Development of a spatially resolving detector for ultracold neutrons

Dissertation zur Erlangung des Grades
"Doktor der Naturwissenschaften"

am Fachbereich Physik, Mathematik und Informatik
der Johannes Gutenberg-Universität
in Mainz

Konrad Guntermann geb. Franz
geboren in Frankfurt am Main

Mainz, den 18.06.2026

Eigenständigkeitserklärung

Hiermit erkläre ich, dass ich die vorliegende Arbeit "Development of a spatially resolving detector for ultracold neutrons" selbstständig verfasst, keine anderen als die angegebenen Quellen und Hilfsmittel benutz und von der Ordnung zur Sicherung guter wissenschaftlicher Praxis in Forschung und Lehre und zum Verfahren zum Umgang mit wissenschaftlichem Fehlverhalten Kenntnis genommen habe. Sämtliche wörtlichen oder sinngemäßen Übernahmen und Zitate sind kenntlich gemacht und nachgewiesen. Im Anhang H habe ich die verwendeten KI-Tools dokumentiert. Auf Kollaborationen mit Kolleg:innen habe ich an den entsprechenden Stellen hingewiesen. Mit der Abgabe der vorliegenden Leistung übernehme ich die Verantwortung für das eingereichte Gesamtprodukt. Ich verantworte damit auch jegliche KI-generierten Inhalte, die ich in meine Arbeit übernommen habe. Die Richtigkeit übernommener Aussagen und Inhalte habe ich nach bestem Wissen und Gewissen geprüft. Ich habe die Arbeit nicht zum Erwerb eines anderen Leistungsnachweises in gleicher oder ähnlicher Form eingereicht. Ich habe die Arbeit oder Teile davon auch nicht an anderer Stelle als Dissertation eingereicht. Mir ist bekannt, dass ein Verstoß gegen die genannten Punkte prüfungsrechtliche Konsequenzen hat.

K. Guntermann

Mainz, den 18.06.2026

1. Gutachter: Dr. Dieter Ries
2. Gutachter: Prof. Dr. Martin Fertl
Datum der mündlichen Prüfung: 17.08.2026

Abstract

Precision measurements of the free neutron β-decay parameters enable tests of the unitarity of the Cabibbo-Kobayashi-Maskawa (CKM) matrix and constrain Big Bang nucleosynthesis (BBN) models through the neutron lifetime τ_n. Current measurements of τ_n obtained using beam and bottle methods differ by 4.6σ, which is commonly referred to as the "neutron lifetime puzzle". Improved neutron lifetime measurements with high precision and improved control of systematic effects are therefore required. Among these efforts is the τSPECT experiment, which traps ultracold neutrons (UCNs) in a fully magnetic trap. In addition, the trap design is scalable to a deeper trapping potential, which could enable a determination of the neutron lifetime beyond the current level of precision. To achieve this goal, systematic effects associated with the trap design must be studied in great detail. As part of this thesis, two such effects were investigated using experimental data. The first effect is caused by so-called marginally trapped neutrons, which possess sufficient energy to escape the trap but remain confined for extended periods due to their specific trajectories. The second effect is related to the evolution of the phase space of the trapped ensemble, causing the mean detection time to differ after different storage times. While the investigation of the effect related to marginally trapped neutrons ultimately remained inconclusive due to limited statistical precision, it provided valuable guidance for future studies. Phase space related effects were observed for the first time in this study and a correction procedure was successfully established. To improve the understanding of both systematic effects and to provide additional information about the trapped ensemble, a spatially resolving UCN detector was developed and characterized in the second part of the thesis. The detector consists of thin conversion and scintillation layers, converting incoming UCNs into light pulses, a light guide and an array of silicon photomultipliers (SiPMs). Spatial resolution is achieved by distributing the scintillation light among the SiPMs. As part of the development, the SiPM spacing and the light guide geometry were optimized by modeling the light pulses as point like light sources. During the characterization, three approaches for reconstructing the event positions were investigated. However, all approaches were found to exhibit a reconstruction bias. This was attributed to the point light source model apparently not adequately representing the true light pulses and to the low average number of active SiPM channels per event. Using the approach with the smallest reconstruction bias, the spatial resolution was estimated through dedicated test measurements, yielding a resolution of 2.3(2) mm. While this confirmed that sub-pixel resolution was achieved, the previously estimated target resolution of approximately 1 mm was not achieved with the first prototype. In addition to the spatial resolution, parameters such as the rate capability, γ-sensitivity and intrinsic background rate were investigated and the prototype was compared to both the current τSPECT detector and a commercial detector. Together, the presented studies provide insight into systematic effects in magnetic UCN lifetime experiments, demonstrate the viability of the detector concept and provide recommendations for design improvements.

