated in Figure 5.10, which compares the distributions of matrix average noise rate (Figure 5.10a) and difference between matrix minimum and maximum noise rate (Figure 5.10b) of the two SiPM groups. The first group, with matrix IDs 01-36 (in dotted-blue), has a mean average noise rate of 159.3 kHz with a standard deviation of 11.6 kHz, which is higher than the mean of 135.2 kHz (standard deviation 14.4 kHz) of the second group comprised of IDs 37-62 in dashed-yellow. The behaviour of the spread, on the other hand, is flipped, with the first group having a much smaller average min-max spread of 53.9 kHz with a standard deviation of 21.6 kHz than the second group at an average min-max spread of 98.3 kHz with a standard deviation of 29.5 kHz. Whether these differences between the two groups stem from the SiPM production, the matrix assembly at the electronics workshop, or the QC measurement system, is not clear. However, it was notable already during the data taking that the QC setup would sometimes pick up electronic noise, possibly from the climate chamber,



(a) Average noise rate per matrix of the two SiPM groups.

(b) Difference between matrix minimum and maximum noise rate of the two SiPM groups.

Figure 5.10: Comparison of the matrices produced from the two groups of SiPMs. In (a), the matrix average noise rate from Figure 5.9 is shown as a histogram of all matrices in black, and separated into the two SiPM groups in blue dotted and yellow dashed lines. The mean average noise rate of the first group of matrices is higher than that of the second group. The opposite is true for the spread of the matrix noise rate $\Delta_{\min-\max}$, shown in (b), which is smaller for the first group than for the second.

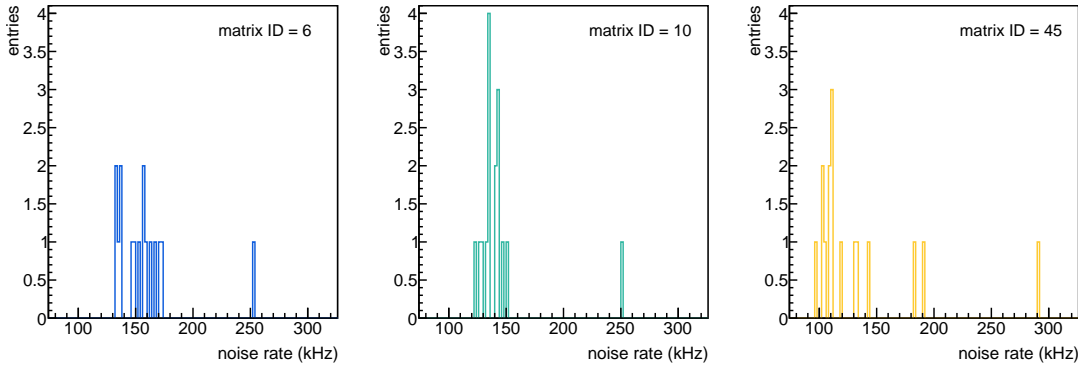


Figure 5.11: Comparison of the noise rate QC measurements from matrices 6, 10 and 45. All three matrices have a notably high maximum noise rate compared to a relatively small matrix average, revealed here to be caused by a single outlier. In all three examples, the outlier channel is a different one, excluding a problem in the ASIC or the electronics connection, and the noise rate spectrum of the outlier does not show any unusual features.

and that the protection of the setup from stray light was not always of the same quality. Even though the deviations between the two SiPM groups may therefore originate from the measurement process alone, they are still homogeneous enough to not repeat the QC procedure. Furthermore, the SiPM QC is at this stage simply a pass/fail process, with no subsequent use of the data in further QC or calibration measurements.

3. Some matrices, e.g. the matrices with IDs 6, 10 and 45 in Figure 5.9, have a relatively small average value compared to an unusually high maximum value. Looking at the noise rate distributions of the individual matrices reveals that this is due to only one outlier, as illustrated in Figure 5.11. It should be noted that the outlier channel is always a different channel for the three matrices, so a systematic error from the electronic channel or the flex print connection can be excluded. In all three cases, the outlier channel was investigated by checking the respective noise rate spectrum (as shown exemplary for matrix ID = 30 in Figure 5.8a); however, in none of the examples does the shape of the spectrum indicate that the electronic noise floor has been hit, or that the channel in question is in any way unusual otherwise. Another example is matrix ID = 50, which has a comparably large min-max spread. Further investigation yields a wide spread of all channels in general, not just two outliers at the minimum and maximum, as shown in Figure 5.12; however, when checking the individual noise rate spectra, all channels show typical behaviour.

The validation of the 62 matrices demonstrates that a general improvement of the setup in terms of light and noise shielding as well as the stability of operation would be beneficial for the reliability of the QC measurements. In addition, the examples of




Figure 5.12: Comparison of the noise rate QC measurements from matrix ID = 30 and matrix ID = 50. Compared to the distribution of matrix 30 in blue, the noise rates of matrix 50 in yellow are notably wider spread; however, none of the individual noise rate spectra show any signs of electronic noise or unusual behaviour.

irregular channels shown above illustrate that not all of the possible SiPM issues can be investigated with the presented QC setup. Therefore, all experience and insights gained from the QC of the preproduction matrices are invested in an improved version discussed in Section 5.3.3, which will be in place for the full Tile Detector production. It provides additional data taking and analysis features regarding SiPM properties such as breakdown voltage or gain, which allow a more detailed investigation in case of irregular or unusual behaviour of the SiPMs.

5.3 Development of a light yield-based matrix quality control system

The last production step of the tile matrix line in Figure 5.1 is the gluing of the scintillator tiles to the SiPMs. The corresponding QC system is designed to evaluate the performance of the tile matrices after the gluing procedure. Measuring the light yield of the tile-SiPM system is a very well suited method to investigate quality of the gluing and the uniformity of the detector response, which is crucial for later calibrations. The setup designed for the light yield QC measurements is, like the SiPM QC setup, matrix-based with an ASIC for the signal read-out. However, instead of MuTRiG, in this case a different ASIC called KLauS (Kanäle für die Ladungsauslese von Silizium-Photomultipliern) [72] is used. This

chip was also developed at KIP as a low power read-out ASIC for an imaging calorimeter at a future linear collider [73]. The KLauS ASIC design offers two *gain branches* for the charge measurement: the high gain branch, which can amplify single pixel signals, and the low gain branch designed for large signals. Each branch provides two scaling factors, which are selected by a switch, in order to cover a large signal range: high gain (HG) and medium gain (MG) for the high gain branch, and low gain (LG) and ultra-low gain (ULG) for the low gain branch. This makes the KLauS chip very well suited for the concept of the light yield QC, which can be summarised as follows:

1. Induce hits in the scintillator tiles using a laser or radioactive source with a known signature.
2. Map the energy known from the signal source to the digitised ASIC output.
3. Measure the relation between the number of photoelectrons in the SiPMs and the digitised ASIC output.
4. Calibrate the signal in digitised ASIC units to the number of photoelectrons.
5. Thus, map the energy deposited in the scintillator tile to the number of photons detected by the SiPM.
6. Normalise the result to 1 MeV to determine the final light yield for each channel.

The different gain factors of KLauS allow the measurement of the large signals of the source deposited in the scintillating tile, as well as the single-photon spectrum needed to obtain the gain of the SiPM in order to estimate the number of photons. The calibration steps needed to incorporate the ASIC into the light yield QC system are described in the following section, followed by a discussion of the QC setup and measurements in Section 5.3.2. The new common matrix QC system, which allows the SiPM QC and light yield QC steps to be performed using only one setup, is the subject of Section 5.3.3.


5.3.1 Calibration of the KLauS ASIC

Before the KLauS ASIC can be used in the QC process, it must be calibrated. Two calibration steps are described in the following paragraphs: the *gain calibration*, which maps the number of photoelectrons to the digitised ASIC output (in the form of analog-to-digital converter (ADC) bins); and the *interbranch calibration*, which allows the conversion between the different gain factors. The setup used for the calibration measurements is similar to the one shown in Figure 5.7, but with a KLauS ASIC instead of a MuTRiG chip. One SiPM matrix, corresponding to 16 channels in the KLauS ASIC, is used to demonstrate the calibration concept. Again, all measurements described below were performed in the climate chamber already used for the SiPM QC, but in this case close to room temperature at 25 °C.

Gain calibration

In the gain calibration step, the relation between the number of photoelectrons in the SiPM and the ASIC ADC bins is determined. Using a pulsed laser to directly shine light on the SiPMs, a single-photon spectrum (SPS) can be generated for each channel. The distance between the individual peaks of the spectrum yields the gain factor of the ASIC, in units of ADC bins/photon. Tuning the laser prior to the actual calibration data taking showed that using the MG configuration of KLauS, instead of the HG, is sufficient to produce an SPS with well separable peaks. As this saves an additional step in the interbranch calibration (described in the next subsection), the MG configuration was chosen for the data taking. The corresponding SPS measurement is shown for one exemplary channel in Figure 5.13a, with the electronic ASIC offset, the *pedestal*, already subtracted from the data.

The gain can then be extracted by performing a fast Fourier transform (FFT) of the SPS in Figure 5.13a. The resulting FFT spectrum is shown in Figure 5.13b, with the strongest frequency component of the spectrum corresponding to the distance between the individual SPS peaks. The peak position in the frequency domain is obtained using a Gauss fit and determining its mean μ , as indicated in Figure 5.13b. By dividing the full ADC range of 10 bit by the extracted mean value, the gain can be calculated; for the example of the channel in Figure 5.13, this results in a gain of $G \approx 5.26$ ADC bins/photon. The same procedure is performed for all channels of one matrix, resulting in a gain $G(i)$



(a) Single-photon spectrum.

(b) Fast Fourier transform.

Figure 5.13: Gain calibration of the KLauS ASIC. (a) shows a single-photon spectrum, which was generated by shining a pulsed laser directly onto a SiPM. The x axis is already corrected for the electronic offset (pedestal) of the ASIC. The distance between the individual peaks of the SPS yields the gain. (b) By performing an FFT of the spectrum and extracting the position of the highest peak using a Gauss fit, the gain can be calculated by dividing the full ADC range of 10 bit by the peak position.




Figure 5.14: Result of the gain calibration. 16 channels of the KLauS ASIC were calibrated using one SiPM matrix. The spread of the gain values between the channels is due to variations in both the ASIC channels and the SiPMs used in the measurement. The uncertainties on the determined gain values are on the order of $\%$ for most of the channels and therefore not visible in the graph. Only in the case of channels 5, 7 and 8, larger uncertainties are estimated, as the Gauss fits did not describe the data well and the bin with the maximum entries in the FFT spectrum was used instead.

for each channel i , as shown in Figure 5.14. The uncertainties on the gain values, which are propagated from the Gauss fits of the FFT spectra, are in the $\%$ range for most of the channels and too small to be visible in the figure. The exception are channels 5, 7 and 8, for which the Gauss fits did not work well; in these cases, the bin with the highest value in the FFT spectrum was used instead of the mean μ , with the bin width as an estimate for the uncertainty. All 16 SPS and FFT spectra are shown in Appendix B in Figures B.1 and B.2, respectively.

While the uncertainties on the individual gain values are small, a spread of the gain of more than 10% is observable between the different channels in the figure. This is caused by a combination of channel-by-channel variations in the KLauS ASIC and differences in the SiPMs used in the measurement. As the data taking time for a full matrix at a laser frequency of 100 kHz is only ≈ 40 s, the gain measurement is performed for each matrix individually in order to take the SiPM variations into account.

Interbranch calibration

As the amplification of the MG configuration was found to be sufficient for the SPS in the previous section, a calibration between the different branch gain factors is only necessary from the MG configuration to the ULG configuration needed for the energy measurement. In this interbranch calibration, the laser is used again to generate signals in the SiPMs. Using the same laser settings for two gain factors allows the measurement of the shift in the spectrum, from which the scaling factor between the gain factors can be determined. This is illustrated in Figure 5.15 for the MG and the LG settings. However, it is not possible to take data with the ULG and MG settings at the same time due to the limited ADC range and resolution of the ASIC. The calibration is therefore done in two steps: from ULG to LG, and from LG to MG. Afterwards, the results can be merged to get the final interbranch scaling factor B from ULG to MG.

In order to obtain the scaling factor between two gain settings, the automatic gain selection (GS) feature of KLauS can be exploited, as developed in [72]: it allows the ASIC to automatically select the most suitable gain setting for the sampling and digitisation of the analog signal in the ADC. If this feature is actively disabled, the ASIC will take data only with one selected gain configuration. However, the GS information is still accessible via a GS bit, in which the comparator status is encoded. As the comparator

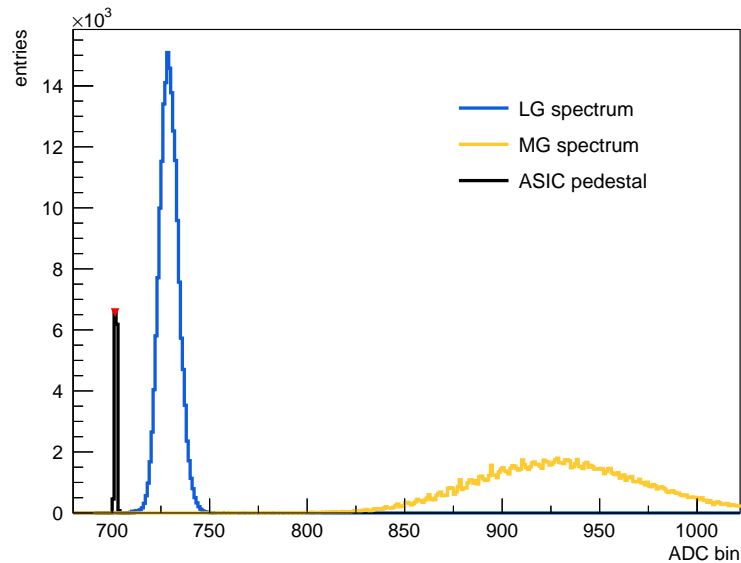
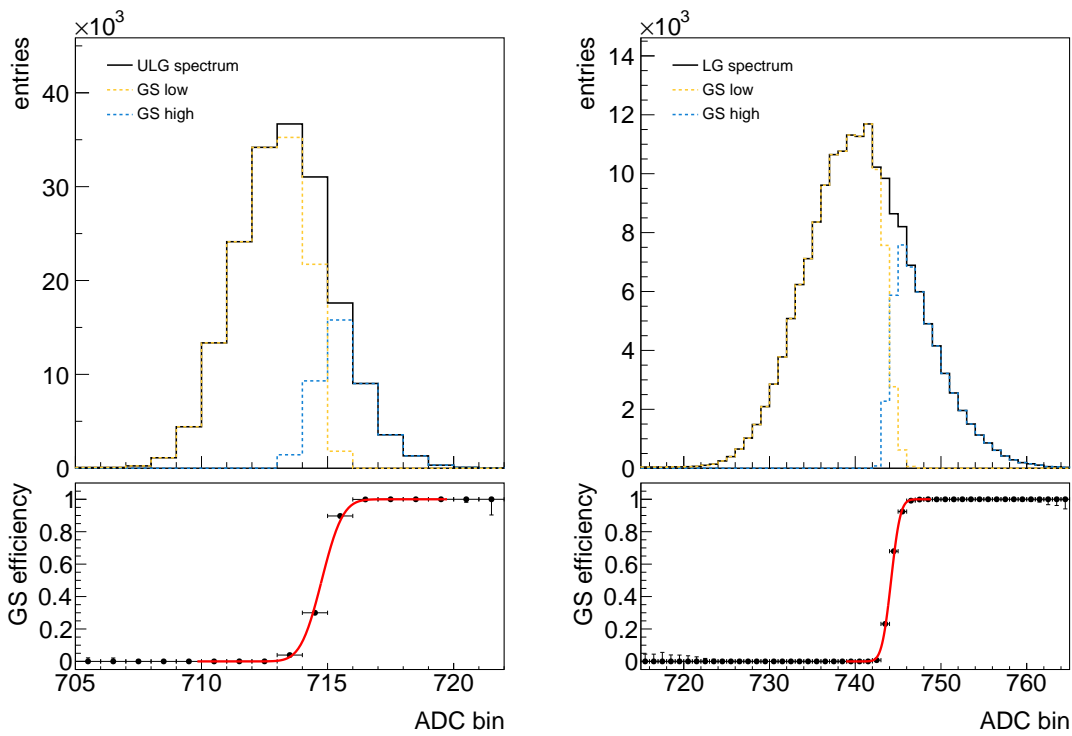


Figure 5.15: Interbranch calibration concept. The laser intensity is tuned such that the laser spectrum can still be resolved when using either the low gain or the medium gain configuration. This allows the extraction of the interbranch scaling factor $B_{\text{LG-MG}}$. The pedestal of the ASIC, which limits the available ADC range, is shown in black.

decision is based on the analog signal current, the GS is triggered at the same physical (charge) threshold independent of the enforced gain configuration. This is illustrated in the upper parts of Figure 5.16, which show two spectra that were taken under the same conditions with the automatic GS disabled: once with the ULG output enforced in Figure 5.16a, and once with the LG output enforced in Figure 5.16b. The dashed yellow (blue) sub-spectrum in each of the figures corresponds to a low (high) comparator status; at the point where the two sub-spectra cross, the GS comparator would be triggered. This is even clearer illustrated in the lower parts of the figure, which show the GS efficiency. Since the triggering point of the comparator is the same whether the ULG or the LG configuration is selected, the scaling factor between the two can be determined by fitting the efficiency curve and extracting the value at 50%.



(a) Ultra-low gain configuration.

(b) Low gain configuration.

Figure 5.16: Interbranch calibration of the ULG and LG configurations. In (a), the spectrum taken with the ULG setting enforced is shown. The dashed blue and yellow lines in the upper part of the figure illustrate the spectra of the automatic gain selection. The corresponding GS efficiency, fitted with Equation 5.2, is shown in the lower part. In (b), the equivalent graphs for the LG setting are presented.

A Gaussian error function of the form

$$f(x) = \frac{1}{2} \cdot (\text{Erf}(\sigma \cdot (x - \mu)) + 1),$$

$$\text{with } \text{Erf}(x) \equiv \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt$$
(5.2)

is used for the fit, with the parameter $\mu \equiv S_{\text{LG/ULG}}$ yielding the threshold position in question. The scaling factor $B_{\text{ULG-LG}}$ can then be calculated as

$$B_{\text{ULG-LG}} = \frac{S_{\text{LG}} - P_{\text{LG}}}{S_{\text{ULG}} - P_{\text{ULG}}},$$
(5.3)

with the threshold positions S_{LG} and S_{ULG} and the ASIC pedestal positions P_{LG} and P_{ULG} of the LG and ULG settings, respectively. The pedestals are obtained for each channel and for both gain configurations by taking a 10 s measurement with the laser turned off and the SiPM bias voltage $V_{\text{op}} = 40$ V set below the breakdown voltage. In the example of the channel shown in Figure 5.16, the triggering point of the gain selection is at $S_{\text{ULG}} \approx 714.8$ ADC bins for the ULG configuration and at $S_{\text{LG}} \approx 744.2$ ADC bins for the LG configuration. After subtracting the pedestal, $P_{\text{LG}} \approx P_{\text{ULG}} \approx 701.9$ ADC bins, the scaling factor is determined to be $B_{\text{ULG-LG}} \approx 3.3$. This procedure is then performed for all 16 channels, resulting in the graph in Figure 5.17a for the ULG-LG interbranch calibration. The individual pedestal measurements, GS spectra and GS efficiency curves of the 16 channels are shown in Appendix B, Section B.1.




Figure 5.17: Results of the interbranch calibration. The scaling factors of the 16 investigated channels are summarised in **(a)** for the conversion between the ULG and LG settings and in **(b)** for the conversion between the LG and MG settings.

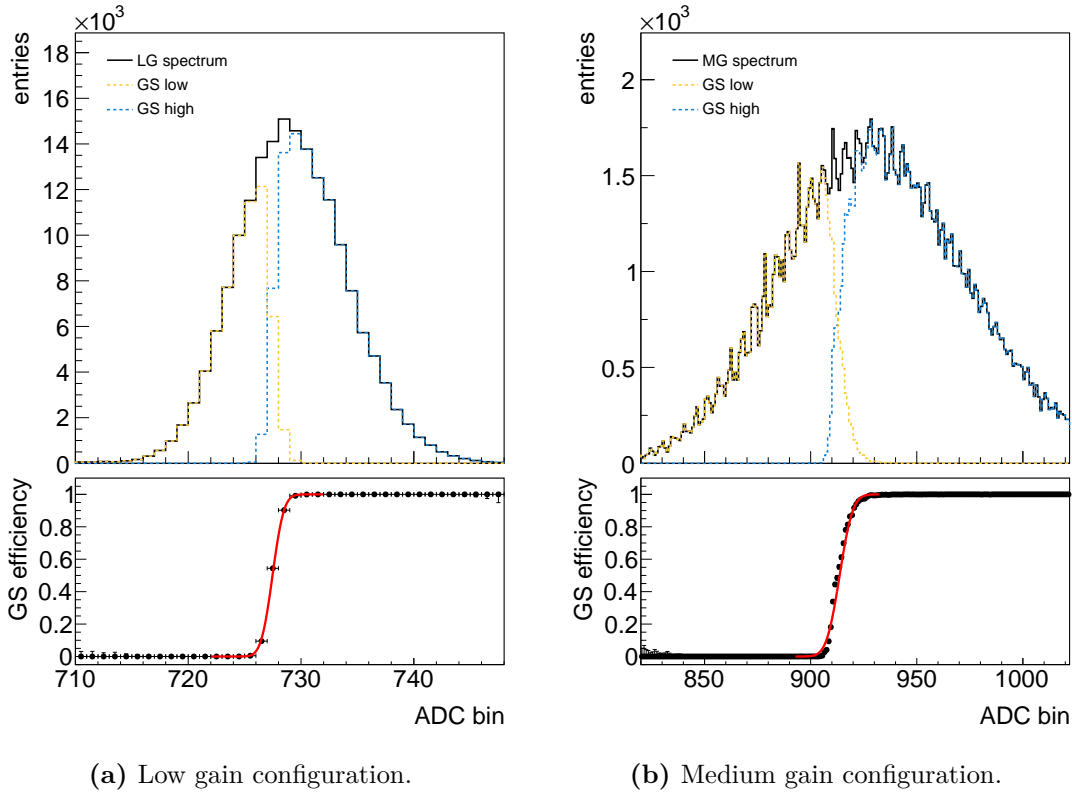


Figure 5.18: Interbranch calibration of the LG and MG configurations. In (a), the spectrum taken with the LG setting enforced is shown. The dashed blue and yellow lines in the upper part of the figure illustrate the spectra of the automatic gain selection. The corresponding GS efficiency, fitted with Equation 5.2, is shown in the lower part. In (b), the equivalent graphs for the MG setting are presented. The laser intensity and the GS setting are slightly adjusted compared to Figure 5.16 in order to fully resolve both the LG spectrum and the MG spectrum.

The same process is repeated for the LG–MG calibration, with a slight adjustment of the laser intensity and the GS setting necessary to fully resolve both spectra. The resulting spectra and GS efficiency curves are shown in Figure 5.18, for the LG configuration in Figure 5.18a and for the MG configuration in Figure 5.18b. After repeating the pedestal measurements, the scaling factor $B_{\text{LG-MG}}$ is calculated analogously to Equation 5.3. Using the channel shown in Figure 5.18 as an example again, the triggering point of the GS is at ≈ 727.5 ADC bins for the LG setting and at ≈ 913.6 ADC bins for the MG setting, resulting in a scaling factor of ≈ 8.5 . The scaling factors of all 16 channels are shown in Figure 5.17b for the LG–MG interbranch calibration.

Combining the results from Figures 5.17a and 5.17b yields the final interbranch scaling factor B ,

$$B \equiv B_{\text{ULG-MG}} = B_{\text{ULG-LG}} \cdot B_{\text{LG-MG}}, \quad (5.4)$$

which is shown for all 16 channels of a matrix in Figure 5.19. The uncertainties on the scaling factors are estimated from the fit uncertainties of the pedestal and threshold values by applying error propagation on Equations 5.3 and 5.4 for Figures 5.17 and 5.19, respectively. Similar to the results of the gain calibration in Figure 5.14, the uncertainties are of the order of ‰ and thus hardly visible in the figures. The interbranch scaling factors in Figure 5.19 are the final results needed for the incorporation of the KLauS ASIC into the light yield setup. As opposed to the gain calibration, the interbranch calibration needs to be performed only once.

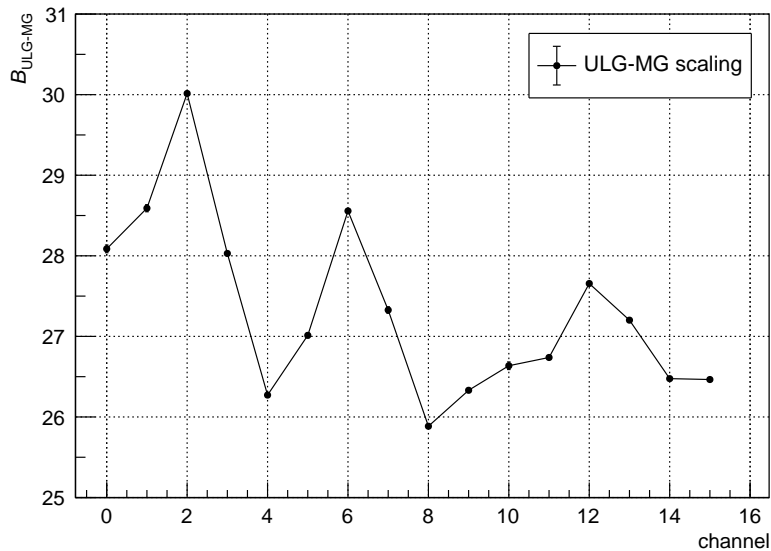


Figure 5.19: Final result of the interbranch calibration. The scaling factors for the conversion between the ULG and MG settings are determined for the 16 channels that were investigated.

5.3.2 Light yield QC

The setup for the light yield QC measurements, shown in Figure 5.20, is an optimised version of the system used for the KLauS calibration. It features a custom adapter board, which connects the tile matrices to a carrier board for the ASIC. The carrier board is in turn connected to an interface board, which directly provides services such as the ASIC power and connects the ASIC to a commercial Raspberry Pi [74] for the data




Figure 5.20: Setup for the light yield measurement. A custom adapter board connects the tile matrices (still assembled on the gluing pedestal) to a carrier board with the KLaUS ASIC. An interface board provides services such as the ASIC power and connects the ASIC to a commercial Raspberry Pi 4 Model B for the data read-out and chip configuration.

read-out and chip configuration. The setup is designed for the matrices to be connected directly after the gluing, i.e. while still assembled on the gluing pedestal, which makes the handling of the matrices easier and safer.

^{22}Na is chosen as a source for the signal generation. It is a radioactive isotope of sodium with a half life of $\approx 2.6\text{ y}$ [75], which decays via a β^+ decay to an excited state of ^{22}Ne . This excited state is very short-lived and decays to the ground state via the emission of a 1 275 keV photon after only 3.7 ps [76]. The positron emitted in the β^+ decay immediately annihilates with an electron, producing two back-to-back photons at 511 keV each. An exemplary spectrum of ^{22}Na , measured with a LYSO:Ce scintillator coupled to a SiPM, is shown in Figure 5.21. While the photons are invisible for the scintillating tiles of the Tile Detector due to the low density of PVT, another feature of the spectrum, the *Compton edge*, can be used. The photons emitted in the ^{22}Na decay scatter on the electrons in the scintillator material, resulting in the plateaus to the left side of both γ peaks. At the maximum possible energy transfer, the plateau is cut off (sharply in the case of the 511 keV peak, less prominently so for the 1 275 keV peak). The photon energy after the scattering can be calculated as [54]:

$$E'_\gamma = \frac{E_\gamma}{1 + (E_\gamma/m_e) \cdot (1 - \cos \theta_\gamma)}$$

For the maximum energy transfer of the 511 keV photons at $\theta_\gamma = 180^\circ$, this yields:

$$E'_{\gamma,\text{max}} = \frac{E_\gamma}{1 + (E_\gamma/m_e) \cdot 2} = \frac{1}{3} E_\gamma \approx 170.3\text{ keV}$$

As the electron is treated at rest prior to being removed from its shell during the Compton scattering process, the maximum electron energy after the scattering is

$$E_{\text{CE}} = E_{\gamma} - E'_{\gamma}{}^{\text{max}} \approx 340.7 \text{ keV}, \quad (5.5)$$

which defines the position of the Compton edge, marked in red in Figure 5.21.

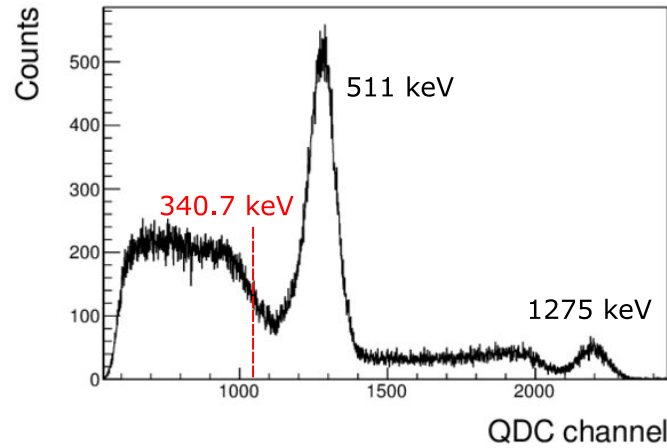


Figure 5.21: Exemplary measurement of the ^{22}Na spectrum, adapted from [77]. The ^{22}Na decays to ^{22}Ne via a β^+ decay, producing the two peaks at 511 keV and 1275 keV marked in the spectrum. To the left of both peaks, plateaus with a cut-off edge are caused by Compton scattering of the photons. The cut-off – also called the Compton edge – belonging to the 511 keV peak is positioned at ≈ 340.7 keV and is used in the light yield QC as a distinguishable feature to map a known energy to the digitised ASIC output.

The concept of the light yield QC is demonstrated in the following paragraphs using the example of the matrix with matrix ID = 02. This lays the groundwork for the finalisation of the light yield QC and its implementation into a common matrix QC system in Section 5.3.3. As a first step, the ^{22}Na spectrum is measured, shown in Figure 5.22a for one of the matrix channels. The ULG setting of the KLauS ASIC is used for this measurement. As discussed above, the two photon peaks are absent, but the Compton edge of the 511 keV γ peak is clearly visible. It is fitted with a convolution of a Gauss error function to fit the edge itself, and a linear function to describe the plateau before the 1275 keV γ peak. The first iteration of the fit function was implemented and described in [78]. The position of the Compton edge μ_{CE} can be extracted from the fit parameters, as shown for the example of the spectrum in Figure 5.22a. This process is performed for all 16 channels of matrix ID = 02, with the resulting Compton edge positions shown in Figure 5.22b. The individual Compton spectra of the 16 channels

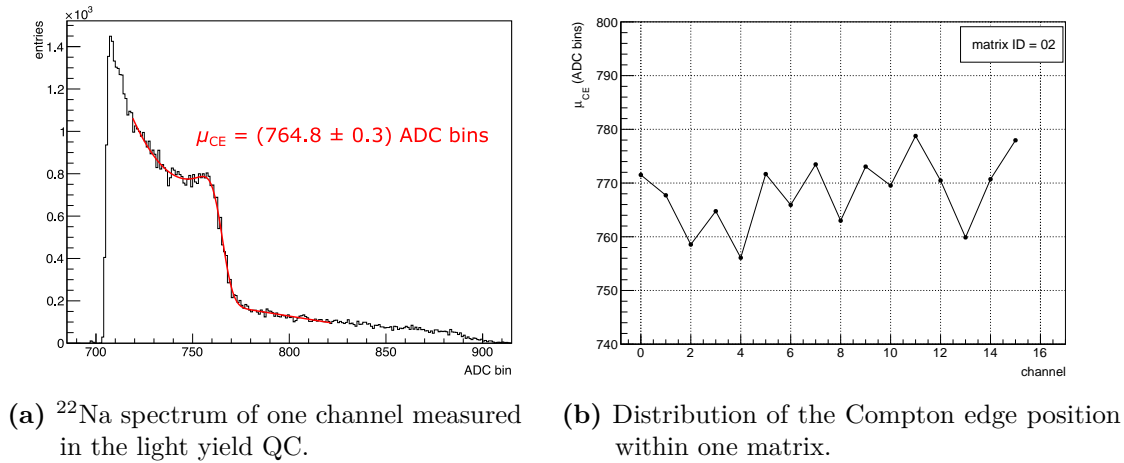



Figure 5.22: Determination of the Compton edge position for the example of matrix ID = 02. In (a), the spectrum of ^{22}Na is fitted with a convolution of an Gauss error function describing the edge and a linear function for the background. The position of the Compton edge μ_{CE} can then be extracted from the fit. This is repeated for all 16 channels of matrix ID = 02, with the resulting Compton edge positions shown in (b).

are shown in Appendix B, Section B.2. The light yield L , which specifies the number of photons/MeV produced in a channel i , is then calculated as:

$$L(i) = (\mu_{\text{CE}}(i) - P_{\text{ULG}}(i)) \cdot \frac{1}{G(i)} \cdot B(i) \cdot \frac{1}{E_{\text{CE}}} \quad (5.6)$$

The Compton edge position μ_{CE} in ADC bins is extracted from the ^{22}Na spectrum as shown in Figure 5.22. The pedestal P_{ULG} , the gain scaling factor G and the interbranch calibration factor B are obtained from the KLauS calibration in Section 5.3.1. Lastly, the energy of the Compton edge as calculated in Equation 5.5, $E_{\text{CE}} = 0.34$ MeV, is needed to normalise the light yield to the number of photons per 1 MeV.