Zusammenfassung

Präzisionsmessungen der Parameter des β-Zerfalls freier Neutronen ermöglichen es, die Unitarität der Cabibbo-Kobayashi-Maskawa (CKM) Matrix zu überprüfen und liefern über die Neutronenlebensdauer τ_n Randbedingungen für Modelle der primordialen Nukleosynthese. Derzeit gibt es eine 4.6σ Abweichung zwischen Messungen von τ_n, die mittels Strahl- und Flaschenmethode durchgeführt wurden, was als das "Neutronenlebensdauerrätsel" bekannt ist. Um Neutronenlebensdauermessungen weiter zu verbessern, sind weitere hochpräzise Experimente mit verbesserter Kontrolle systematischer Effekte nötig. Teil dieser Bemühungen ist das τSPECT Experiment, bei dem ultrakalte Neutronen (UCNs) in einer vollmagnetischen Falle eingefangen werden. Außerdem kann das Fallendesign hin zu einem tieferen Fallenpotential skaliert werden, was eine Bestimmung der Neutronenlebensdauer jenseits des derzeitigen Präzisionsniveaus ermöglichen könnte. Um dieses Ziel zu erreichen, müssen mit dem Fallendesign zusammenhängende systematische Effekte im Detail untersucht werden. Im Rahmen dieser Arbeit wurden zwei solcher Effekte anhand experimenteller Daten untersucht. Der erste Effekt wird durch sogenannte marginal gespeicherte Neutronen verursacht, die über ausreichend Energie verfügen, um der Falle zu entkommen, aber aufgrund ihrer Trajektorien für längere Zeiträume in der Falle verbleiben. Der zweite Effekt wird durch Veränderungen im Phasenraum des gespeicherten Ensembles verursacht, wodurch sich die mittleren Detektionszeiten nach verschiedenen Speicherzeiten unterscheiden. Auch wenn letztlich die Untersuchungen des durch marginal gespeicherte Neutronen verursachten Effekts aufgrund der eingeschränkten statistischen Sensitivität uneindeutig blieben, lieferten sie wertvolle Anhaltspunkte für zukünftige Studien. Durch Phasenraumveränderungen verursachte Effekte wurden als Teil dieser Studie erstmals beobachtet und eine Korrekturmethode für diese konnte erfolgreich etabliert werden. Um das Verständnis beider systematischer Effekte zu verbessern und weitere Informationen über das gespeicherte Ensemble zu erhalten, wurde im zweiten Teil der Arbeit ein ortsauflösender UCN-Detektor entwickelt und charakterisiert. Der Detektor besteht aus dünnen Konversions- und Szintillationsschichten, die einfallende UCNs in Lichtpulse umwandeln, einem Lichtleiter und einem Gitter aus Silizium-Photomultipliern (SiPMs). Die Ortsauflösung wird dadurch erreicht, dass das Szintillationslicht auf verschiedene SiPMs verteilt wird. Im Rahmen der Entwicklung wurden die Lichtleitergeometrie und der Abstand zwischen den SiPMs optimiert, indem Lichtpulse als Punktlichtquellen modelliert wurden. Während der Charakteriserung wurden drei verschiedene Ansätze zur Rekonstruktion der Eventpositionen untersucht. Allerdings zeigte sich, dass alle Ansätze Rekonstruktionsverzerrungen aufwiesen. Dies wurde darauf zurückgeführt, dass das Punktlichtquellenmodell offenbar keine gute Beschreibung der tatsächlichen Lichtpulse darstellt und dass die mittlere Anzahl aktiver SiPM-Kanäle pro Event gering war. Mithilfe des Ansatzes mit der geringsten Rekonstruktionsverzerrung und dedizierter Testmessungen wurde die Ortsauflösung des Detektors abgeschätzt, woraus sich eine Auflösung von 2.3(2) mm ergab. Dies bestätigte zwar, dass eine Auflösung unterhalb der Pixelgröße erreicht wurde, jedoch erreichte der erste Prototyp damit nicht das zuvor abgeschätzte Ziel einer Auflösung von etwa 1 mm. Neben der Ortsauflösung wurden Parameter wie Ratenkapazität, γ-Sensitivität und intrinsische Untergrundrate untersucht und der Prototyp sowohl mit dem aktuellen τSPECT Detektor als auch mit einem kommerziellen Detektor verglichen. Insgesamt liefern die präsentierten Studien Einblicke in die systematischen Effekte magnetischer UCN-Lebensdauerexperimente, zeigten, dass das Detektorkonzept umsetzbar ist und geben Empfehlungen für Verbesserungen des Designs.

Acknowledgements

Diese Seite wurde aus Datenschutzgründen entfernt.

Diese Seite wurde aus Datenschutzgründen entfernt.

Contents

List of Figures

List of Tables

List of Abbreviations

ADC	**A**nalog-to-**D**igital **C**onverter
AFP	**A**diabatic **F**ast **P**assage
ASIC	**A**pplication-**S**pecific **I**ntegrated **C**ircuit
BBN	**B**ig **B**ang **N**ucleosynthesis
CCD	**C**harge **C**oupled **D**evice
CDF	**C**umulative **D**istribution **F**unction
CF flange	**C**on**F**lat flange
CMOS	**C**omplementary **M**etal-**O**xide-**S**emiconductor
CKM matrix	**C**abbibo-**K**obayashi-**M**askawa matrix
DAC	**D**igital-to-**A**nalog **C**onverter
DAQ	**D**ata **A**c**Q**uisition
DC	**D**irect **C**urrent
EDM	**E**lectric **D**ipole **M**oment
FPGA	**F**ield **P**rogrammable **G**ate **A**rray
GPS	**G**lobal **P**ositioning **S**ystem
GPIO	**G**eneral **P**urpose **I**nput/**O**utput
HFS	**H**igh **F**ield **S**eeking
HIPA	**H**igh **I**ntensity **P**roton **A**ccelerator
ISO	**I**nternational **O**rganization for **S**tandardization
IV-curve	Current (**I**)-**V**oltage-curve
JGU	**J**ohannes **G**utenberg **U**niversität
J-PARC	**J**apan-**P**roton **A**ccelerator **R**esearch **C**omplex
LFS	**L**ow **F**ield **S**eeking
NIST	**N**ational **I**nstitute of **S**tandards and **T**echnology
NLK	**N**eutronen**L**eiter**K**lappe
NMR	**N**uclear **M**agnetic **R**esonance
NTP	**N**etwork **T**ime **P**rotocol
OS	**O**perating **S**ystem
PCB	**P**rinted **C**ircuit **B**oard
PCG	**P**ermuted **C**ongruential **G**enerator
PDG	**P**article **D**ata **G**roup
PI controller	**P**roportional-**I**ntegral controller
PSI	**P**aul **S**cherrer **I**nstitute
PWM	**P**ulse **W**idth **M**odulation
QDC	Charge (**Q**)-to-**D**igital **C**onverter
RF	**R**adio **F**requency
RPC	**R**emote **P**rocedure **C**all
SF	**S**pin-**F**lipping unit
SiPM	**Si**licon **P**hoto**M**ultiplier
SPAD	**S**ingle **P**hoton **A**valanche **D**iode
SRIM	The **S**topping and **R**ange of **I**ons in **M**atter
SSH	**S**ecure **SH**ell

STL	**St**ereo**L**itography
TDC	**T**ime-to-**D**igital **C**onverter
ToT	**T**ime **o**ver **T**hreshold
TPC	**T**ime **P**rojection **C**hamber
TRIGA	**T**raining, **R**esearch, **I**sotopes, **G**eneral **A**tomics
UCN	**U**ltra**C**old **N**eutrons

Chapter 1

Introduction

In 1930, Walther Bothe and Herbert Becker performed experiments in which they irradiated different chemical elements with α-radiation. In some cases they observed radiation with unusually high penetrating power, which they interpreted as high-energy γ-radiation [BB30]. Two years later, Irène and Frédéric Joliot-Curie found that this radiation could eject protons from hydrogen rich materials and attributed this effect to Compton scattering of γ-radiation on hydrogen nuclei [CJ32]. After repeating and refining these experiments as well as performing model calculations that challenged this interpretation, James Chadwick explained the observations by postulating a new particle, the neutron [Cha32]. For his interpretation and experiments Chadwick was awarded the Nobel Prize in Physics in 1935 "for the discovery of the neutron"[1]. The discovery of the neutron had profound implications for nuclear and particle physics, as it provided the missing component needed to explain the structure of atomic nuclei [NV17]. Furthermore, it enabled the discovery of nuclear fission [NV17], ultimately leading to the development of nuclear weapons and nuclear reactors, which continue to influence international politics to this day. Today, neutron sources are used in a wide range of applications, from neutron scattering in materials science [Boo20] to the production of isotopes for medical purposes [Qai12; Wan+22]. Beyond these applications, the fundamental properties of the neutron remain an active area of research, aimed at deepening our understanding of nuclear and particle physics [DM21]. Among the experimental efforts in this field are precision measurements of the free neutron β-decay [DM21; GS23]. These measurements contribute to tests of the unitarity of the Cabibbo-Kobayashi-Maskawa (CKM) quark mixing matrix, providing a precision low-energy test of the Standard Model of particle physics that is not limited by theory uncertainties associated with nuclear structure corrections [Wie18; DM21; Hay24]. A key parameter in this context is the free neutron lifetime. To date, the most precise measurement of this quantity has been obtained using a magneto-gravitational trap for ultracold neutrons (UCNs), in which the number of surviving neutrons is measured after different storage times to extract the lifetime [Mus+25]. One approach to further improve the precision of these measurements and thereby contribute to a more stringent test of CKM unitarity is the use of a deeper, fully magnetic trap for UCNs. In the τSPECT experiment, a trap geometry suitable for such a deep magnetic trap is currently under investigation [Aul+24; Aul+25]. Achieving higher precision, however, also requires an improved understanding of systematic effects. One such effect is the evolution of the phase space distribution between different storage times, which can lead to variations in the temporal distribution of detected neutrons and thereby bias the extracted lifetime. Although this effect has been negligible in the currently most precise measurement of the neutron lifetime [Mus+25], its impact may become

[1] The Nobel Prize in Physics 1935: `https://www.nobelprize.org/prizes/physics/1935/summary/`

significant in experiments employing different trap geometries or aiming at higher precision. Another effect is the presence of so-called marginally trapped neutrons, which are neutrons with energies above the nominal trap edge that nevertheless remain confined for extended periods due to their trajectories [Wie18; WG11]. These neutrons can act as an effective additional loss channel biasing the extracted lifetime. A detector that is compatible with high magnetic fields and cryogenic ultra high vacuum environments, while providing both spatial and temporal resolution, can significantly enhance the investigation of these effects. In combination with the magnetic field of a τSPECT-like trap, such a detector also enables the extraction of information about the energy spectrum of the trapped neutron ensemble. This capability makes it possible to monitor the energy spectrum and thereby assess whether associated systematic effects remain stable over many experimental runs or vary with changes in the energy distribution.

This thesis addresses both the investigation of marginally trapped neutrons and phase space evolution in the τSPECT experiment as well as the development and characterization of a first prototype of a spatially resolving, high magnetic field compatible detector for UCNs. Chapter 2 provides an introduction to UCNs, the motivation for neutron lifetime measurements and an overview of the current experimental status. Chapter 3 introduces the τSPECT experiment and presents an analysis of marginally trapped neutrons and phase space evolution based on experimental data. The development and characterization of the detector prototype are described in Chapter 4 and Chapter 5, respectively.

Chapter 2

Motivation and theory

2.1 Ultracold neutrons

Ultracold neutrons (UCNs) are neutrons with a kinetic energy on the order of 100 neV, corresponding to a velocity of roughly 5 m/s. Their defining property is that they can be stored in material vessels and confined by magnetic fields. At these low energies, all four fundamental interactions, namely electromagnetic, strong, weak and gravitational interaction must be considered. The relevance of each of these interactions is explained in the following sections.

2.1.1 Fermi potential

The reflection of neutrons from material walls was first reported by Fermi and Zinn in 1946 [FZ46]. This phenomenon can be described in terms of the so-called neutron optical potential, also known as the Fermi potential. It arises from the fact that the de Broglie wavelength of UCNs is on the order of 100 nm, which is significantly larger than both the size of a single nucleus and the inter-nuclear spacing in solids. The following paragraphs summarize the derivation of this potential as presented in [GRL91]. At low energies, the interaction between a neutron and a nucleus can be approximated by a square well potential with a depth on the order of 10 MeV and a spatial extent at the femtometer scale. In this regime, the total wavefunction outside the potential well is given by the sum of the incident plane wave and the scattered wave

$$\psi = \mathrm{e}^{i\vec{k}\cdot\vec{r}} + f(\theta)\frac{\mathrm{e}^{ikr}}{r}, \tag{2.1}$$

where $\vec{k}$ denotes the neutron wave vector, $f(\theta)$ the so-called scattering amplitude and $\vec{r}$ the position relative to the nucleus. The magnitude of the wave vector k is proportional to the inverse of the de Broglie wavelength. Since the de Broglie wavelength of UCNs is large compared to the size of a nucleus, only s-wave scattering (angular momentum $l = 0$) contributes significantly to the scattered wave. Because the s-wave scattering is isotropic, i.e., there is no preferred direction, $f(\theta)$ is independent of the angle and can be written as

$$f(\theta) = -b, \tag{2.2}$$

where b is a constant known as the scattering length. Because the neutron wavelength is large compared to the relevant length scales, the condition $k \ll 1$ applies, yielding

$$\psi \approx \frac{r-b}{r} \tag{2.3}$$

for the wavefunction outside the nucleus. Thus, b corresponds to the position at which ψ approximately vanishes. The scattering length can be related to the differential scattering cross section by noting that the latter is defined as the ratio of the scattered probability flux into a solid angle element $d\Omega$ to the incident flux. This yields

$$\frac{d\sigma}{d\Omega} = b^2 \tag{2.4}$$

for the differential scattering cross section. The numerical value of b depends on the boundary conditions at the edge of the nuclear potential and can be either positive or negative. Since the wavefunction inside the nucleus typically undergoes multiple oscillations, its phase at the potential edge is highly sensitive to both the depth of the nuclear potential and the size of the nucleus. Consequently, b varies significantly between different nuclei. Because nuclear theory is typically not sufficiently precise to calculate b, the scattering length is usually determined experimentally. To describe the effective interaction of low-energy neutrons with a nucleus, the actual nuclear potential can be replaced by the Fermi potential

$$V_{\mathrm{F}} = \frac{2\pi\hbar^2 b}{m_{\mathrm{n}}}\delta^{(3)}(\vec{r} - \vec{R}). \tag{2.5}$$

Here, $\vec{R}$ is the position of the nucleus and m_{n} is the neutron mass. In the first-order Born approximation, this effective potential reproduces the correct scattering behavior. Due to their long wavelength compared to typical inter-nuclear distances, UCNs interacting with a material wall effectively interact with an ensemble of nuclei rather than a single nucleus. In this case, the matrix element describing the transition from initial to final state is given by

$$f\left(\vec{k}_0 \to \vec{k}_{\mathrm{f}}\right) = \frac{2\pi\hbar^2}{m_{\mathrm{n}}}\sum_i b_i e^{i\vec{Q}\cdot\vec{R}_i}, \tag{2.6}$$

where $\vec{Q} = \vec{k}_0 - \vec{k}_{\mathrm{f}}$ denotes the wave vector transfer during scattering. This expression represents the superposition of spherical wave amplitudes from individual nuclei located at positions $\vec{R}_i$. For unpolarized neutrons, the scattering cross section must be averaged over both the isotopic composition of the ensemble of nuclei and their spin distribution. This yields

$$\frac{d\sigma}{d\Omega} = \overline{\left|f\left(\vec{k}_0 \to \vec{k}_{\mathrm{f}}\right)\right|^2} = \sum_{i,j} e^{i\vec{Q}\cdot(\vec{R}_j - \vec{R}_i)}\overline{b_i^* b_j}. \tag{2.7}$$