By inserting the Compton edge positions of Figure 5.22b, the gain factors of Figure 5.14 and the interbranch scaling factors of Figure 5.19 into Equation 5.6, the normalised light yield of matrix ID = 02 can be determined. The results are shown in Figure 5.23, as a 1D graph in Figure 5.23a and as a 2D colour map in Figure 5.23b. The uncertainties on the values in Figure 5.23a are propagated from the uncertainties on the factors in Equation 5.6 and are again barely discernible in the figure, as they are on the order of 0.5% to 1%. The resulting light yield is fairly uniform, with an average value of $L_{\text{av}} = (1031 \pm 83)$ photons/MeV.



(a) Light yield, 1D graph.

(b) Light yield, 2D colour map.

Figure 5.23: Results of the light yield measurement of matrix ID = 02, (a) as a one-dimensional graph and (b) as a two-dimensional colour map. Note that the channel numbers refer to the KLauS channels, which are mapped to the SiPM-tile matrix in an unordered fashion and are therefore indicated in the 2D colour map as a guide. The light yield averaged over all 16 channels is fairly uniform, with $L_{av} = (1\,031 \pm 83)$ photons/MeV.

5.3.3 Common setup for the matrix QC

As discussed in Section 5.2, a broader range of features for the SiPM characterisation is advantageous in order to investigate irregular SiPM behaviour. The KLauS ASIC is not only well suited for the light yield measurement, but also allows the study of SiPM properties such as the breakdown voltage V_{br} , the dependence of the gain on the SiPM bias voltage, or the timing threshold behaviour. It was therefore decided to incorporate the SiPM QC concept of Section 5.2 into the light yield system, and to extend it by utilising the available KLauS features. The result is a common matrix QC system, which performs both the SiPM QC and the light yield QC steps illustrated in Figure 5.1. This has several advantages:

- An in-depth SiPM characterisation is possible, allowing the investigation of SiPMs exhibiting unusual behaviour.
- Using one common system, instead of two independent setups, simplifies the measurements and reduces the number of components to maintain.
- No hardware changes of the setup are needed, as the SiPM matrix and the tile matrix are the same in terms of connection and mechanical assembly.
- The data obtained in the SiPM QC can be reused directly in the subsequent QC step for the light yield measurement.

The extension of the light yield system to form a new, common matrix QC system has been performed by the Tile Detector group. This includes the implementation of

new data taking and analysis features customised for the SiPM QC. In addition, the SiPM data is carried over to the light yield QC to keep SiPM-induced variations small, in particular for the gain calibration as discussed in Section 5.3.1. At the time of writing, the new system has been commissioned by the Tile Detector group and is in use for the full Tile Detector production. Figure 5.24 shows an example output of the SiPM QC measurements for one matrix performed with the new system. The full amount of matrices needed for the Tile Detector production have passed through the SiPM QC step, while the light yield QC is currently on-going.




Figure 5.24: SiPM QC measurements performed with the new matrix QC system. The extended SiPM characterisation features include measurements of (a) the SPS, (b) the dependence of the gain on the operating voltage V_{op} , (c) the signal and noise rates depending on the ASIC timing threshold, and (d) the breakdown voltage. The measurements in (a) to (c) are shown for one exemplary channel, while the data in (d) correspond to one full matrix.

5.4 Outlook for the Tile Detector module QC

Following the order of production outlined in Chapter 4, the last step in the QC chain is the qualification of the Tile Detector modules. In the module QC, a full module is evaluated as a unit for the first time, and is studied in order to find settings which allow measurements at PSI to validate the timing performance of the detector. The module QC procedure is therefore tightly interconnected with the calibration of the modules. At the writing of this thesis, several QC steps preceding the module QC are under development. They concern the evaluation of the MuTRiG ASICs, the TMB, and the cooling circuit, and are outlined in the following paragraphs.

MuTRiG

As discussed in Section 4.4, the MuTRiG chips are produced and packaged by two external companies before being assembled to the TMB at the KIP electronics workshop. The first step in the MuTRiG QC is therefore a visual inspection prior to the assembly, which allows sorting out chips with a faulty packaging. Following the visual examination, electronic tests of the ASICs are performed next. A custom QC board with a spring-loaded socket was developed for this purpose, which allows the testing of the packaged chips without the need of soldering them to the TMB (and unsoldering them again, should they fail the tests). However, the socket is not fully reliable in establishing the electrical connections, making an in-depth investigation of MuTRiG parameters difficult. Therefore, only basic functionality tests are foreseen for this step, including

- powering of the ASICs,
- ASIC configuration,
- burn-in tests to identify possible component failures due to thermal stress on the BGA package or the bonding wires.

MuTRiGs passing the above QC measures are cleared for the TMB assembly in the next step. A more detailed investigation of the MuTRiG parameters is foreseen at a later stage during the module calibration.

TMB

As for the MuTRiG chips, the TMB QC foresees tests concerning the basic functionality of the TMB. They are performed using the full Mu3e DAQ chain – that is, tile DAB, FEB and filter farm PC as described in Section 3.2.3 – and include

- powering of the board and the ASICs,
- sending and receiving data of all 13 MuTRiGs,
- synchronisation of clock and reset signals,

- slow control functionality.

Parts of the TMB QC were already developed and evaluated at the testbeam with the preproduction modules discussed in Section 4.4, for example in [68]. At the writing of this thesis, the full QC procedure is in the process of being developed and implemented.

Cooling circuit

Following the production of the aluminium support plates and the gluing of the stainless steel pipes, the final cooling circuits are connected to the Mu3e cooling test system currently set up at the KIP mechanical workshop. Besides mechanical inspections concerning for example connection tightness and possible leakage, the test system allows the investigation of the cooling plates under conditions close to the final Mu3e experiment. In particular, temperature tests down to $-25\text{ }^{\circ}\text{C}$ and long-term temperature monitoring are possible and can be used to validate the cooling circuits. If a circuit has been successfully evaluated, it is transferred to the module production.

Module

Once all individual parts of a module have passed their respective QC tests, they can be assembled and tested as a unit. The matrices are connected to the TMB via the low-height FPC connectors to check

- electrical connections,
- responsiveness of all 416 channels,
- read-out of temperature sensors.

These measurements are supplemented with calibrations of the MuTRiG parameters relevant for the detector performance. First studies of calibration measurements in the laboratory have been performed in [69], using a ^{22}Na source in a setup similar to the light yield QC system. The calibration algorithms were further refined in the testbeam campaign of the preproduction modules presented in Section 4.4.

Figure 5.25 shows an example of the MuTRiG energy threshold calibration for one channel as it was performed during the testbeam. A broad range of E threshold values was scanned and the ToT as a measure for the deposited energy was recorded for each threshold step. The resulting 2D spectrum is shown in Figure 5.25a. The x axis shows the DAC setting of the E threshold in MuTRiG, with a low value of DAC_{ETH} corresponding to a high E threshold, and vice versa. A pattern of two almost horizontal bands is visible in the figure. The upper one at higher ToT values corresponds to the energy deposition of electrons which fully traverse the tile, sometimes also referred to as the *Landau peak* as the deposited energy follows a Landau distribution. The lower band is attributed to electrons passing only through a part of the tile, which also need to be recorded as they are necessary for a meaningful clustering of hits in a later step. At high E thresholds,

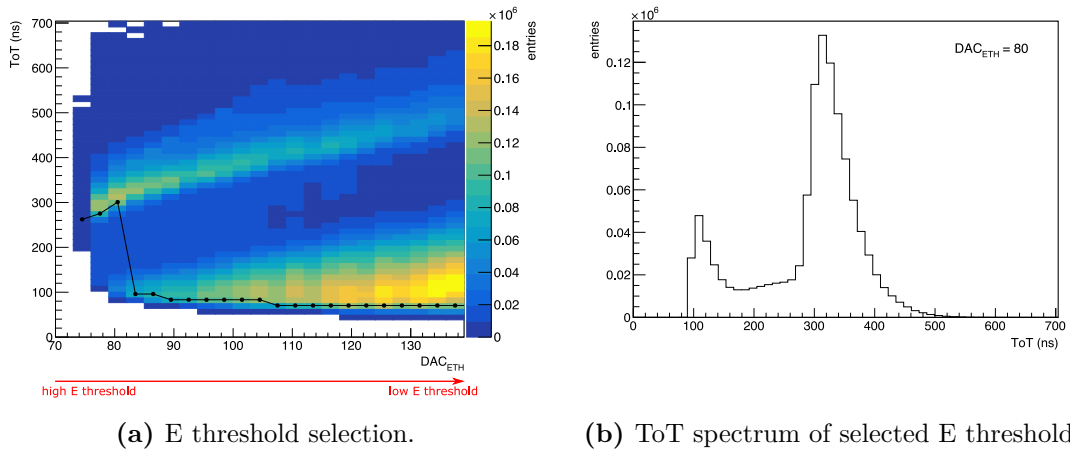


Figure 5.25: E threshold calibration of one channel of a Tile Detector module. **(a)** shows the measured ToT depending on the selected E threshold in terms of MuTRiG DAC values, where a low value of DAC_{ETH} corresponds to a high E threshold, and vice versa. The calibration algorithm selects the DAC_{ETH} value where the Landau peak is still dominating, but the lower peak of particles only grazing the tile is detectable as well. The step in the black graph marks the position where the E threshold is selected by the algorithm. **(b)** shows the ToT spectrum at the selected threshold, $\text{DAC}_{\text{ETH}} = 80$. Figure adapted from [69].

the upper band still has a higher number of hits attributed to it; however, at decreasing E thresholds, the rising number of detected hits on the tile edges increases the entries in the lower band. Ultimately, at very low E thresholds, the ASIC also triggers on the SiPM noise and the lower band dominates the spectrum. The optimum E threshold setting is therefore defined as high enough to yield a dominating Landau peak and to suppress any SiPM noise, but at the same time low enough to still detect the partial energy deposition of the particles on the tile edges.

The calibration algorithm developed in [69] utilises this threshold-dependent behaviour of the ToT spectrum by searching for a two-peak structure with a dominating and a smaller peak. For each setting of the E threshold, the individual ToT spectra are extracted and the highest number of entries in the respective spectrum is identified. Each spectrum is then scanned for the first ToT bin with a value of 50% of the maximum number of entries, which is marked by the black graph in Figure 5.25a. Thus, it is possible to gauge the optimal E threshold by identifying the step in the graph, at which the height of the peak containing the partial energy deposition has increased to at least 50% of the Landau peak. Accordingly, Figure 5.25b shows the ToT spectrum at the DAC_{ETH} value which was selected by the calibration algorithm. The algorithm was automated and applied to all 416 channels of a module during the testbeam campaign, leading to the measurements shown in Figures 4.23 and 4.22 in Section 4.4. At the writing of this

thesis, the calibration algorithms are being optimised for a full module calibration in the laboratory at KIP. Complimentary to the laboratory qualification, further DESY testbeam campaigns for the module calibration are considered as well.

Chapter 6

Detector cooling studies

Over the course of this thesis, the design of the Tile Detector has been continuously adapted to incorporate new insights in the detector assembly or reflect changes in the infrastructure of the Mu3e experiment (see Section 4.1.1 for a detailed discussion). One component that was repeatedly affected by the design revisions was the cooling unit of the Tile Detector module, consisting of the aluminium support plate and the embedded U-shaped pipe. The thermal simulations presented in this chapter allowed studying the effects of the modifications on the performance and the validity of the Tile Detector cooling system. Section 6.1 discusses the initial implementation of the cooling simulation using the first technical prototype, which allowed the tuning of the simulation to laboratory measurements. The resulting simulation setup was then used to study changes of the cooling concept before their physical implementation, as shown in Section 6.2 for different design phases of the cooling circuit. As the irradiation studies performed in [51] found that a much lower temperature than originally foreseen is needed to compensate the increased dark current of the irradiated SiPMs, Section 6.3 summarises the final cooling requirements of the Tile Detector and gives an outlook on the implementation of the final cooling system.

6.1 Cooling simulation and validation of the technical prototype

During the construction and commissioning of the first technical prototype (Section 3.3.2), a test setup was assembled at KIP to evaluate the performance of the Tile Detector cooling circuit. In parallel, a computational fluid dynamics simulation was set up using the CAD implementation of the prototype, with the measurements of the test setup as input. In addition, the test setup also provided parameters for comparison, making it possible to tune the simulated output to the measured data with the goal of having a reliable simulation to study future design changes. The test setup is shown in Figure 6.1 and uses the design of the technical prototype: the packaged STiC 3.1 ASIC – the predecessor of MuTRiG – is assembled on the submodule PCB, which in turn is mounted on the original aluminium support plate (before its height reduction, as discussed in Section 4.1.1). The cooling pipes embedded in the plate are made from copper with an inner diameter of $d_{\text{in}} = 2.5$ mm, and are water-cooled at this stage of the cooling circuit design. The water temperature was set to $T_{\text{inlet}} = 15$ °C, which was considered




Figure 6.1: Test setup for the evaluation of the Tile Detector cooling circuit concept. The design of the first technical prototype is used, with one packaged STiC ASIC assembled to a submodule PCB. Temperature sensors monitor the temperature of the package surface, the cooling plate, the water inlet and outlet as well as the environment.

low enough to see a clear effect of the cooling at room temperature, but high enough to avoid condensation which could damage the electronics. In addition, the volume flow of the water was measured as $\dot{V} \approx 4.7 \cdot 10^{-6} \text{ m}^3/\text{s}$. The STiC chip was operated in three different configurations, which correspond to three different measured power consumption values. These values serve as an input for the simulation and are summarised in Table 6.1.

Several temperature sensors were installed on the test setup to monitor the temperature of the water at the pipe inlet and outlet, of the cooling plate, and of the ASIC package surface. In addition, one sensor remained unattached to the test setup to monitor the ambient temperature. Figure 6.2 shows the resulting temperature measurements for a data taking time of ≈ 40 min, with configuration changes of STiC according to Table 6.1. Prior to the start of the measurement window, the cooling circuit was already turned on to cool down the support plate. The different power consumption levels of the chip are mirrored in the temperature curve of the STiC package surface, with a strong increase of the temperature after changing from 0.9 W to 1.7 W, and a less pronounced increase after the change to 2.1 W. The temperature of the water and the cooling plate, on the

Table 6.1: Simulation input obtained from measurements using the cooling test setup.

Property	Value	Description
T_{inlet}	15 °C	water inlet temperature
v_{inlet}	1 m/s	water flow speed (corresponding to a volume flow of $\dot{V} = 4.7 \cdot 10^{-6} \text{ m}^3/\text{s}$)
T_{a}	21 °C	ambient temperature
P_{outlet}	101 325 Pa	outlet pressure
P_{STiC}^1	0.9 W	STiC power consumption (low)
P_{STiC}^2	1.7 W	STiC power consumption (medium)
P_{STiC}^3	2.1 W	STiC power consumption (high)

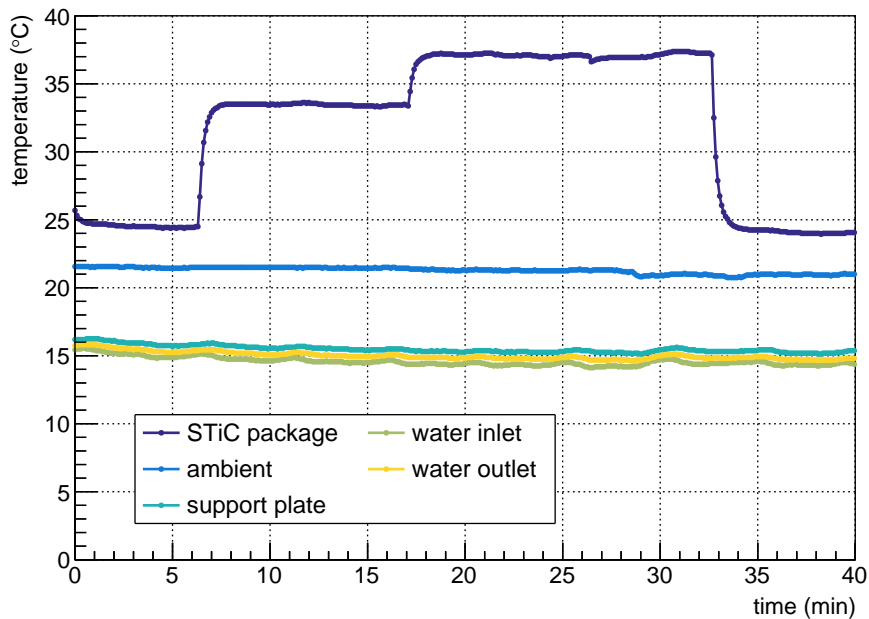


Figure 6.2: Temperatures measured with the cooling test setup. Prior to the start of data taking, the cooling circuit was switched on to cool the support plate down to ≈ 15 °C. During the measurement, the STiC ASIC was operated in three configuration modes corresponding to different power consumption levels, which translate to the varying temperatures measured on the package surface.

other hand, are practically not affected by the increase in power consumption. This can be explained by calculating the expected increase in temperature [79],

$$\Delta T = \frac{P_{\text{STiC}}}{\dot{m} \cdot c_p}, \quad \text{with } \dot{m} = \dot{V} \cdot \rho_W, \quad (6.1)$$

where P_{STiC} is the power consumption of STiC, $c_p = 4.19 \text{ kJ}/(\text{kg K})$ is the specific heat capacity of water, $\dot{V} = 4.7 \cdot 10^{-6} \text{ m}^3/\text{s}$ is the measured volume flow, and $\rho_W = 997 \text{ kg}/\text{m}^3$ is the density of water. Even when inserting the highest power consumption of STiC into Equation 6.1, $P_{\text{STiC}}^3 = 2.1 \text{ W}$, the resulting temperature increase is only $\Delta T \approx 0.1 \text{ K}$. The heat produced by the ASIC is therefore completely absorbed by the cooling liquid with a negligible increase in the water temperature. The small but visible difference between the water inlet and outlet temperatures in Figure 6.2 is more likely to stem from cooling the bulk of the support plate, which is still in contact with the ambient temperature of $T_a \approx 21 \text{ }^\circ\text{C}$.