Assuming b_i and b_j are uncorrelated, $\overline{b_i^* b_j}$ can be written as

$$\overline{b_i^* b_j} = \begin{cases} \left|\overline{b}\right|^2 & \text{if } i \neq j \\ \overline{|b|^2} & \text{if } i = j. \end{cases} \tag{2.8}$$

Expanding the sum in Eq. (2.7) then yields

$$\frac{d\sigma}{d\Omega} = \left|\overline{b}\right|^2 \left|\sum_i e^{i\vec{Q}\cdot\vec{R}_i}\right|^2 - \left|\overline{b}\right|^2 N + \overline{|b|^2} N = b_{\mathrm{coh}}^2 \left|\sum_i e^{i\vec{Q}\cdot\vec{R}_i}\right|^2 + b_{\mathrm{inc}}^2 N, \tag{2.9}$$

where N denotes the total number of scattering nuclei and b_{coh} and b_{inc} are referred to as coherent and incoherent scattering length, respectively. If the nuclei are arranged periodically, the sum in the coherent term collapses to N^2 when $\vec{Q} \cdot \vec{R}_i$ equals an integer multiple of 2π, resulting in Bragg diffraction. The second term is proportional to the variance of the scattering lengths of the scattering nuclei and produces a diffuse background. However, because the inter-nuclear spacing is much smaller than the UCN wavelength, Bragg scattering does not occur for UCNs. The effective potential for an ensemble of nuclei is given by

$$V(\vec{r}) = \frac{2\pi\hbar^2}{m_\mathrm{n}} \sum_i b_i \delta(\vec{r} - \vec{R}_i). \tag{2.10}$$

In the continuum limit, where the UCN wavelength is much larger than the inter-nuclear spacing and only coherent forward scattering contributes, the discrete sum can be replaced by a volume-averaged potential, given by

$$V_\mathrm{F} = \frac{2\pi\hbar^2}{m_\mathrm{n}} n b_\mathrm{coh}, \tag{2.11}$$

where n denotes the number density of the ensemble of nuclei. This potential effectively acts as a potential step for a neutron incident on a material surface. If the kinetic energy of the neutron is smaller than the Fermi potential of the surface, it is reflected under all angles of incidence. Consequently, UCNs can be confined in vessels made of suitable materials and material guides can be used to transport UCNs from a source to an experimental setup. In addition, losses of UCNs at material surfaces can be incorporated within this framework by introducing an imaginary component to the coherent scattering length. The Fermi potential then becomes

$$V_\mathrm{F} = V - iW = \frac{2\pi\hbar^2}{m_\mathrm{n}} n(b_r - ib_i), \tag{2.12}$$

where W denotes the imaginary component of the Fermi potential and b_r and b_i are the real and the imaginary component of the scattering length, respectively. This approach accounts for loss processes that are irreversible or reversible with negligible probability. However, this approximation is well justified for neutron capture and inelastic upscattering related losses, which are the dominant loss mechanisms in the UCN regime. A list of calculated Fermi potentials for all materials that are of relevance for this work is provided in Appendix A.

2.1.2 Magnetic interaction

This section summarizes the interaction of UCNs with magnetic fields as described in [GRL91]. Although the neutron is electrically neutral, it possesses a magnetic dipole moment of magnitude [Moh+25]

$$\mu_\mathrm{n} = 9.662\,365\,3(23) \times 10^{-27}\,\mathrm{J/T} \approx 60.31\,\mathrm{neV/T} \tag{2.13}$$

and has a spin of 1/2. The magnetic potential for a neutron in a magnetic field $\vec{B}(\vec{r})$ is therefore given by

$$V_\mathrm{m} = -\vec{\mu}_\mathrm{n} \cdot \vec{B}(\vec{r}) \tag{2.14}$$

and depends on the relative orientation of the magnetic moment and the magnetic field. As a consequence, a neutron in an inhomogeneous magnetic field experiences a Stern-Gerlach force

$$\vec{F}_\mathrm{m} = -\nabla V_\mathrm{m} = \pm\mu_\mathrm{n} \nabla B(\vec{r}), \tag{2.15}$$

which is either attractive or repulsive, depending on whether the magnetic moment is aligned or anti-aligned with the field. Accordingly, neutrons attracted to regions of high magnetic fields are referred to as high field seeking (HFS), whereas those repelled by high magnetic fields are referred to as low field seeking (LFS). Within this framework, it is implicitly assumed that the velocity of the neutron is sufficiently small such that its magnetic moment maintains the same orientation with respect to the magnetic field while traversing it. This is also known as the adiabaticity condition and is satisfied when the temporal change of the magnetic field experienced by the neutron is much smaller than its Larmor frequency. The Larmor frequency, defined as the precession frequency of the magnetic moment in the magnetic field, is given by

$$\omega_{\mathrm{L}} = 2\frac{\vec{\mu_{\mathrm{n}}} \cdot \vec{B}(\vec{r})}{\hbar}. \tag{2.16}$$

Formally, the adiabaticity condition can therefore be expressed as

$$\frac{1}{B}\left|\frac{d\vec{B}}{dt}\right| \ll \omega_{\mathrm{L}}. \tag{2.17}$$

Since the magnetic potential associated with fields on the order of 1 T is comparable to the kinetic energy of UCNs, magnetic fields can be used to confine UCNs in magnetic traps.

2.1.3 Interaction with gravity

The gravitational potential for a neutron at the Earth's surface is, to good approximation, given by

$$V_{\mathrm{g}} = m_{\mathrm{n}} g h, \tag{2.18}$$

where $m_{\mathrm{n}} = 1.674\,927\,500\,56(85) \times 10^{-27}\,\mathrm{kg}$ [Moh+25] denotes the neutron mass, h the height and g the standard acceleration due to gravity. A neutron moving upwards in the Earth's gravitational field loses approximately 102.5 neV/m of kinetic energy, which is of the same order as the typical kinetic energy of UCNs. Therefore, gravitational effects must be taken into account in the design of UCN experiments. Vertical guide sections can e.g. be used to shift the UCN energy spectrum by exploiting this effect. Additionally, gravity can be utilized to trap UCNs vertically by loading them into a sufficiently high trap.

2.1.4 Free neutron decay

Because the mass of the free neutron exceeds the combined mass of proton, electron and electron anti-neutrino, free neutrons are unstable and undergo β^--decay described by

$$n \rightarrow p + e^- + \overline{\nu}_e + 782\,\mathrm{keV}. \tag{2.19}$$

The released energy corresponds to the mass difference between the neutron and its decay products, calculated using the recommended values in [Moh+25]. At the quark level, one of the down valence quarks is converted into an up quark via the emission of a W^- boson, thereby transforming the neutron into a proton. The corresponding Feynman diagram is shown in Fig. 2.1. This intrinsic instability ultimately limits the storage time of UCNs in experiments. The current status of measurements of the neutron lifetime arising from the β-decay is discussed in detail in Section 2.3.