The temperatures of the three STiC operating modes can be extracted from Figure 6.2 for direct a comparison with the thermal simulation. For this purpose, the cooling test setup of Figure 6.1 was mirrored in the CAD software, and the parameters of Table 6.1 were inserted as boundary conditions and to define the heat source of the simulation. In order to incorporate the heat exchange with the environment, a large box was placed around the prototype model. It is defined as a second fluid domain (alongside the cooling water in the pipes) with air as a medium and a wall temperature of $T_a = 21 \text{ }^\circ\text{C}$ as a boundary condition. The full model used for the simulation is shown in Figure 6.3a, with a zoom on the prototype model in Figure 6.3b. The volume of the





Figure 6.3: Implementation of the test setup of Figure 6.1 in the thermal simulation. (a) shows the full setup, which includes a large air-filled box to model the environment of the laboratory. In (b), a zoomed-in view of the prototype model is shown. The water temperature and flow speed at the pipe inlet are set as boundary conditions for the simulation, and the package of the STiC ASIC is defined as a volume heat source.



(a) Simulated heat transport across the full model. (b) Simulated heat profile on the STiC package surface.

Figure 6.4: Example output of the thermal simulation of $P_{\text{STiC}}^3 = 2.1 \text{ W}$. (a) shows the heat transport across the full surface. The temperature gradient from the cooled support plate in the centre to the walls of the box bound to $21 \text{ }^\circ\text{C}$ is visible. The surface of the STiC heat source is shown in (b), with the size of the package marked for orientation. The black square in the centre sketches the approximate size and orientation of the temperature sensor, which was placed on the package surface in the test setup. To compare the measured and simulated data, the temperature of the package surface is extracted by averaging over the area marked by the black square.

STiC package was implemented as a heat source, and three separate simulations for each STiC power consumption level were performed. In addition, the expected MuTRiG power consumption during normal operation, $P_{\text{MuTRiG}} = 1.2 \text{ W}$, was simulated as well.

In Figure 6.4, the simulation output for the highest power consumption level of STiC is illustrated. Figure 6.4a shows a front view of the full CAD model, including the box imitating the environment. The temperature gradient from the $15 \text{ }^\circ\text{C}$ cooling circuit in the centre to the environment walls fixed at $21 \text{ }^\circ\text{C}$ is clearly visible. The STiC package in the centre reaches up to $47 \text{ }^\circ\text{C}$. A top down view of the package is shown in Figure 6.4b, with the package size ($15 \text{ mm} \times 15 \text{ mm}$) marked in red. As a guide to the eye, the size of the temperature sensor ($\approx 5 \text{ mm} \times 5 \text{ mm}$) is sketched in black. The temperature of the package surface is extracted by averaging over the area of the temperature sensor, i.e. the area marked by the black square. The resulting temperatures obtained from the four simulations are shown in Figure 6.5, with the three temperatures measured with the cooling test setup added for comparison. The uncertainties of the simulated results originate from the averaging method, while for the measured data, the measurement error of the temperature sensor ($\Delta T_{\text{sensor}} = 0.5 \text{ }^\circ\text{C}$) is the dominating uncertainty. In general, there is a good agreement between the simulated and measured temperatures within their uncertainties. The simulation is therefore considered to be adequately tuned




Figure 6.5: Comparison of the simulated and measured temperatures on the STiC package surface as a function of the ASIC power. The results are in good agreement within their respective uncertainties. A fourth simulation point was added for the MuTRiG power consumption, which is used as a reference for future simulations. Adapted from [5].

and ready to be used to evaluate design changes and extensions to the full module. The simulated MuTRiG power consumption, which was added in Figure 6.5 for comparison to the STiC results, from now on serves as the default power consumption setting.

6.2 Extension and adaptation of the simulation

Following the establishment of a tuned thermal simulation using the technical prototype, any subsequent fundamental changes affecting the Tile Detector cooling were simulated to evaluate the effects on the cooling performance. This section provides an overview of the various simulations accompanying the design modifications over the course of this thesis, and thus spans a bridge between the technical prototype in Section 6.1 and the final Tile Detector cooling system in Section 6.3. Table 6.2 summarises the most important parameters and changes implemented in the simulation. Here, the MuTRiG setting of the cooling simulation discussed in the previous section serves as a starting point (and is henceforth referred to as *prototype*). It can be directly compared to the simulation of *configuration I*, which evaluates the effect of the first major rework of the cooling

circuit, including a thinning of the support plate and reduction of the pipe diameter as discussed in Section 4.1.1. While only one MuTRiG is simulated in configuration I to be comparable to the prototype simulation, *configuration II* extends the model to 14 MuTRiGs (corresponding to the default number of submodules at the time). Lastly, the simulation of *configuration III* is mechanically and electronically closest to the final design of the Tile Detector, with a shorter support plate, 13 MuTRiGs and a more detailed estimate of the module power consumption, including the TMB and the SiPMs.

Each of the simulations will be shortly discussed in a dedicated paragraph, with differences with respect to previous results pointed out where applicable. In order to gauge the effect of the cooling circuit modifications, the maximum MuTRiG temperature is chosen as an evaluation parameter. It is therefore also set as a convergence criterion for the simulation, with the requirement that the temperature must be stable within 0.3 °C for 100 iterations for the simulation to converge.

Table 6.2: Sequence of intermediate simulations leading from the technical prototype to the final Tile Detector cooling (Section 6.3).

Name	Power	Temperature	Description
prototype	1.2 W (one MuTRiG)	$T_{\text{inlet}} = 15\text{ °C}$ $T_{\text{a}} = 21\text{ °C}$	first technical prototype (Sections 3.3.2 and 6.1)
configuration I	1.2 W (one MuTRiG)	$T_{\text{inlet}} = 15\text{ °C}$ $T_{\text{a}} = 21\text{ °C}$	thinned support plate (see Section 4.1.1)
configuration II	16.8 W (14 MuTRiGs)	$T_{\text{inlet}} = 15\text{ °C}$ $T_{\text{a}} = 21\text{ °C}$	extension of configura- tion I to 14 submodules
configuration III	40 W (full module)	$T_{\text{inlet}} = 5\text{ °C}$ $T_{\text{a}} = 21\text{ °C}$	first simulation with detailed module power consumption

Configuration I: Thinned cooling plate with one MuTRiG

As briefly mentioned in Section 4.1.1, the thickness of the support plate was reduced by approximately 2 mm to free space for the services on the beam pipe. A direct comparison of the original plate and the thinned down version is illustrated in Figure 6.6. As the modified support plate in Figure 6.6b shows, another design change involved the switch of the top and bottom layer of the ASIC PCB. This implementation allowed a further reduction of the Tile Detector thickness, and added the advantage of the ASIC directly touching the cooling plate for a better heat conduction. However, a cut-out of the support plate was required for both the ASIC package and the components on the PCB, which in





Figure 6.6: Comparison of the cooling circuit design of (a) the technical prototype and (b) the thinned down support plate (\equiv configuration I). The design change also affects the pipe position and size, with a reduction of the inner diameter from 2.5 mm to 1.2 mm as a consequence.

turn affected the pipe position and size. The size was thus reduced to an inner (outer) diameter of $d_{\text{in}} = 1.2$ mm ($d_{\text{out}} = 2.2$ mm). Driven by the availability of physical pipes and pipe sizes, the pipe material was also changed to brass.


In the simulation of configuration I, the new cooling circuit design was first implemented with one submodule assembled, allowing for a direct comparison with the technical prototype. The simulation parameters remain the same as for the prototype listed in Table 6.1. The heat source is simulated using the MuTRiG power consumption, $P_{\text{MuTRiG}} = 1.2$ W. Two exemplary results of the simulation are shown in Figure 6.7. The maximum temperature of the ASIC package volume has increased to ≈ 46.4 °C (see Figure 6.7a), compared to a maximum temperature of 33.4 °C reached for the prototype simulation with MuTRiG settings. Furthermore, a rise of the outlet water temperature is observable in comparison with the prototype: while the difference between water outlet to inlet was $\Delta T \approx 0.1$ °C for the prototype with MuTRiG settings, it has now increased to $\Delta T \approx 1.2$ °C for the thin plate design. As Figure 6.7b shows, the reduced volume flow of the cooling water due to decreased pipe diameter is not enough to fully cool down the support plate, resulting in a greater influence of the ambient temperature. However, even at ≈ 46 °C, the MuTRiG temperature does not affect normal operation.

Configuration II: Extension to 14 MuTRiGs

After the successful implementation of the modified cooling design, the next step was the extension of configuration I to a full module (amounting to 14 submodules at that time). Accordingly, 14 heat sources with $P_{\text{MuTRiG}} = 1.2$ W each were implemented in the simulation, while keeping all other parameters the same. The results are shown in Figure 6.8 and can be directly compared to Figure 6.7. The change from one to 14 heat sources is present in the temperature difference between water inlet and outlet, which has increased to $\Delta T \approx 3.3$ °C. In addition, the average temperature of the support plate has increased from 16.0 °C (single heat source) to 19.7 °C. Evidently, the whole module has significantly warmed up, with the maximum MuTRiG temperature in Figure 6.8a




(a) Simulated heat transport, front view.




(b) Simulated heat transport, support plate.

Figure 6.7: Simulation results of configuration I. (a) shows the heat transport at the centre of the MuTRiG package. The colour scale ranges from the minimum to the maximum simulated temperature. In (b), the colour scale was adjusted by hand to range from the water inlet temperature to the ambient temperature to illustrate the effect of the environment on the cooling.

now reaching 50.9°C . Figure 6.8b also confirms that the reduced volume flow cannot transport away all the heat as it did in the prototype design, and as a consequence the cooling plate is heating up. While temperatures around 50°C are still in a safe range for the MuTRiG operation, it must be noted that this first simulation of a module does not yet include any further read-out electronics besides the 14 ASICs. The next simulation stage, configuration III, contains a more detailed estimate of the full module power consumption.



(a) Simulated heat transport, front view.



(b) Simulated heat transport, support plate.

Figure 6.8: Simulation results of configuration II, to be directly compared to the results of configuration I in Figure 6.7. (a) shows the heat transport at the centre of a MuTRiG assembled in the middle of the support plate. The colour scale ranges from the minimum to the maximum simulated temperature. In (b), the colour scale was adjusted by hand to range from the water inlet temperature to the ambient temperature to illustrate the effect of the environment on the cooling.

Configuration III: Detailed module power consumption

Between the simulation of configuration II and configuration III, another set of significant modifications of the module design, which also affected the cooling circuit, was introduced. First of all, the submodules were replaced by the tile matrices and the TMB, as explained in Section 4.1.1. As the merging of the submodule's MuTRiG PCB with a dedicated read-out board increased the available radial space, it was possible to move the MuTRiG

back to the radially inward side of the TMB, where it was wire-bonded at this early stage (instead of packaged, as in the final design). The cut-out in the cooling plate was however partially kept for other TMB components. In order to still have a good cooling contact, 5 mm × 5 mm aluminium blocks were spared from the cut-out to cool the MuTRiGs from the backside. Accordingly, solid copper tubes were implemented in the TMB design to increase the thermal conductance through the TMB layers to the ASICs. Figure 6.9a shows the CAD implementation of the design at this stage. In order to reduce the complexity and computing time of the simulation, the flex prints for the matrix-TMB connection were suppressed, and the individual scintillating tiles were merged into one large envelope in this model. Furthermore, by optimising the pipe position and the cut-out in the plate, the brass pipes could be replaced by aluminium pipes with very thin walls (250 µm thickness) and a larger inner (outer) diameter of $d_{\text{in}} = 2.0$ mm ($d_{\text{out}} = 2.5$ mm), allowing for a larger volume flow than in configuration II.


For the simulation of the newly modified design, some of the input parameters were adapted. While the ambient temperature, the flow speed and the outlet pressure remained unchanged, the inlet water temperature was set to 5 °C (which was considered the default value for the Tile Detector cooling system at the time). In addition, the full module power consumption was estimated in more detail for the first time, since a finalised TMB design was available, and the TMB is directly cooled by the support plate in the new design. The following assumptions were made for the module power:

- $P_{\text{MuTRiG}} = 13 \cdot 1.5$ W
- $P_{\text{TMB}} = 11.2$ W for all other TMB components
- $P_{\text{SiPM}} = 416 \cdot 6$ mW


The estimate for the MuTRiGs originates from the maximum ASIC power consumption with a safety factor of 10 %. For the SiPMs, the following argument is made: assuming a maximum current of 1.5 mA per matrix is provided by the HV power supplies, the current per SiPM would amount to ≈ 100 µA. With a maximum voltage of 60 V, the highest power consumption of a single SiPM is estimated to be 6 mW¹. The three contributions to the module power consumption were modelled as individual volume sources, similarly to the implementation of the 14 MuTRiGs in the previous simulation.

Figure 6.9b shows one example of the simulation output. The cooling of the MuTRiGs is not as effective as in the previous configurations, since there is less contact of the PCB and the support plate (compared to the prototype design), and the ASIC has no direct contact with the plate either (compared to the design used in configurations I and II). In addition, the TMB itself is now a heat source, and is evidently better cooled in the areas directly touching the support plate. The maximum MuTRiG temperature has

¹It should be noted that this estimation was performed prior to the irradiation studies in [51], which resulted in an estimate of ≈ 6.3 mW per SiPM (assuming a 30 % recovery through annealing). In fact, as a direct consequence of the results presented in [51], the number of designated Tile Detector HV supplies was increased in order to cope with the increased current of the irradiated SiPMs.



(a) Simulation CAD model.



(b) Simulated heat transport, front view.

Figure 6.9: Simulation of configuration III. In (a), the CAD implementation of configuration III is shown. Note that at this stage of the module design, the bare (unpacked) MuTRiGs were wire-bonded to the TMB on the far side of the support plate. The ASICs are modelled accordingly, with a size of $5\text{ mm} \times 5\text{ mm} \times 0.2\text{ mm}$. To simplify the simulation, the scintillating tiles are merged into a tile envelope and the flex prints connecting the tile matrices to the TMB are suppressed. (b) shows the heat transport at the centre of a MuTRiG assembled in the middle of the support plate. The colour scale ranges from the minimum to the maximum simulated temperature.

therefore increased to 59.2°C , despite the fact that the water temperature was decreased compared to the previous simulations. This explains why the average temperature of the support plate is at 13.5°C lower than in both configuration I and II, but the maximum ASIC temperature is higher. In Figure 6.10, the maximum temperatures of the 13 individual MuTRiGs are compared, sorted by the position of the ASIC on the support plate: position 0 is here defined as closest to the cooling pipe inlet and outlet, while position 12 is closest to the target. A clear position-dependent increase in temperature is visible, with a temperature difference of $\approx 3.5^\circ\text{C}$ between the first and the last ASIC. Still, all ASICs can be sufficiently cooled for safe operation.




Figure 6.10: Comparison of the MuTRiG maximum temperatures depending on their position on the support plate. MuTRiG 0 is assembled closest to the water inlet and outlet, and therefore has the lowest maximum temperature. On the other end of the scale are MuTRiGs 11 and 12 at the target side of the module, which are $\approx 3.5^\circ\text{C}$ warmer.

The previous paragraphs demonstrate how the thermal simulation allowed the evaluation of necessary design modifications, and thus accompanied several iterations of the Tile Detector cooling circuit. The last major impact on the detector cooling was the decision to move to a liquid cooling system running at $\approx -20^\circ\text{C}$ for both the SciFi detector and the Tile Detector. This required one more iteration of the cooling circuit design, which resulted in the design presented in Section 4.1. An outlook on the final cooling system, including a discussion of the requirements and the implementation, will be presented in the next section.

6.3 Outlook on the final cooling system

Irradiation studies of the SiPMs performed for both the SciFi detector [50] and the Tile Detector [51] showed that it is not possible to operate the SiPMs of both systems for the full run time of Mu3e phase I without any means of mitigation: in the case of the SciFi detector, too high noise rates make it impossible to resolve single photons, while for the Tile Detector, the high dark current will exceed the current limit of the HV supplies. One consequence is the establishment of a common cooling system, which provides liquid cooling for both timing detectors at temperatures of -20°C to -25°C . Section 6.3.1

summarises the finalised requirements posed by the Tile Detector on the common cooling system, and Section 6.3.2 then gives an outlook on its development and implementation.

6.3.1 Final cooling requirements

Similarly to the first detailed estimate of the module power consumption in the previous section, the final power consumption per module can be determined by separating the contributors: first, the nominal (maximum) MuTRiG power consumption is 1.2 W (1.5 W), including a safety factor of 10 %. Second, the power consumption of the remaining TMB components is estimated as $3.6 \text{ V} \cdot 0.8 \text{ A} \approx 3 \text{ W}$. Last, the SiPM power consumption per module is expected to reach 2.63 W, including dark counts (1.88 W) and signal (0.75 W) at -10°C and 55 V operating voltage [51]. A recovery of the SiPMs of 30 % through annealing is assumed. This results in $\approx 25 \text{ W}$ plus 20 % safety margin, yielding a power consumption of 30 W per module. The following requirements are thus posed by the Tile Detector on the common cooling system:


1. The Tile Detector cooling system is required to cool 210 W per station, or consequently 420 W for the full detector.
2. The nominal temperature at the SiPMs is defined as -10°C to mitigate the irradiation effects [51]. This translates to the requirement of -20°C at the inlet of the Tile Detector modules to ensure adequate cooling of the SiPMs.
3. The temperature stability is required to be 0.2°C over time.

6.3.2 System development and implementation


A diagram of the final Tile Detector cooling system is shown in Figure 6.11a: the two Tile Detector stations are connected to two manifolds, which distribute the cooling liquid to the modules. The manifolds are in turn connected via a pair of pipes to the in- and outlet of the cooling unit of the liquid detector cooling system, a Huber Unistat 425W-spec. The flow in the upstream and downstream stations of the Tile Detector can be adjusted via plug valves, which are installed between the Unistat and the manifold. A CAD implementation of the cooling circuit of one station, including the custom designed manifold, is shown in Figure 6.11b.

Using the thermal simulation established above, a parametric study of the final Tile Detector cooling circuit was conducted in order to investigate whether the required temperature of -10°C can be achieved. A silicon oil ² was chosen as a cooling medium, since it offers a relatively low viscosity at low temperatures. The volume flow rate was varied at two different inlet temperatures, $T = -20^\circ\text{C}$ and $T = -25^\circ\text{C}$, which correspond to two different kinematic viscosities of the silicon oil, $\nu = 7.2 \text{ mm}^2/\text{s}$ and $\nu = 9 \text{ mm}^2/\text{s}$. Figure 6.12a shows the pressure difference between the pipe inlet and

²Huber SilOil M80.055.03




(a) Diagram of the Tile Detector cooling system.



(b) CAD implementation of the Tile Detector cooling system.


Figure 6.11: Final Tile Detector cooling circuit. (a) shows a diagram of the cooling circuit connection. In (b), the CAD model of the circuit of one station is shown.

outlet of one module depending on the volume flow rate for both temperature settings. In Figure 6.12b, the maximum temperature of the support plate is shown depending on the pressure difference, likewise for the two temperature settings. The requirement of -10°C is sketched in this figure as well. For the lower inlet temperature, $T = -25^{\circ}\text{C}$, the requirement can be fulfilled at a pressure difference of ≈ 500 mbar, corresponding to a volume flow of roughly 200 mL/min. Figure 6.13 shows the resulting simulation output at these setting, with a front view of the module in Figure 6.13a and the close-up of the support plate in Figure 6.13b. The required temperature of -10°C is reached for a large fraction of the cooling plate, with a relatively uniform distribution. As the Unistat is capable of delivering higher flow rates and pressure differences than assumed in this simulation, the cooling requirement is expected to be accomplishable.




(a) Volume flow rate vs pressure difference. (b) Pressure difference vs support temperature.

Figure 6.12: Parametric simulation study of the final Tile Detector cooling circuit. In (a), the pressure difference between the pipe inlet and outlet of a module is shown depending on the volume flow rate. (b) shows the dependence of the maximum support temperature on the same pressure difference. For a minimum pressure difference of ≈ 500 mbar, the requirement of -10°C , drawn as a grey dashed line, is fulfilled.



(a) Simulated heat transport, front view.



(b) Simulated heat transport, support plate.

Figure 6.13: Simulation results of the final Tile Detector cooling circuit. (a) shows the heat transport in the support plate, with the colour scale ranging from -25°C to 25°C for better visibility. In (b), the colour scale was adjusted to demonstrate the uniformity of the required temperature of -10°C .




Figure 6.14: Test system for the liquid detector cooling. The Tile Detector cooling circuit is assembled on a cooling pipe, which is mounted on a mock-up of the detector cage with the correct Mu3e dimensions. The Tile Detector distribution manifold is mounted on the cage as well, but partially hidden in the photo behind the low voltage connections.

A test system for the liquid detector cooling was set up at the KIP workshop using the Unistat and the silicon oil as well as the (partial) detector cooling circuits, including the new Tile Detector manifold. Furthermore, the setup was equipped with sensors to monitor the temperature, pressure, and flow rate. A photo of the Tile Detector cooling circuits assembled on the test setup is shown in Figure 6.14. In an early cooling test using the test system, the Unistat was set to -20°C and the temperature of the Tile Detector cooling circuit was monitored. As the tests in the KIP workshop were performed at room temperature and in normal atmosphere, the cooling of the surroundings and the formation of ice on the cooling elements could not be prevented. Still, the temperature at the Tile Detector pipe was able to reach -8°C ³. Moreover, through extensive optimisation of the flow in the full cooling circuit by utilising valves and bypasses, the flow rate required in Figure 6.12b could be validated. At the writing of this thesis, the liquid detector cooling system is being tested and prepared in Heidelberg for the long-term installation at PSI.

³Since the Mu3e experiment will be operated in a dry helium atmosphere, no condensation is expected in the detector volume.

Chapter 7

Conclusion

The Mu3e experiment is designed to search for the charged lepton flavour violation decay $\mu \rightarrow eee$ with an unprecedented single event sensitivity of $2 \cdot 10^{-15}$ in phase I, and ultimately 10^{-16} in phase II. As the decay is strongly suppressed in the Standard Model (SM) with a branching ratio of $\mathcal{B}(\mu \rightarrow eee) < 10^{-54}$, an observation of the process $\mu \rightarrow eee$ would be a clear sign of physics beyond the SM. In order to achieve the projected signal sensitivity, background suppression is crucial and requires excellent vertex, momentum and timing capabilities of the detector systems. With a timing resolution of less than 100 ps, the Tile Detector is designed to provide the most accurate time information in the Mu3e experiment and is vital for the suppression of combinatorial background. It utilises a plastic scintillator with a short rise time and silicon photomultipliers (SiPMs) for the scintillating light detection as active detector material, with a custom ASIC, the MuTRiG, for the signal read-out.

In this work, the design and construction of the Tile Detector for Mu3e phase I is described. The first technical prototype of the Tile Detector was manufactured to evaluate the feasibility of the assembly procedure, the effectiveness of the cooling circuit, and the performance, in particular with regard to the timing resolution. The prototype was validated during two measurements campaigns at the testbeam facility of the Deutsches Elektron-Synchrotron (DESY), which yielded a single-channel time resolution of ≈ 47 ps, clearly below the requirement of 100 ps.

With the prototype design as the baseline, the Tile Detector design was continuously adapted to improve the construction procedure, and to enable the mechanical integration into the full experiment. The finalised Tile Detector design was developed and implemented over the course of this thesis. From the experience gained during the prototype assembly, a manufacturing line for the tile matrix production was designed, constructed and evaluated with 62 matrices, which were subsequently used for the assembly of two full preproduction modules. In a measurement campaign at the DESY testbeam facility, the preproduction modules were successfully commissioned and calibrated. Excluding the channels with known faulty connections yielded more than 96 % of fully functional channels.

For each step in the manufacturing line, quality control (QC) procedures were established. The QC measurements of the scintillating tiles resulted in the optimisation of the scintillator production method, yielding deviations of less than $\pm 10 \mu\text{m}$ in the length and

width of the tiles, and less than $\pm 20 \mu\text{m}$ in the tile height. The first version of the SiPM QC concept was evaluated using the 62 preproduction matrices. All matrices could be operated and were well below the defined noise rate limit of 500 kHz, thus passing the QC criteria ensuring a smooth operation of the detector. Lastly, the light yield setup for the tile matrix QC was tested and improved. The calibration of the KLauS ASIC, which is used in this QC setup, was performed and first measurements were undertaken to show that the system operates as expected. These efforts, combined with the experience gained during the extensive operation of the SiPM QC, laid the foundation for a new common QC system which combines the light yield measurements with improved SiPM performance evaluation capabilities. This new system has been commissioned and further improved by the Tile Detector group, and is currently in use for the full Tile Detector matrix production.

In parallel to the production and assembly of the first technical prototype, a thermal simulation was set up to evaluate the performance of the cooling circuit. The input for the simulation was extracted from laboratory measurements with the prototype setup, which also provided parameters for the simulation tuning. A very good agreement between simulation and data was achieved. Advancements in the detector production and the integration of the Tile Detector into the Mu3e experiment repeatedly required modifications of the Tile Detector cooling circuit. The tuned simulation was used to evaluate the effect of these modifications on the cooling performance. Three configurations of intermediate design stages were validated, leading to the finalised cooling circuit which is now implemented in the liquid cooling system for the two timing detectors of Mu3e.

Outlook

At the time of writing, the Tile Detector production is in full operation. All tile matrices have been produced, with the QC procedures currently in progress. The new common QC system for both the SiPM matrices and the tile matrices has been commissioned and is being further optimised while already being in operation. The full amount of MuTRiG ASICs has been received and is being validated, with a QC system in place to test the basic functionality of the ASICs. The TMB assembly, and the accompanying QC procedures, are being prepared and are expected to start in the following months. All of the cooling circuits needed for two detector stations have been produced. In addition, the liquid cooling system for the two timing detectors has been commissioned in Heidelberg, and will be installed in the Mu3e beam area at PSI this year. Lastly, the first Tile Detector modules are foreseen to be commissioned and installed at PSI at the end of this year, together with the vertex detector and the Scintillating Fibre detector. First calibrations and engineering runs are foreseen at the end of this year and throughout the next year, with the Tile Detector expected to be fully installed by mid-2025.

Appendices

Appendix A

Supplementary material for the Tile Detector dimensions

This chapter provides detailed information concerning the Tile Detector sizes discussed in Chapter 4. In Table A.1, an extensive list of dimensions is compiled, including a short description and an explanation concerning the chosen size. For a better visualisation, sketches of the Tile Detector are provided in Figure A.1, with the relevant dimensions marked.

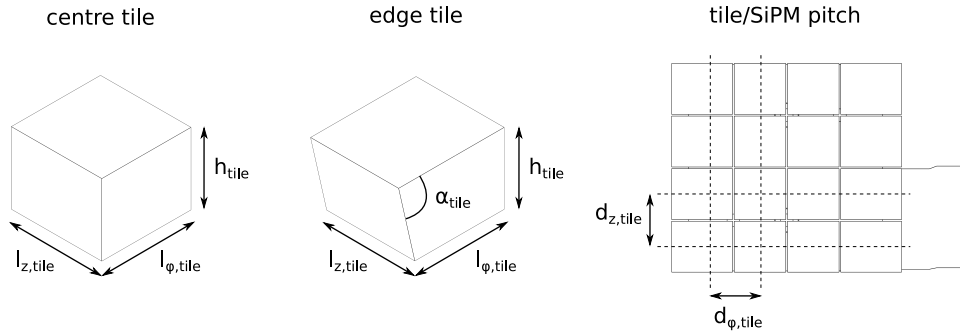
The calculation of the full station length listed in Table A.1 is given by

$$\begin{aligned} l_{z,\text{TD}} &= 13 \cdot l_{z,\text{matrix}} + 12 \cdot w_{z,\text{inter}} + 2 \cdot w_{\text{endring}} + 2 \cdot 0.1 \text{ mm} \\ &= 342.32 \text{ mm}, \end{aligned} \tag{A.1}$$


with the matrix length in z direction $l_{z,\text{matrix}}$, the z gap width between matrices $w_{z,\text{inter}}$, and the endring thickness w_{endring} as defined in the table. Additionally, $2 \cdot 100 \mu\text{m}$ are added to the support plate length in order to have a bit more margin on each end of the sensitive tile volume. The resulting $l_{z,\text{TD}} \approx 342 \text{ mm}$ is the length quoted in Section 4.1.

Table A.1: List of Tile Detector dimensions, with explanations of the individual sizes.
For a visual representation of the dimensions, refer to Figure A.1.


Symbol	Description	Value	Explanation
$d_{\varphi,\text{tile}}$	tile pitch in φ direction	6.50 mm	SiPM pitch in φ , defined by the technical prototype
$d_{z,\text{tile}}$	tile pitch in z direction	6.41 mm	SiPM pitch in z , defined by the technical prototype
h_{tile}	tile height	5.00 mm	defined by the technical prototype
$l_{\varphi,\text{tile}}$	tile length in φ direction	6.30 mm	tile pitch minus intra gap size in φ
$l_{z,\text{tile}}$	tile length in z direction	6.21 mm	tile pitch minus intra gap size in z
α_{tile}	edge tile angle	77.14°	calculated as $90^\circ - 360^\circ/14 \cdot 2$
$w_{\varphi,\text{intra}}$	φ gap width within a matrix	0.20 mm	compensation for foil width ($2 \cdot 80 \mu\text{m}$)
$w_{\varphi,\text{inter}}$	φ gap width between matrices	0.15 mm	additional margin for mechanical assembly
$w_{z,\text{intra}}$	z gap width within a matrix	0.20 mm	compensation for foil width ($2 \cdot 80 \mu\text{m}$)
$w_{z,\text{inter}}$	z gap width between matrices	0.15 mm	additional margin for mechanical assembly
$l_{\varphi,\text{matrix}}$	matrix length in φ direction	26.0 mm	given by 4 times the φ tile pitch, i.e. $4 \cdot 6.5 \text{ mm}$
$l_{z,\text{matrix}}$	matrix length in z direction	25.64 mm	given by 4 times the z tile pitch, i.e. $4 \cdot 6.41 \text{ mm}$
$d_{z,\text{matrix}}$	matrix pitch in z direction	25.79 mm	matrix length plus inter gap width in z
$l_{z,\text{TD}}$	full station length	342.32 mm	given by Equation A.1
r_{out}	outer radius	64.11 mm	restricted by pixel recur layers
r_{in}	inner radius	47.19 mm	restricted by helium ducts
w_{endring}	endring thickness	3.5 mm	inherited from technical prototype



(a) Scintillating tile dimensions and pitch.



(b) Tile matrix dimensions and gap widths.



(c) Tile Detector envelope dimensions and tile matrix pitch.

Figure A.1: Sketch of the Tile Detector dimensions. In (a), the sizes of the two types of tiles, the edge tile angle, and the tile/SiPM pitch in both the direction of the azimuthal angle φ and the beam axis z are labelled. (b) defines the size of a matrix, as well as the gap widths *within* a matrix and *between* two neighbouring matrices, both in φ and in z . In (c), the matrix pitch in z and the envelope-defining dimensions, i.e. length and inner/outer radius, are labelled. The final sizes are summarised in Table A.1.

Appendix B

Supplementary material for the quality control

This chapter compiles additional information on the development of the the quality control procedures in Chapter 5. Section B.1 provides supplementary diagrams of all 16 matrix channels for the gain and interbranch calibration steps of the KLauS ASIC discussed in Section 5.3.1. In Section B.2, additional diagrams for the light yield quality control in Section 5.3.2 are shown.

B.1 KLauS calibration

B.1.1 Gain calibration




Figure B.1: Single-photon spectra of the 16 channels of the matrix used for the KLauS gain calibration. As opposed to the spectrum shown in Figure 5.13a, the ASIC pedestal has not been subtracted for the shown spectra.