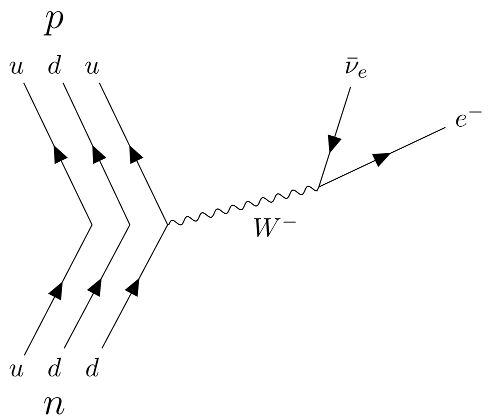

FIGURE 2.1: Feynman diagram of the free neutron β⁻-decay.

2.2 Ultracold neutrons in the Standard Model

In the last two decades UCNs have become an important tool for probing the Standard Model of particle physics. Among the experiments that have attracted the most attention are measurements of the neutron electric dipole moment (EDM), which, if observed at a level above the value predicted by the Standard Model, would provide evidence for physics beyond the Standard Model contributing to the baryon asymmetry observed in the universe [Abe+20]. Next to neutron EDM measurements, measurements of the neutron lifetime τ_n are of particular interest. The relevance of neutron lifetime measurements for both the Standard Model and cosmology is outlined in the following sections.

2.2.1 Unitarity of the CKM Matrix

The Cabibbo-Kobayashi-Maskawa (CKM) matrix describes the mixing between quark mass eigenstates and weak eigenstates and is given by

$$\begin{pmatrix} \mathrm{d}' \\ \mathrm{s}' \\ \mathrm{b}' \end{pmatrix} = \begin{pmatrix} V_{\mathrm{ud}} & V_{\mathrm{us}} & V_{\mathrm{ub}} \\ V_{\mathrm{cd}} & V_{\mathrm{cs}} & V_{\mathrm{cb}} \\ V_{\mathrm{td}} & V_{\mathrm{ts}} & V_{\mathrm{tb}} \end{pmatrix} \cdot \begin{pmatrix} \mathrm{d} \\ \mathrm{s} \\ \mathrm{b} \end{pmatrix}, \tag{2.20}$$

where d', s' and b' denote the weak eigenstates and d, s and b the mass eigenstates [Tho13, sec. 14.2-14.3]. Cabibbo originally introduced the Cabibbo mixing angle, which is equivalent to a 2×2 submatrix involving down and strange quarks, to explain the observed differences in coupling strength between up-down and up-strange weak interaction vertices [Cab63]. The Cabibbo mechanism was later extended to three quark generations by Kobayashi and Maskawa [KM73], leading to the full CKM matrix. Here, the quantities $|V_{ij}|^2$ are proportional to the transition probabilities between the corresponding quark flavors [Pov+14, p. 152]. According to the Standard Model, the CKM matrix is unitary and can be parametrized by four independent parameters [Tho13, sec. 14.3]. Consequently, the squared absolute values of each row and column must sum to unity, implying that all possible weak interaction transitions for a given quark are fully accounted for within the matrix. Testing the unitarity of the CKM matrix therefore constitutes a stringent probe of the internal consistency of the Standard Model. A violation of unitarity could e.g. indicate the existence of additional quarks, the existence of additional bosons or a hidden decay channel [BBB20]. To test unitarity, precise determinations of the individual matrix elements are required. For the first row, unitarity imposes the condition

$$|V_{\mathrm{ud}}|^2 + |V_{\mathrm{us}}|^2 + |V_{\mathrm{ub}}|^2 = 1. \tag{2.21}$$

Measurements of V_{us} are obtained from kaon, hyperon and tau decays [Nav+24, p. 904], with the most precise values at present coming from kaon decays. The particle data group (PDG) currently recommends $V_{\mathrm{us}} = 0.224\,31(85)$ [Nav+24, p. 905]. The most precise measurements of V_{ub} are obtained from B-meson decays and the PDG recommends $V_{\mathrm{ub}} = 3.82(20) \times 10^{-3}$ [Nav+24, p. 967]. The matrix element V_{ud} provides the largest contribution to the first row. Its most precise determination is obtained from 15 superallowed $0^+ \to 0^+$ β-decays [Nav+24, p. 904]. These transitions are pure Fermi decays and depend only on the vector component of the weak interaction [HT20]. The vector coupling constant G_{V}, related to V_{ud} via

$$G_{\mathrm{V}} = G_{\mathrm{F}} V_{\mathrm{ud}}, \tag{2.22}$$

can be extracted by measurements of the lifetime and the transition energies of these decays [HT20]. Here, G_{F} denotes the Fermi constant, which is determined with high precision from muon lifetime measurements [Nav+24, p. 186][MuL+11]. However, the extraction of V_{ud} from superallowed $0^+ \to 0^+$ β-decays requires additional theoretical input, including radiative and nuclear structure-dependent corrections [HT20]. At present, uncertainties associated with these corrections dominate the overall uncertainty budget. Consequently, the PDG reports them separately and recommends

$$V_{\mathrm{ud}} = 0.97367(11)_{\mathrm{exp}}(13)_{\mathrm{R}}(27)_{\mathrm{NS}}(32)_{\mathrm{total}}, \tag{2.23}$$

where "exp" denotes the experimental uncertainty, "R" the radiative correction uncertainty and "NS" the nuclear structure correction uncertainty [Nav+24, p. 904]. The total uncertainty is obtained by summing the squared individual contributions. In contrast to superallowed $0^+ \to 0^+$ β-decays, neutron β-decay does not require nuclear structure corrections. However, it involves both Fermi and Gamow-Teller components, meaning that the neutron couples to both the vector and axial-vector components of the weak interaction [Pov+14, pp. 274-276]. The neutron decay rate $\Gamma_{\mathrm{n}} = 1/\tau_{\mathrm{n}}$ in natural units is given by

$$\Gamma_{\mathrm{n}} = \frac{G_{\mathrm{F}}^2 \, |V_{\mathrm{ud}}|^2 \, m_{\mathrm{e}}^5}{2\pi^3} (1 + 3\lambda^2)(1 + \mathrm{RC}) f, \tag{2.24}$$

where m_{e} is the electron mass, f the so-called phase space factor, $\lambda = G_{\mathrm{A}}/G_{\mathrm{V}}$ the ratio of axial-vector to vector coupling and RC denotes the radiative corrections [CMS04; DM21]. The parameter λ accounts for the fact that the neutron interacts with both the vector and the axial-vector component. Due to parity violation in the weak interaction, spin-momentum correlations arise, leading to an asymmetric angular distribution of the decay products. The magnitude of this asymmetry depends on λ [Nic09]. Consequently, measurements of τ_{n} and λ, combined with theoretical calculations of radiative corrections, enable a determination of V_{ud} that is free from nuclear structure uncertainties. Based on the most recent theoretical evaluation of the neutron β-decay, V_{ud} can be expressed as