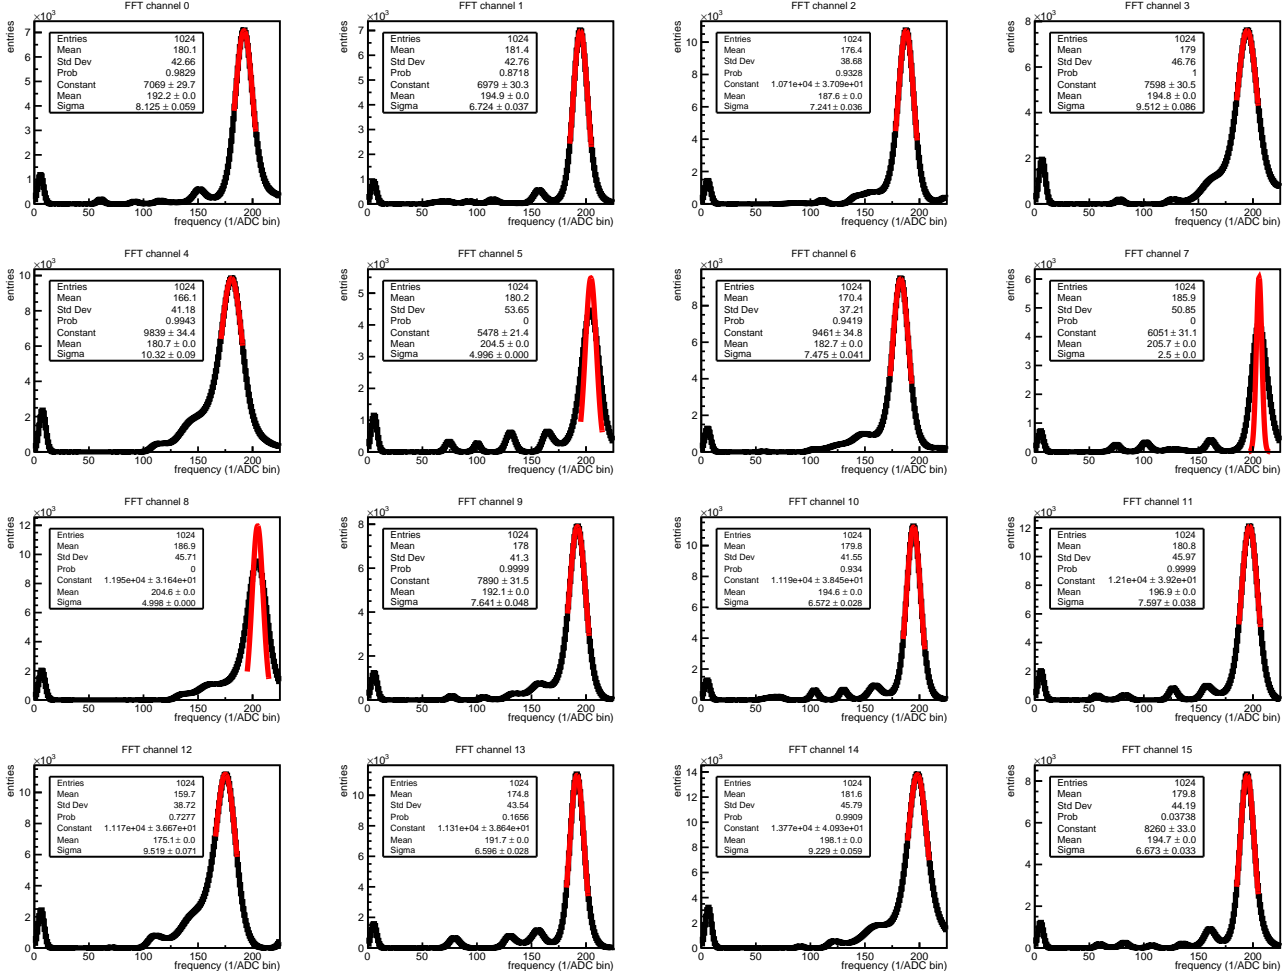


Figure B.2: Fast Fourier transform spectra of the single-photon spectrum (SPS) shown in Figure B.1. A Gauss function is fitted to extract the position of the highest peak. For channels 5, 7 and 8, the fit does not describe the measured data well; the bin with the maximum entries of the spectrum is used instead.

B.1.2 Interbranch calibration

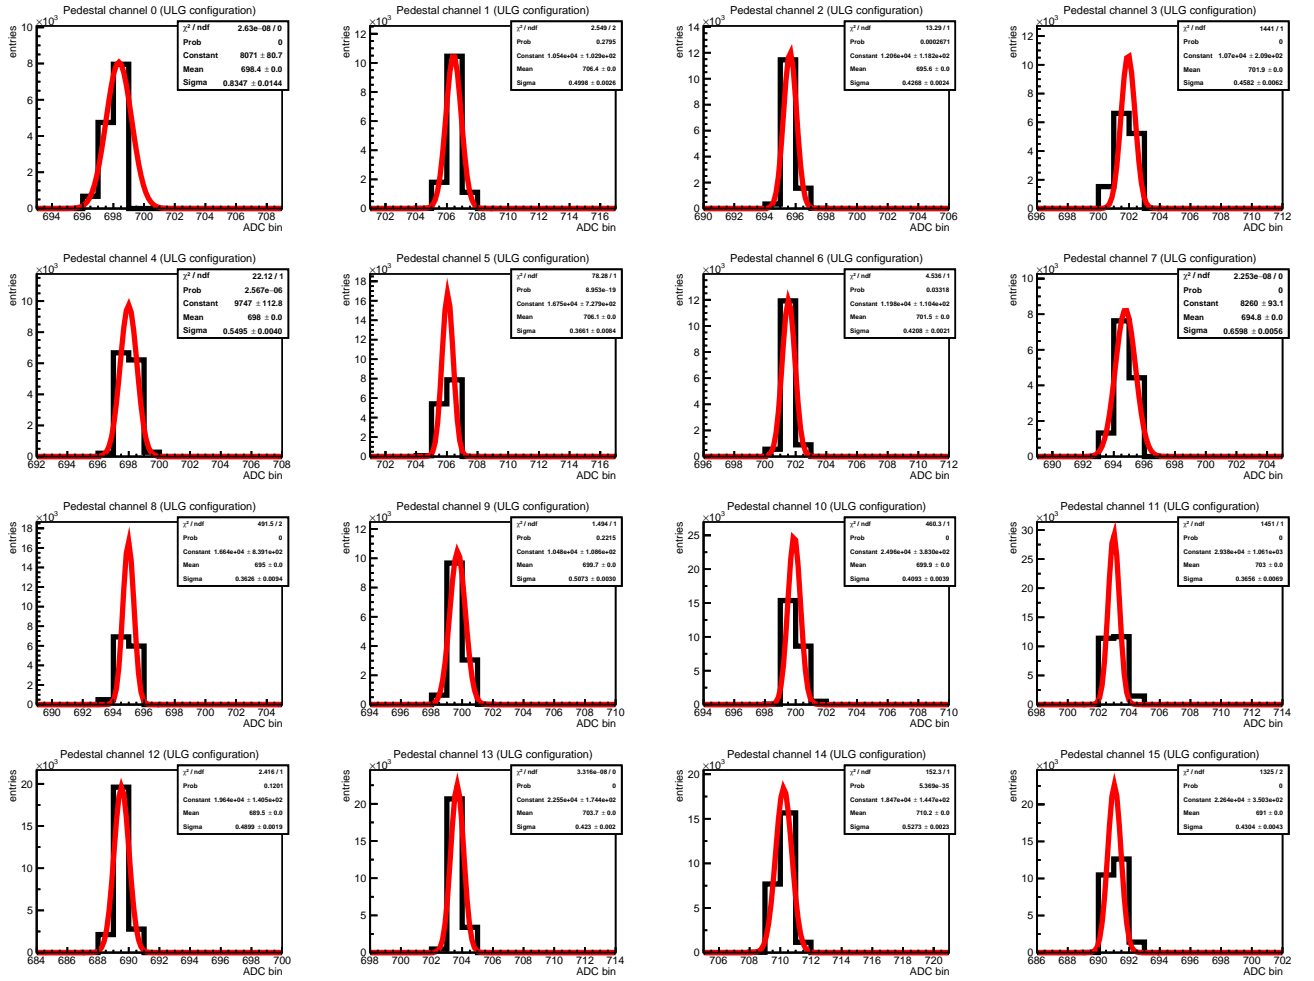


Figure B.3: Ultra-low gain (ULG) configuration pedestal measurements for all 16 channels of the matrix used for the interbranch calibration.




Figure B.4: Low gain (LG) configuration pedestal measurements for all 16 channels of the matrix used for the interbranch calibration.

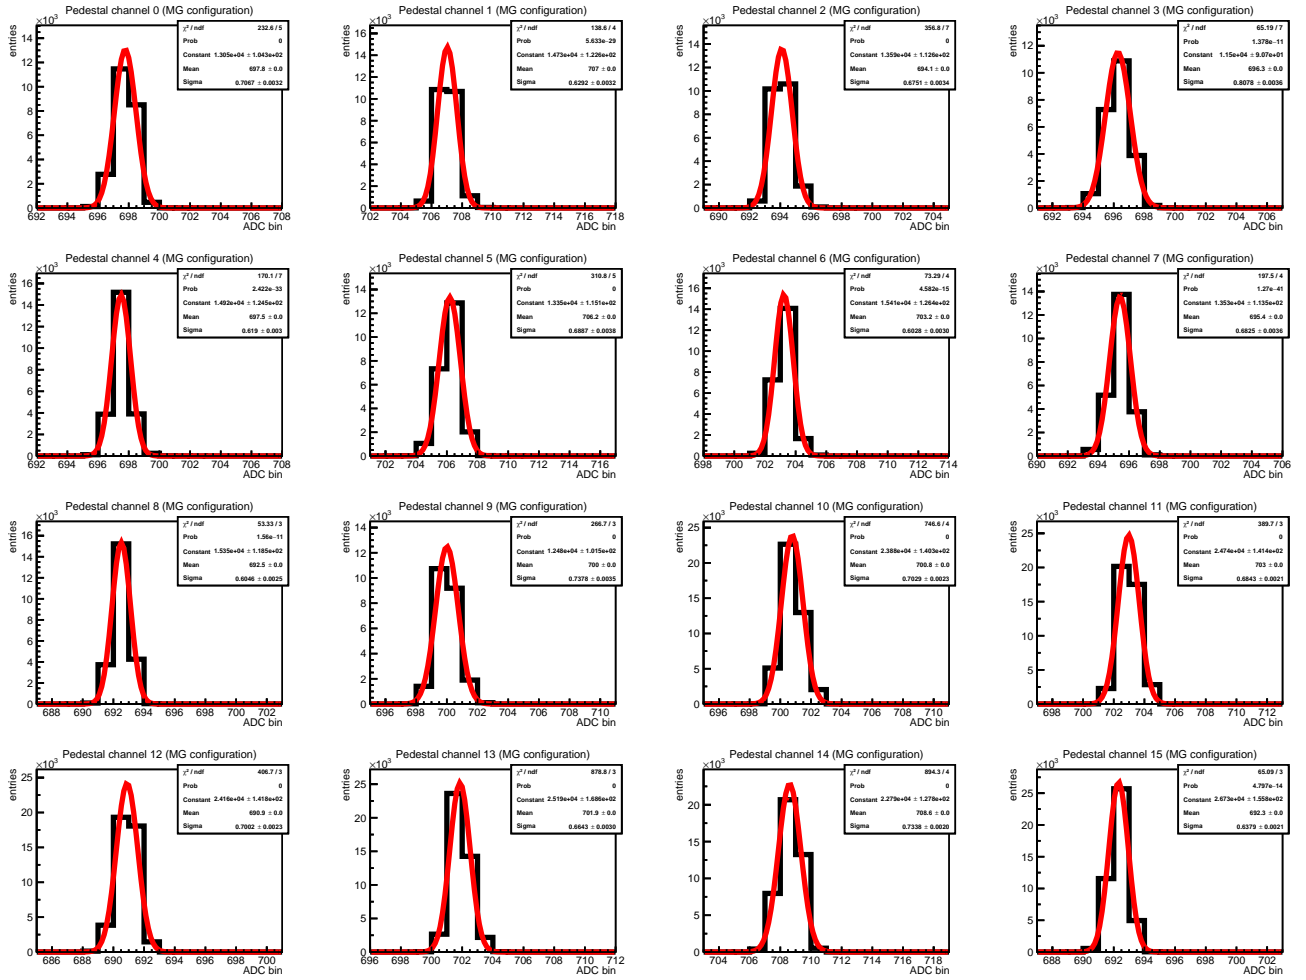


Figure B.5: Medium gain (MG) configuration pedestal measurements for all 16 channels of the matrix used for the interbranch calibration.




Figure B.6: Gain selection spectra for the ULG configuration of all 16 channels of the matrix used for the ULG-LG interbranch calibration.




Figure B.7: Gain selection efficiency for the ULG configuration of all 16 channels of the matrix used for the ULG-LG interbranch calibration.

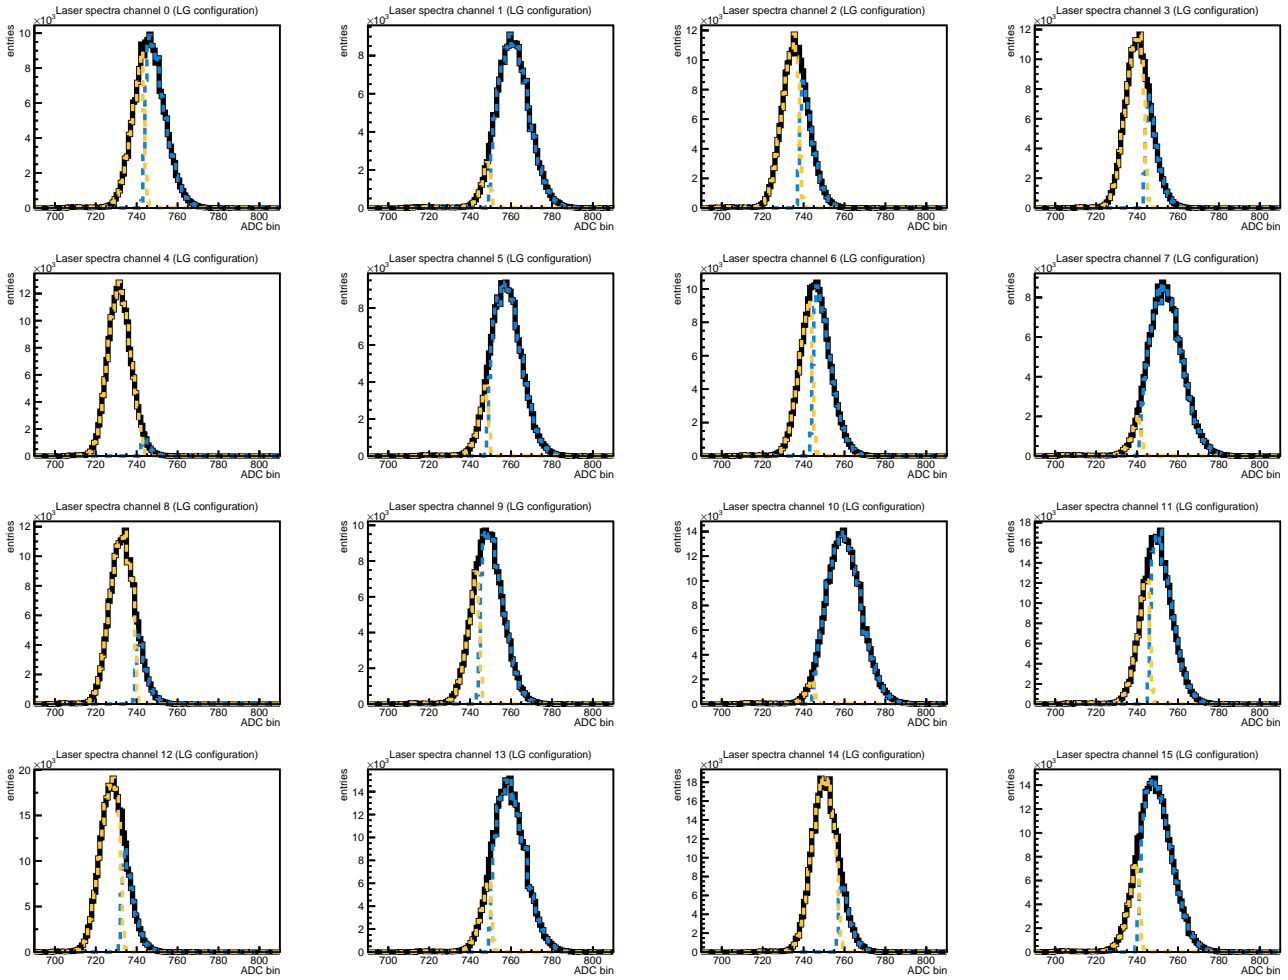


Figure B.8: Gain selection spectra for the LG configuration of all 16 channels of the matrix used for the ULG-LG interbranch calibration.