$$|V_{\mathrm{ud}}|^2 = \frac{5263.284(17)\,\mathrm{s}}{\tau_{\mathrm{n}}(1 + 3\lambda^2)(1 + \Delta_{\mathrm{R}})(1 + \delta)}, \tag{2.25}$$

where Δ_R denotes short-distance electroweak and hadronic radiative corrections and δ long distance radiative corrections from quantum electrodynamics [Van+25]. The recommended values are $\delta = 27.04(7) \times 10^{-3}$ and $\Delta_R = 45.37(27) \times 10^{-3}$ [Van+25], where Δ_R is based on [Cir+23a] and adapted to the renormalization scheme used in [Van+25]. In recent years, λ has been determined from measurements of the correlation coefficient A between neutron spin and electron momentum [Mär+19; Bro+18b] as well as the correlation coefficient a between electron and electron antineutrino momenta [Bec+20; Has+21]. The PDG currently recommends $|\lambda| = 1.2756(13)$ [Nav+24, p. 904], while the single most precise result $\lambda = -1.276\,41(56)$ was obtained by the PERKEO III experiment [Mär+19]. For the neutron lifetime, the PDG currently recommends $\tau_n = 879.4(6)\,\mathrm{s}$ [Nav+24, p. 904], while the most precise single result, $\tau_n = 877.83(22)_{\mathrm{stat}}(^{+0.22}_{-0.17})_{\mathrm{syst}}\,\mathrm{s}$ was obtained by the UNCτ experiment [Mus+25]. A more detailed overview of neutron lifetime measurements is provided in Section 2.3. Combining these inputs yields

$$\left|V_{\mathrm{ud}}^{\mathrm{n,PDG}}\right| = 0.97356(33)_\tau(82)_\lambda(13)_{\Delta_R}(3)_\delta(90)_{\mathrm{total}} \tag{2.26}$$

for PDG averages and

$$\left|V_{\mathrm{ud}}^{\mathrm{n,best}}\right| = 0.97392(17)_\tau(35)_\lambda(13)_{\Delta_R}(3)_\delta(41)_{\mathrm{total}} \tag{2.27}$$

for the most precise individual results of τ_n and λ. Although the precision of V_{ud} from neutron β-decay is approaching that obtained from superallowed β-decays, further improvements in both τ_n and λ are required to fully exploit the absence of nuclear structure corrections. Currently, the first row of the CKM unitarity exhibits a tension at the level of 2.3σ. Using the most precise neutron decay results instead of the superallowed β-decay value for V_{ud} reduces this tension. However, significant discrepancies remain among different determinations of V_{us} and V_{ud}. These tensions are discussed in [Cir+23b; Hay24; GS23] and motivate further complementary measurements. The experimental tensions in case of the neutron lifetime are discussed in detail in Section 2.3.3.

2.2.2 Big Bang nucleosynthesis

The neutron lifetime is also a parameter of interest in Big Bang nucleosynthesis (BBN), which describes the formation of light nuclei in the early universe prior to the formation of the first stars. Standard BBN assumes that the neutron to proton ratio n/p was in thermal equilibrium when the temperature of the universe was $T \gtrsim 1\,\mathrm{MeV}$ [ABG48][Esp+99]. During this phase, equilibrium is mediated by the following reactions:

$$n + \nu_e \longleftrightarrow p + e^- \tag{2.28}$$
$$n + e^+ \longleftrightarrow p + \overline{\nu}_e \tag{2.29}$$
$$n \longleftrightarrow p + e^- + \overline{\nu}_e \tag{2.30}$$

The corresponding equilibrium ratio is given by

$$n/p = \exp(-\Delta m/T), \tag{2.31}$$

where Δm denotes the neutron proton mass difference and T the temperature [YOF23]. The weak reaction rate Γ_{wk}, which determines the rate of these processes, scales as $G_F^2 T^5$, where G_F is the Fermi constant. Maintaining thermal equilibrium requires this rate to exceed the rate of change of the thermodynamic environment. In the early universe, the latter is governed by cosmic expansion, characterized by the Hubble parameter $H \sim \sqrt{G_\mathrm{N}}T^2$, where G_N is Newton's constant. Therefore, the so-called weak freeze out occurs approximately when the condition

$$\Gamma_\mathrm{wk} \approx H \tag{2.32}$$

is satisfied, corresponding to a temperature of approximately 0.84 MeV [YOF23]. Beyond this point, the weak interactions can no longer maintain equilibrium and their rates decrease further until only neutron β-decay remains relevant. However, even after weak freeze out, nuclei cannot form immediately. The reason is that the formation of light nuclei proceeds through deuterium, which has a relatively small binding energy of approximately $B_\mathrm{D} = 2.23\,\mathrm{MeV}$ [Nav+24, p. 474]. Given the baryon to photon ratio $\eta \sim 10 \times 10^{-9}$, the high energy tail of the photon energy distribution rapidly destroys forming deuterium nuclei by photodisintegration, preventing the buildup of a stable deuterium abundance [Cyb+16]. This effect is known as the deuterium bottleneck. The bottleneck is lifted once the condition [Cyb+16]

$$\eta \sim \exp -B_\mathrm{D}/T, \tag{2.33}$$

is satisfied, which occurs approximately three minutes after the Big Bang [MKS05; YOF23]. Subsequently, light nuclei such as ^{2}D, ^{3}He, ^{4}He and ^{7}Li are formed [Fie+20]. The final abundances of these nuclei depend on the neutron to proton ratio at the time the deuterium bottleneck is overcome. Since the conversion of protons to neutrons effectively ceases after the weak freeze out and the deuterium bottleneck lasts for a significant fraction of the neutron lifetime, the neutron lifetime becomes a key parameter in BBN theory [YOF23]. Additionally, the neutron lifetime enters the normalization of the weak interaction matrix element used to calculate Γ_wk, thereby influencing the weak freeze out condition and the initial neutron to proton ratio [YOF23]. Almost all neutrons present after the deuterium bottleneck are incorporated into ^{4}He, while the abundances of ^{2}D and ^{3}He are of order 1×10^{-5} and that of ^{7}Li is of order 1×10^{-10} [Cyb+16]. Accordingly, the ^{4}He mass fraction is in good approximation given by

$$Y_\mathrm{p} \approx \frac{2n/p}{1 + n/p}, \tag{2.34}$$

since ^{4}He consists of two protons and two neutrons of approximately equal mass [Cyb+16]. The primordial helium mass fraction Y_p is the light element fraction most sensitive to the neutron lifetime and its relative uncertainty is approximately given by [YOF23]

$$\frac{\Delta Y_\mathrm{p}}{Y_\mathrm{p}} \approx \frac{\Delta \tau_\mathrm{n}}{\tau_\mathrm{n}}. \tag{2.35}$$

Using the current PDG average of the neutron lifetime [Nav+24, p. 93], this corresponds to a relative theoretical uncertainty of $\Delta Y_\mathrm{p}/Y_\mathrm{p} = 4 \times 10^{-4}$. Experimentally, Y_p is most precisely determined from recombination emission lines in extragalactic HII regions [ITG14; Ave+21]. The PDG recommends $Y_\mathrm{p} = 0.245(3)$ [Nav+24, p. 475], corresponding to a relative uncertainty of approximately 1×10^{-2}. This is significantly larger than the relative uncertainty of the theory predictions. Consequently, more precise measurements of Y_p are required to rigorously test BBN predictions. Nevertheless, more precise measurements of the neutron lifetime still enhance the predictive power of BBN by reducing the theoretical uncertainty.

2.3 Neutron lifetime measurements

Currently, the most precise measurements of the neutron lifetime are obtained using two distinct approaches: the beam method and the bottle method [Wie24]. The two methods are described in detail in the following sections.

2.3.1 Beam method

For the beam method, a beam of slow neutrons is guided through a well-known decay volume. The instantaneous decay rate dN/dt is then related to the neutron lifetime τ_n by

$$\frac{dN}{dt} = -\frac{N}{\tau_\mathrm{n}}, \tag{2.36}$$

where N is the number of neutrons in the decay volume [WG11]. Thus, determining the neutron lifetime using the beam method requires measurements of both the absolute decay rate and the absolute number of neutrons within the decay volume. To determine the number of neutrons in the decay volume, beam type experiments typically monitor the neutron flux downstream of the decay volume [WG11]. However, in some implementations, the neutron flux is instead measured directly within the decay volume using a time projection chamber (TPC) containing a small amount of ^{3}He [Sch+89; Sum+21]. Different experimental realizations have employed a variety of techniques to detect the decay products, targeting the electrons or the protons created in the free neutron decay. In one of the first dedicated beam type measurement of the neutron lifetime conducted in the 1950s, the protons were collected using an electrostatic collection system and detected in coincidence with the electrons to suppress background [Rob51]. Later approaches used static magnetic fields to guide decay electrons into the detector [Chr+72]. In the TPC based designs, both neutron flux monitoring and decay product detection are performed within the same detector [Sch+89; Sum+21]. The most precise beam-method implementations employ a Penning trap to confine and subsequently count decay protons [Byr+90; Nic+05]. In this case, the neutron lifetime is related to the measured neutron and decay rates by

$$\tau_\mathrm{n} \propto \frac{R_\mathrm{n}\epsilon_\mathrm{p}}{R_\mathrm{p}\epsilon_\mathrm{n}}, \tag{2.37}$$

where R_n and R_p denote the neutron and decay rates and ϵ_n and ϵ_p are the corresponding detection efficiencies [Wie18]. The inclusion of detection efficiencies is necessary because neutron flux and decay products are measured with separate detectors or the detection exhibits different efficiencies. Although the measurement principle is conceptually straightforward, a range of systematic effects must be considered to extract a final result. These include uncertainties related to the neutron absorption in the detector, the trap geometry and the spatial profile of the neutron beam [Nic+05]. The most precise result, obtained from a reanalysis of [Nic+05], yields $\tau_\mathrm{n} = 887.7(1.2)_\mathrm{stat}(1.9)_\mathrm{syst}\,\mathrm{s}$ [Yue+13].

2.3.2 Bottle method

The bottle method provides a complementary approach to the beam method. Instead of counting decay products from a neutron beam, UCNs are stored in a trap. By measuring the number of surviving neutrons, after different storage times, the exponential decay law can be reconstructed and the neutron lifetime inferred [Wie24]. In principle, it is sufficient to measure at two storage times t_1 and t_2. The neutron lifetime is then given by

$$\tau_\mathrm{n} = \frac{t_2 - t_1}{\ln\left(\frac{N_1}{N_2}\right)}, \tag{2.38}$$

where N_1 and N_2 denote the number of neutrons remaining after storage times t_1 and t_2, respectively [Wie24]. Conceptually, this method is straightforward and requires only a relative measurement of N_1 and N_2 as well as the difference between the storage times. However, it relies critically on the assumption that neutrons are exclusively lost through β-decay. Any additional loss mechanisms therefore introduce biases [Wie24]. Early bottle-type experiments used material traps with high Fermi potentials. These systems exhibited several additional loss mechanisms, including [WG11; Wie18]:

- Inelastic scattering of UCNs on the walls,
- absorption on the walls,
- absorption on residual gas,
- escape through gaps or slits in the wall and
- loss of marginally trapped neutrons, i.e., neutrons with energies exceeding the trap potential but long residence times.

Assuming these processes occur at constant rates and affect the entire population uniformly, the effective storage time constant can be written as

$$\frac{1}{\tau_\mathrm{store}} = \frac{1}{\tau_\mathrm{n}} + \frac{1}{\tau_\mathrm{inel}} + \frac{1}{\tau_\mathrm{abs}} + \frac{1}{\tau_\mathrm{gas}} + \frac{1}{\tau_\mathrm{gap}} + \frac{1}{\tau_\mathrm{marg}} + \frac{1}{\tau_\mathrm{other}}, \tag{2.39}$$

where τ_other accounts for loss mechanism not listed above [Wie18]. In practice, many of these loss mechanisms depend on the neutron velocity and therefore on the evolving velocity spectrum during storage [WG11]. To mitigate these effects, bottle-type experiments typically employed highly polished surfaces and maintained high vacuum conditions. Additional strategies include specialized wall coatings with high Fermi potentials and low absorption cross sections [Mam+89; Kha+89; Arz+00; Ser+05] as well as cooling of trap walls to suppress thermal upscattering [Kha+89; Ser+05]. Despite these efforts, wall related losses cannot be completely eliminated. Therefore, experiments often vary the surface to volume ratio of the trap and extrapolate measured storage times to the limit of zero loss [Wie18]. Additionally, techniques such as pre-storage volumes with absorbing lids are used to remove marginally trapped neutrons [WG11; Pic+10]. The most precise result obtained with material bottles is $\tau_\mathrm{n} = 878.5(7)_\mathrm{stat}(3)_\mathrm{syst}\,\mathrm{s}$ [Ser+05].

More recent bottle-type neutron lifetime experiments employ magnetic trapping instead of material confinement. This approach eliminates wall interactions and their associated loss mechanisms [Wie18]. However, only neutrons in the LFS spin state can be confined in this way and neutron polarization must be maintained throughout the storage period. This is achieved by ensuring that the trapping field satisfies the adiabaticity condition (see Eq. (2.17)) across the entire trap volume. Depolarization related losses are discussed in detail in [Ste+17; SG20]. At present, two experiments have achieved competitive precision. In both experiments a so-called magneto-gravitational trap was employed, in which UCNs are confined radially and from below by permanent magnets, while gravity provides vertical confinement. The result reported in [Ezh+18], $\tau_n = 878.3(1.6)_{stat}(1.0)_{syst}$ s, required a systematic correction of 3.7 s due to depolarization effects. In contrast, the magnetic field configuration of the UCNτ experiment was designed such that no such corrections were required. The most recent result $\tau_n = 877.83(22)_{stat}(^{+0.22}_{-0.17})_{syst}$ s [Mus+25] currently represents the most precise measurement of the neutron lifetime.