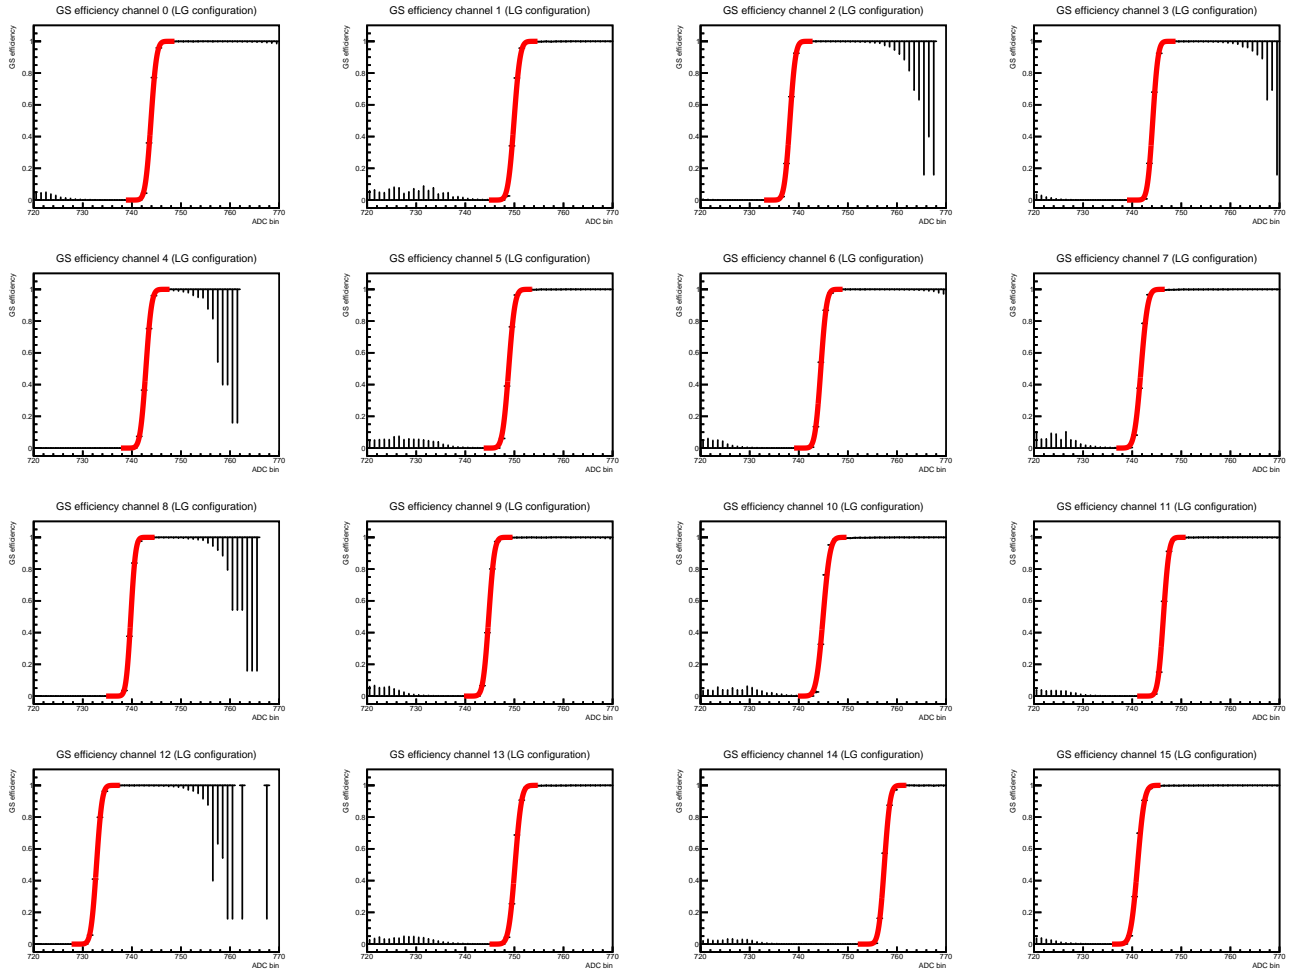


Figure B.9: Gain selection efficiency for the LG configuration of all 16 channels of the matrix used for the ULG-LG interbranch calibration.




Figure B.10: Gain selection spectra for the LG configuration of all 16 channels of the matrix used for the LG-MG interbranch calibration.




Figure B.11: Gain selection efficiency for the LG configuration of all 16 channels of the matrix used for the LG-MG interbranch calibration.




Figure B.12: Gain selection spectra for the MG configuration of all 16 channels of the matrix used for the LG-MG interbranch calibration.




Figure B.13: Gain selection efficiency for the MG configuration of all 16 channels of the matrix used for the LG-MG interbranch calibration.

B.2 Light yield QC




Figure B.14: Determination of the Compton edge position for the all channels of matrix ID = 02. The spectra taken with the ^{22}Na source are fitted with a convolution of a Gaussian error function describing the edge and a linear function for the background to extract the position of the Compton edge μ_{CE} .




Figure B.15: Distribution of the light yield results of matrix ID = 02.

Appendix C

Lists

C.1 List of Figures

2.1	Visualisation of the Standard Model particles.	6
2.2	Feynman diagrams of the $\mu \rightarrow eee$ decay.	11
2.3	History of lepton flavour violation experiments using muons.	13
3.1	Sketches of the Mu3e signal and background.	16
3.2	Contamination of the signal region with internal conversion events.	17
3.3	Conceptual design of the Mu3e experiment.	19
3.4	CAD model of the $\pi E5$ beam line and CMBL.	20
3.5	Stopping target for the Mu3e experiment.	21
3.6	Innermost layer of the vertex detector prototype.	23
3.7	Scintillating Fibre detector.	23
3.8	CAD rendering of the detector cage.	25
3.9	Simulated performance of the Mu3e experiment for phase I.	27
3.10	Original Tile Detector design.	28
3.11	Base ingredients of the Tile Detector.	29
3.12	Single-photon spectrum of a SiPM.	31
3.13	Principle of the signal discrimination in the MuTRiG chip.	32
3.14	Testbeam measurements with the first technical prototype of the Tile Detector.	34
4.1	CAD model of a Tile Detector station.	38
4.2	CAD model of a Tile Detector module.	39
4.3	Photo of a Tile Module Board (TMB).	40
4.4	Photo of a Tile Detector cooling circuit.	40
4.5	Smallest modular building block of the Tile Detector.	41
4.6	Cross section of the Mu3e CAD model (downstream side).	42
4.7	CAD based sketches showing the assembly differences of the technical prototype submodule and the new tile matrix design.	44
4.8	Definition of the dimensions of the scintillating tiles and the tile matrices.	46
4.9	Scintillating tile production.	48
4.10	Comparison of the two tile production methods.	49

4.11	Sketch of the setup used to investigate the different tile types at the DESY testbeam facility.	50
4.12	Comparison of the timing resolution achieved with the old prototype tiles and new ice-vice tiles at the DESY testbeam.	51
4.13	Distribution of the operating voltage V_{op} of the 1000 SiPMs used for the production of the two preproduction modules.	52
4.14	Photo of an assembled SiPM matrix.	53
4.15	Examples of tile wrapping.	55
4.16	Tile wrapping tools.	56
4.17	Photo of a gluing tool.	57
4.18	Gluing procedure using the custom-made gluing tool.	58
4.19	Fully assembled preproduction module.	60
4.20	Commissioning setup of the preproduction modules at the DESY testbeam facility.	62
4.21	Commissioning setup of the preproduction modules at the DESY testbeam facility, with the PVC pipe removed.	63
4.22	Hit map of a preproduction module after calibration during the testbeam campaign at DESY, $\theta = 90^\circ$	64
4.23	Hit map of a preproduction module after calibration during the testbeam campaign at DESY, $\theta = 0^\circ$	65
5.1	Sketch of the production and QC workflow.	67
5.2	Measured deviations of the specified length $l_{z,tile}$ and width $l_{\varphi,tile}$ of 16 scintillator tiles, produced as a 4×4 matrix.	69
5.3	Measured deviations of the tile height of 16 scintillator tiles, produced as a 4×4 matrix.	70
5.4	Measured deviations of the tile height of 4×28 scintillator tiles, produced as 4×4 matrix pieces.	71
5.5	Measured deviations of the specified length $l_{z,tile}$ and width $l_{\varphi,tile}$ of 16 scintillator tiles, produced as four 2×2 matrix pieces.	72
5.6	Measured deviations of the tile height of 16 scintillator tiles, produced as four 2×2 matrix pieces.	72
5.7	Sketch and photo of the first version of the SiPM QC measurement setup.	73
5.8	SiPM QC results from matrix ID = 30.	74
5.9	Summary of the QC results from all 62 SiPM matrices of the preproduction modules.	75
5.10	Comparison of the matrices produced from the two groups of SiPMs.	76
5.11	Comparison of the noise rate QC measurements from matrices 6, 10 and 45.	77
5.12	Comparison of the noise rate QC measurements from matrix ID = 30 and matrix ID = 50.	78
5.13	Gain calibration of the KLauS ASIC.	80
5.14	Result of the gain calibration.	81
5.15	Interbranch calibration concept.	82

5.16	Interbranch calibration of the ULG and LG configurations.	83
5.17	Results of the interbranch calibration.	84
5.18	Interbranch calibration of the LG and MG configurations.	85
5.19	Final result of the interbranch calibration.	86
5.20	Setup for the light yield measurement.	87
5.21	Exemplary measurement of the ^{22}Na spectrum.	88
5.22	Determination of the Compton edge position for the example of matrix ID = 02.	89
5.23	Results of the light yield measurement of matrix ID = 02.	90
5.24	SiPM QC measurements performed with the new matrix QC system. . .	91
5.25	E threshold calibration of one channel of a Tile Detector module.	94
6.1	Test setup for the evaluation of the Tile Detector cooling circuit concept.	98
6.2	Temperatures measured with the cooling test setup.	99
6.3	Implementation of the test setup in the thermal simulation.	100
6.4	Example output of the thermal simulation of $P_{\text{STiC}}^3 = 2.1 \text{ W}$	101
6.5	Comparison of the simulated and measured temperatures on the STiC package surface as a function of the ASIC power.	102
6.6	Comparison of the cooling circuit design of the technical prototype and the thinned down support plate.	104
6.7	Simulation results of configuration I.	105
6.8	Simulation results of configuration II.	106
6.9	Simulation results of configuration III.	108
6.10	Comparison of the MuTRiG maximum temperatures depending on their position on the support plate.	109
6.11	Final Tile Detector cooling circuit.	111
6.12	Parametric simulation study of the final Tile Detector cooling circuit. . .	112
6.13	Simulation results of the final Tile Detector cooling circuit.	112
6.14	Test system for the liquid detector cooling.	113
A.1	Sketch of the Tile Detector dimensions.	123
B.1	Single-photon spectra of the matrix used for the KLauS gain calibration.	126
B.2	Fast Fourier transform spectra of the SPS shown in Figure B.1	127
B.3	Ultra-low gain configuration pedestal measurements.	128
B.4	Low gain configuration pedestal measurements.	129
B.5	Medium gain configuration pedestal measurements.	130
B.6	Gain selection spectra for the ULG configuration (ULG-LG interbranch calibration)	131
B.7	Gain selection efficiency for the ULG configuration (ULG-LG interbranch calibration)	132
B.8	Gain selection spectra for the LG configuration (ULG-LG interbranch calibration)	133

B.9	Gain selection efficiency for the LG configuration (ULG-LG interbranch calibration)	134
B.10	Gain selection spectra for the LG configuration (LG-MG interbranch calibration)	135
B.11	Gain selection efficiency for the LG configuration (LG-MG interbranch calibration)	136
B.12	Gain selection spectra for the MG configuration (LG-MG interbranch calibration)	137
B.13	Gain selection efficiency for the MG configuration (LG-MG interbranch calibration)	138
B.14	Determination of the Compton edge position for the all channels of matrix ID = 02.	139
B.15	Distribution of the light yield results of matrix ID = 02.	140

C.2 List of Tables

2.1	Definition of the lepton flavour.	10
2.2	Key Standard Model parameters of the muon.	12
3.1	Properties of the EJ-228 plastic scintillator.	29
3.2	Properties of the MPPC S13360-3050VE	31
6.1	Simulation input obtained from measurements using the cooling test setup.	99
6.2	Sequence of intermediate simulations leading from the technical prototype to the final Tile Detector cooling.	103
A.1	List of Tile Detector dimensions, with explanations of the individual sizes.	122

Acronyms

ADC	analog-to-digital converter
APD	avalanche photodiode
ASIC	application-specific integrated circuit
BGA	ball grid array
CAD	computer-aided design
CL	confidence level
cLFV	charged lepton flavour violation
CMBL	compact muon beam line
CTR	coincidence time resolution
DAB	detector adapter board
DAC	digital-to-analog converter
DAQ	data acquisition
DCR	dark count rate
DESY	Deutsches Elektron-Synchrotron
ESR	enhanced specular reflector
FEB	front-end board
FFT	fast Fourier transform
FPC	flexible printed circuit
FPGA	field programmable gate array
G-APD	Geiger avalanche photodiode
GPU	graphics processing unit

GS gain selection

HDI high-density interconnect

HG high gain

HIMB High-Intensity Muon Beams

HIPA High Intensity Proton Accelerator

HV high voltage

HV-MAPS high-voltage monolithic active pixel sensor

I2C inter-integrated circuit

KIP Kirchhoff Institute for Physics

KLauS Kanäle für die Ladungsauslese von Silizium-Photomultipliern

LG low gain

LHC Large Hadron Collider

LV low voltage

MEG Mu to Electron Gamma

MEG II Mu to Electron Gamma II

MG medium gain

MuTRiG Muon-Timing Resolver including Gigabit link

PCB printed circuit board

PEEK polyether ether ketone

PSI Paul Scherrer Institut

PVT polyvinyl toluene

QC quality control

QCD Quantum Chromodynamics

QDC charge-to-digital converter

QED Quantum Electrodynamics

QFT Quantum Field Theory

SciFi Scintillating Fibre

SES single event sensitivity

SiPM silicon photomultiplier

SM Standard Model

SPAD single photon avalanche diode

SPS single-photon spectrum

SSW service support wheel

STiC Silicon Photomultiplier Timing Chip

SUSY supersymmetry

TDC time-to-digital converter

TMB Tile Module Board

ToT time-over-threshold

ULG ultra-low gain

WIMP weakly interacting massive particle

Bibliography

- [1] Y. Fukuda et al., Super-Kamiokande Collaboration. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562).
- [2] SNO Collaboration. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 89 (2002), p. 011301. DOI: [10.1103/PhysRevLett.89.011301](https://doi.org/10.1103/PhysRevLett.89.011301).
- [3] K. Arndt et al., Mu3e Collaboration. “Technical design of the phase I Mu3e experiment”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1014 (2021), p. 165679. DOI: [10.1016/j.nima.2021.165679](https://doi.org/10.1016/j.nima.2021.165679).
- [4] P. Eckert. “The Mu3e Tile Detector”. PhD thesis. Heidelberg University, 2015.
- [5] H. Klingenmeyer et al. “Measurements with the technical prototype for the Mu3e tile detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 958 (2020), p. 162852.
- [6] M. Thomson. *Modern Particle Physics*. Cambridge: Cambridge University Press, 2013.
- [7] D. Griffiths. *Introduction to Elementary Particles*. Second, Revised Edition. Wiley-VCH, 2008.
- [8] D. Galbraith and C. Burgard. *Standard Model of Physics*. <http://davidgalbraith.org/portfolio/ux-standard-model-of-the-standard-model/>. Accessed: 2024-03-04.
- [9] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Physics Letters B* 716.1 (2012), pp. 1–29. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020).
- [10] CMS Collaboration. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Physics Letters B* 716.1 (2012), pp. 30–61. DOI: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021).

- [11] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Phys. Rev. Lett.* 13 (1964), pp. 321–323. DOI: [10.1103/PhysRevLett.13.321](https://doi.org/10.1103/PhysRevLett.13.321).
- [12] P. W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Phys. Rev. Lett.* 13 (1964), pp. 508–509. DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508).
- [13] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. “Global Conservation Laws and Massless Particles”. In: *Phys. Rev. Lett.* 13 (1964), pp. 585–587. DOI: [10.1103/PhysRevLett.13.585](https://doi.org/10.1103/PhysRevLett.13.585).
- [14] N. Aghanim et al., Planck Collaboration. “Planck 2018 results - VI. Cosmological parameters”. In: *A&A* 641 (2020). DOI: [10.1051/0004-6361/201833910](https://doi.org/10.1051/0004-6361/201833910).
- [15] The KATRIN Collaboration. “Direct neutrino-mass measurement with sub-electronvolt sensitivity”. In: *Nature Physics* 18 (2022), pp. 160–166. DOI: [10.1038/s41567-021-01463-1](https://doi.org/10.1038/s41567-021-01463-1).
- [16] A. de Gouvêa and P. Vogel. “Lepton flavor and number conservation, and physics beyond the standard model”. In: *Progress in Particle and Nuclear Physics* 71 (2013), pp. 75–92. DOI: <https://doi.org/10.1016/j.pnpnp.2013.03.006>.
- [17] E. P. Hincks and B. Pontecorvo. “Search for Gamma-Radiation in the 2.2-Microsecond Meson Decay Process”. In: *Phys. Rev.* 73 (3 1948), pp. 257–258. DOI: [10.1103/PhysRev.73.257](https://doi.org/10.1103/PhysRev.73.257).
- [18] R. D. Sard and E. J. Althaus. “A Search for Delayed Photons from Stopped Sea Level Cosmic-Ray Mesons”. In: *Phys. Rev.* 74 (10 1948), pp. 1364–1371. DOI: [10.1103/PhysRev.74.1364](https://doi.org/10.1103/PhysRev.74.1364).
- [19] S. T. Petcov. “The Processes $\mu \rightarrow e + \gamma, \mu \rightarrow e + \bar{e}, \nu' \rightarrow \nu + \gamma$ in the Weinberg-Salam Model with Neutrino Mixing”. In: *Sov. J. Nucl. Phys.* 25 (1977). [Erratum: *Sov.J.Nucl.Phys.* 25, 698 (1977), Erratum: *Yad.Fiz.* 25, 1336 (1977)], p. 340.
- [20] G. Hernández-Tomé, G. López Castro, and P. Roig. “Flavor violating leptonic decays of τ and μ leptons in the Standard Model with massive neutrinos”. In: *Eur. Phys. J. C* 79.1 (2019). [Erratum: *Eur.Phys.J.C* 80, 438 (2020)], p. 84. DOI: [10.1140/epjc/s10052-019-6563-4](https://doi.org/10.1140/epjc/s10052-019-6563-4).
- [21] R. H. Bernstein and P. S. Cooper. “Charged Lepton Flavor Violation: An Experimenter’s Guide”. In: *Phys. Rept.* 532 (2013), pp. 27–64. DOI: [10.1016/j.physrep.2013.07.002](https://doi.org/10.1016/j.physrep.2013.07.002).
- [22] J. P. Ellis. “TikZ-Feynman: Feynman diagrams with TikZ”. In: *Computer Physics Communications* 210 (2017), pp. 103–123. DOI: [10.1016/j.cpc.2016.08.019](https://doi.org/10.1016/j.cpc.2016.08.019).

- [23] F. Renga. “Experimental searches for muon decays beyond the Standard Model”. In: *Reviews in Physics* 4 (2019), p. 100029. DOI: <https://doi.org/10.1016/j.revip.2019.100029>.
- [24] F. Cei. “Status and perspectives of Lepton Flavour Violation experiments with muons”. In: *J. Phys. Conf. Ser.* 1526 (2020), p. 012020. DOI: [10.1088/1742-6596/1526/1/012020](https://doi.org/10.1088/1742-6596/1526/1/012020).
- [25] Particle Data Group et al. “Review of Particle Physics”. In: *Progress of Theoretical and Experimental Physics* 2022.8 (2022), p. 083C01. DOI: [10.1093/ptep/ptac097](https://doi.org/10.1093/ptep/ptac097).
- [26] A. M. Baldini et al., MEG Collaboration. “Search for the lepton flavour violating decay $\mu^+ \rightarrow e^+\gamma$ with the full dataset of the MEG experiment”. In: *The European Physical Journal C* 76 (2016), p. 434. DOI: [10.1140/epjc/s10052-016-4271-x](https://doi.org/10.1140/epjc/s10052-016-4271-x).
- [27] SINDRUM Collaboration. “Search for the decay $\mu^+ \rightarrow e^+e^+e^-$ ”. In: *Nuclear Physics B* 299.1 (1988), pp. 1–6. DOI: [10.1016/0550-3213\(88\)90462-2](https://doi.org/10.1016/0550-3213(88)90462-2).
- [28] SINDRUM II Collaboration. “A search for $\mu-e$ conversion in muonic gold”. In: *Eur. Phys. J. C* 47 (2006), pp. 337–346. DOI: [10.1140/epjc/s2006-02582-x](https://doi.org/10.1140/epjc/s2006-02582-x).
- [29] W. J. Marciano, T. Mori, and J. M. Roney. “Charged Lepton Flavor Violation Experiments”. In: *Annual Review of Nuclear and Particle Science* 58 (2008), pp. 315–341. DOI: [10.1146/annurev.nucl.58.110707.171126](https://doi.org/10.1146/annurev.nucl.58.110707.171126).
- [30] K. Afanaciev et al., MEG II collaboration. “A search for $\mu^+ \rightarrow e^+\gamma$ with the first dataset of the MEG II experiment”. In: *arXiv preprint arXiv:2310.12614* (2023). DOI: [10.48550/arXiv.2310.12614](https://doi.org/10.48550/arXiv.2310.12614).
- [31] S. Middleton, M. Lee, and Y. Seiya. “Experimental Searches for Muon to Electron Conversion in a Nucleus: COMET, DeeMe, and Mu2e. A Contributed paper for Snowmass 21”. In: *arXiv preprint arXiv:2203.07089* (2022). DOI: [10.48550/arXiv.2203.07089](https://doi.org/10.48550/arXiv.2203.07089).
- [32] Mu2e Collaboration. “Mu2e technical design report”. In: *arXiv preprint arXiv:1501.05241* (2015). DOI: [10.48550/arXiv.1501.05241](https://doi.org/10.48550/arXiv.1501.05241).
- [33] COMET collaboration. “COMET Phase-I technical design report”. In: *Progress of Theoretical and Experimental Physics* 2020.3 (2020), p. 033C01. DOI: [10.1093/ptep/ptz125](https://doi.org/10.1093/ptep/ptz125).
- [34] A. Blondel et al. *Letter of Intent for an Experiment to Search for the Decay $\mu \rightarrow eee$* . 2012.
- [35] A. Blondel et al. “Research Proposal for an Experiment to Search for the Decay $\mu \rightarrow eee$ ”. In: *arXiv preprint arXiv:1301.6113* (2013). DOI: [10.48550/arXiv.1301.6113](https://doi.org/10.48550/arXiv.1301.6113).

- [36] Y. Kuno and Y. Okada. “Muon decay and physics beyond the standard model”. In: *Rev. Mod. Phys.* 73 (1 2001), pp. 151–202. DOI: [10.1103/RevModPhys.73.151](https://doi.org/10.1103/RevModPhys.73.151).
- [37] S. Corrodi. “A Timing Detector based on Scintillating Fibres for the Mu3e Experiment: Prototyping, Simulation and Integration”. PhD thesis. ETH Zurich, 2018.
- [38] M. Mühlwinkel. “Kinematic and Timing Studies on the Mu3e Detector”. Bachelor’s thesis. Heidelberg University, 2022.
- [39] J. Grillenberger, C. Baumgarten, and M. Seidel. “The High Intensity Proton Accelerator Facility”. In: *SciPost Phys. Proc.* (2021), p. 002. DOI: [10.21468/SciPostPhysProc.5.002](https://doi.org/10.21468/SciPostPhysProc.5.002).
- [40] D. Kiselev et al. “The Meson Production Targets in the high energy beamline of HIPA at PSI”. In: *SciPost Phys. Proc.* (2021), p. 003. DOI: [10.21468/SciPostPhysProc.5.003](https://doi.org/10.21468/SciPostPhysProc.5.003).
- [41] F. Berg. “CMBL-A High-Intensity Muon Beam Line & Scintillation Target with Monitoring System for Next-Generation Charged Lepton Flavour Violation Experiments”. PhD thesis. ETH Zurich, 2017.
- [42] M. Aiba et al. “Science Case for the new High-Intensity Muon Beams HIMB at PSI”. In: *arXiv preprint arXiv:2111.05788* (2021). DOI: [10.48550/arXiv.2111.05788](https://doi.org/10.48550/arXiv.2111.05788).
- [43] G. Dal Maso et al. “Future facilities at PSI, the High-Intensity Muon Beams (HIMB) project”. In: *EPJ Web Conf.* 282 (2023), p. 01012. DOI: [10.1051/epjconf/202328201012](https://doi.org/10.1051/epjconf/202328201012).
- [44] A. Weber. “Development of integrated circuits and smart sensors for particle detection in physics experiments and particle therapy”. PhD thesis. Heidelberg University, 2021.
- [45] H. Augustin. “Development of a novel slow control interface and suppression of signal line crosstalk enabling HV-MAPS as sensor technology for Mu3e”. PhD thesis. Heidelberg University, 2021.
- [46] H. Augustin et al. “Mupix10: first results from the final design”. In: *Proceedings of the 29th International Workshop on Vertex Detectors (VERTEX2020)*. 2021, p. 010012.
- [47] T. Rudzki et al. “An ultra-light helium cooled pixel detector for the Mu3e experiment”. In: *Journal of Instrumentation* 18.10 (2023), p. C10022. DOI: [10.1088/1748-0221/18/10/C10022](https://doi.org/10.1088/1748-0221/18/10/C10022).
- [48] T. Rudzki. “The Mu3e vertex detector - construction, cooling, and first prototype operation”. PhD thesis. Heidelberg University, 2022.

- [49] A. Bravar et al. “Development of the Scintillating Fiber Timing Detector for the Mu3e Experiment”. In: *arXiv preprint arXiv:2208.09906* (2022). DOI: [10.48550/arXiv.2208.09906](https://doi.org/10.48550/arXiv.2208.09906).
- [50] L. Gerritzen. “The Mu3e Scintillating Fibre Detector: Light Attenuation, Radiation Damage, and Data Acquisition”. PhD thesis. ETH Zurich, 2022.
- [51] T. Zhong. “Irradiation Studies for the Mu3e Tile Detector”. PhD thesis. Heidelberg University, 2023.
- [52] H. Augustin et al. “The Mu3e Data Acquisition”. In: *IEEE Transactions on Nuclear Science* 68.8 (2021), pp. 1833–1840. DOI: [10.1109/TNS.2021.3084060](https://doi.org/10.1109/TNS.2021.3084060).
- [53] N. Berger et al. “A new three-dimensional track fit with multiple scattering”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 844 (2017), pp. 135–140. DOI: [10.1016/j.nima.2016.11.012](https://doi.org/10.1016/j.nima.2016.11.012).
- [54] H. Kolanoski and N. Wermes. *Particle Detectors: Fundamentals and Applications*. Oxford University Press, 2020. DOI: [10.1093/oso/9780198858362.001.0001](https://doi.org/10.1093/oso/9780198858362.001.0001).
- [55] Eljen Technology. *EJ-228, EJ-230*. <https://eljentechnology.com/products/plastic-scintillators/ej-228-ej-230>. Accessed: 2024-03-04.
- [56] Hamamatsu Photonics K.K. *S13360-2050VE/-3050VE/-6050VE*. https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/ssd/s13360-2050ve_etc_kapd1053e.pdf. Accessed: 2024-03-04.
- [57] Patrick Eckert et al. “Characterisation studies of silicon photomultipliers”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 620.2 (2010), pp. 217–226. DOI: <https://doi.org/10.1016/j.nima.2010.03.169>.
- [58] H. Chen. “A Silicon Photomultiplier Readout ASIC for the Mu3e Experiment”. PhD thesis. Heidelberg University, 2018.
- [59] T. Harion. “The STiC ASIC: High Precision Timing with Silicon Photomultipliers”. PhD thesis. Heidelberg University, 2015.
- [60] W. Shen. “Development of High Performance Readout ASICs for Silicon Photomultipliers (SiPMs)”. PhD thesis. Heidelberg University, 2012.
- [61] T. Hartwig. “Messung der Zeitaufösung eines Szintillator-SiPM-Systems für den Kacheldetektor des Mu3e-Experiments”. Bachelor’s thesis. Heidelberg University, 2013.

- [62] C. Licciulli. “Präzise Zeitmessung für das Mu3e-Experiment”. Master’s thesis. Heidelberg University, 2013.
- [63] C. Graf. “A Calibration Scheme for the Mu3e Tile Detector”. Master’s thesis. Heidelberg University, 2015.
- [64] J. Urban. “Time Resolution Studies for the Mu3e Experiment”. Bachelor’s thesis. Heidelberg University, 2015.
- [65] L. Bereska. “Optical Crosstalk in the Mu3e Tile Detector”. Bachelor’s thesis. Heidelberg University, 2015.
- [66] R. Diener et al. “The DESY II test beam facility”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 922 (2019), pp. 265–286. DOI: [10.1016/j.nima.2018.11.133](https://doi.org/10.1016/j.nima.2018.11.133).
- [67] Witte Barskamp GmbH & Co. KG. *Gefrierspanntechnik*. <https://www.witte-barskamp.de/vakuum-spanntechnik/gefrier-spanntechnik/>. Accessed: 2024-03-04.
- [68] J. Küpperbusch. “Module-Wide Synchronization Studies on the Mu3e Tile Detector”. Bachelor’s thesis. Heidelberg University, 2023.
- [69] A. Junkermann. “Calibration of the Mu3e Tile Detector”. Master’s thesis. Heidelberg University, 2022.
- [70] L. Lauer. “Slow Control and Monitoring of the Mu3e Tile Detector”. Bachelor’s thesis. Heidelberg University, 2023.
- [71] Deutsches Elektronen-Synchrotron DESY. *Generation of the DESY Test Beams*. https://particle-physics.desy.de/test_beams_at_desy/e252106/e252211. Accessed: 2024-03-04.
- [72] K. Briggel. “Silicon Photomultiplier Readout Electronics for Imaging Calorimetry Applications”. PhD thesis. Heidelberg University, 2018.
- [73] K. Briggel et al. “KLauS: a low power Silicon Photomultiplier charge readout ASIC in 0.18 UMC CMOS”. In: *Journal of Instrumentation* 11.03 (2016), p. C03045. DOI: [10.1088/1748-0221/11/03/C03045](https://doi.org/10.1088/1748-0221/11/03/C03045).
- [74] The Raspberry Pi Foundation. *Raspberry Pi 4 Model B*. <https://www.raspberrypi.com/products/raspberry-pi-4-model-b/>. Accessed: 2024-03-04.
- [75] F. G. Kondev et al. “The NUBASE2020 evaluation of nuclear physics properties”. In: *Chinese Physics C* 45.3 (2021), p. 030001. DOI: [10.1088/1674-1137/abddae](https://doi.org/10.1088/1674-1137/abddae).

- [76] LD DIDACTIC. γ -spectrum of Na-22. <https://www.ld-didactic.de/software/524221en/Content/Appendix/Na22.htm>. Accessed: 2024-03-04.
- [77] E. Auffray et al. “Characterization studies of silicon photomultipliers and crystals matrices for a novel time of flight PET detector”. In: *Journal of Instrumentation* 10.06 (2015), P06009. DOI: [10.1088/1748-0221/10/06/P06009](https://doi.org/10.1088/1748-0221/10/06/P06009).
- [78] A. Junkermann. “Detector Parametrisation Studies of the Mu3e Tile Detector”. Bachelor’s thesis. Heidelberg University, 2020.
- [79] W. Demtröder. *Mechanics and Thermodynamics*. Springer Cham, 2017. DOI: [10.1007/978-3-319-27877-3](https://doi.org/10.1007/978-3-319-27877-3).

Acknowledgements

At last, it is time to thank all the people without whom this thesis would not have been possible. Off we go:

To Hans-Christian Schultz-Coulon, for not only agreeing to be my supervisor, but for also supporting me throughout the past six years and encouraging me to never give up. You have my sincerest gratitude for everything you have done for this thesis to succeed.

To Loredana Gastaldo, for agreeing to be my second thesis referee, and for always showing genuine interest in my work and having an open door and a kind word for me.

To my beloved colleagues of the Tile Detector team, who spent most of the past six years with me and survived various testbeams (although it was a close call sometimes...): Yonathan, for being my mentor and for always believing in me, especially when I could not; Konrad, for still going on testbeams with me and for giving everything to help my thesis succeed; and Tiancheng, for not only sharing an office with me for all those years, but also your culture, your food, and your kindness.

To all my colleagues in the Mu3e collaboration, for all the discussions, the workshops, the kindness, and the wonderful in-person meetings (even though I don't really ski). It was a pleasure working with you all.

To the people behind the scenes: those who have read my thesis will have repeatedly stumbled on “the mechanical workshop” or “the electronics workshop”... Here they are in person: Werner Lamadé, David Jansen, Julia Bing, Christian Herdt, Knut Azeroth, Jens Kohlgrüber, and Alexander Dobler. Your expertise, and your patience, has been invaluable to me. A special thanks goes out to Vera Stankova, for her unwavering support and her big, kind heart.

To Petra Pfeifer, for knowing everything, and if she doesn't, for knowing who to call. You make all of our lives so much easier, and you still always spare some time for a coffee with me.

To the wonderful groups F8 and F11, for putting the fun in fundamental physics (and the Ouzo in the fridge).

To my small army of proofreaders, for reading, and reading, and then reading some more: Rainer, Konrad, Yonathan, Martin, Falk, Thomas J., Sebastian, Lisa, Jan, Ralf, and Thomas R., who has now spent more than ten years with me in lectures, tutorials, ALICE and Mu3e, and has made it all so much more cheerful.

To my family and my in-laws, for picking me up when I'm down, and reminding me that life outside of physics exists, and is actually fun.

And lastly, to Ralf. For everything.