2.3.3 The neutron lifetime puzzle

Fig. 2.2 summarizes the most precise neutron lifetime measurements obtained using the beam and bottle methods, including reanalyses of results since 1995. As is evident, a significant discrepancy exists between the results obtained with the two approaches. The PDG currently recommends $\tau_n = 878.4(5)$ s, which does not include results from the beam-type measurements [Nav+24, p. 2091]. Averaging the two most precise independent beam measurements ([Byr+96] and [Yue+13]) yields $\tau_n = 888.0(21)$ s [Yue+13]. That is, the beam and bottle averages differ by 9.6 s, corresponding to a tension of 4.6σ. This discrepancy is referred to as the neutron lifetime puzzle. Since the beam method measures decay products, whereas the bottle method counts surviving neutrons, several new physics explanations have been proposed [Wie24]. Among

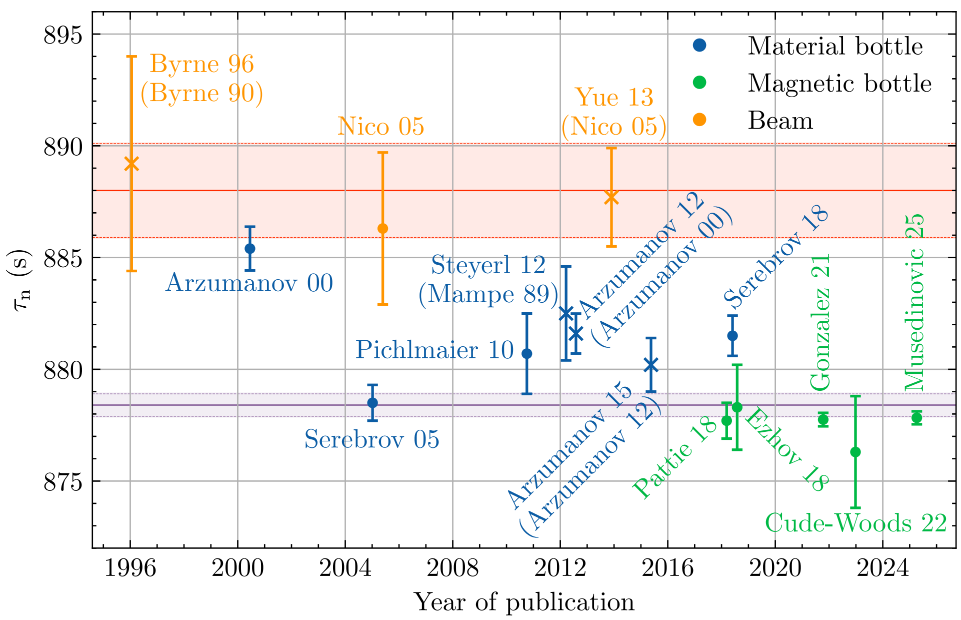

FIGURE 2.2: Most precise published measurements of the neutron lifetime as a function of publication year, including results obtained with beam and bottle methods as well as reanalyses since 1995. Cross markers indicate reanalyses of earlier results, which are listed in brackets in the corresponding labels. The violet band indicates the current PDG average [Nav+24, pp. 2091-2092], including the results of [Gon+21] (Gonzalez 21), [Ezh+18] (Ezhov 18), [Pat+18] (Pattie 18), [Ser+18] (Serebrov 18), [Arz+15] (Arzumanov 15), [Ste+12] (Steyerl 12), [Pic+10] (Pichlmaier 10) and [Ser+05] (Serebrov 05). The red band represents the average of the two most precise beam-type measurements, [Byr+96] (Byrne 96) and [Yue+13] (Yue 13). Earlier results superseded by later reanalyses are also shown, including [Arz+00] (Arzumanov 00), [Nic+05] (Nico 05) and [Arz+12] (Arzumanov 12). The original results [Mam+89] (Mampe 89) and [Byr+90] (Byrne 90) are not shown, as they were published before 1995. The recently published results [Cud+22] (Cude-Woods 22) and [Mus+25] (Musedinovic 25) are not included in the PDG average as of now. The result reported in Musedinovic 25 includes that of Gonzalez 21.

these is e.g. a dark neutron decay [FG18], whose proposed decay channels containing visible particles, however, have largely been excluded experimentally [Tan+18; Sun+18; Klo+19]. An alternative explanation is that the discrepancy arises from unaccounted or underestimated systematic effects in either method [Wie24]. This interpretation is supported by the history of reanalyses, which have led to shifts in previously reported results (see Fig. 2.2). To date, however, no systematic effect has been identified that accounts for the observed discrepancy. Consequently, current efforts focus on performing new beam and bottle experiments with improved control of systematic uncertainties as well as on varied approaches with different sets of systematic effects.

2.3.4 Current status of neutron lifetime measurements

As discussed in Section 2.2.1 and Section 2.2.2, precise measurements of the neutron lifetime remain of significant importance for testing the Standard Model and for understanding Big Bang nucleosynthesis. In addition, further measurements are required to resolve the neutron lifetime puzzle (see Section 2.3.3). In this context, new beam-type experiments with improved precision are being developed. The BL2 experiment, a successor to the most precise beam measurement [Yue+13] operating on the same principle, is taking data at NIST and aims to achieve an uncertainty below 2 s [Hoo+19; Cay+25; Wie25]. A further successor, BL3, is currently under development. It is projected to measure the neutron lifetime with a precision of better than 0.3 s, while providing enhanced control over systematic effects [Wie24; Wie25]. A complementary beam-type experiment is taking data at J-PARC. This experiment employs a TPC containing a small fraction of ^{3}He, enabling simultaneous measurements of neutron flux and decay products within a single detector. This approach introduces a different set of systematic effects compared to Penning-trap based experiments [Fuw+24]. At Los Alamos National Laboratory, the UCNProBe experiment is being commissioned. It aims to determine the neutron lifetime with a precision of approximately 2 s and employs a hybrid beam-bottle approach by detecting decay products from stored UCNs. This method is particularly relevant in the context of the neutron lifetime puzzle [Kri25]. To further increase the precision of neutron lifetime measurements, the UCNτ experiment is undergoing an upgrade of its filling mechanism as part of the UCNτ+ project, with the goal of increasing the number of stored UCNs by a factor of 5-10 [Mus+25]. A new bottle-type experiment, τSPECT, is currently taking data at PSI with the goal of achieving a precision better than 0.3 s. It employs a fully magnetic trap, providing a complementary approach to magneto-gravitational traps. Additionally, the trap design is scalable to a deeper trapping potential, which could enable a determination of the neutron lifetime with sufficient precision for a competitive determination of V_{ud} [Aul+24; Aul+25]. The τSPECT experiment is described in detail in the subsequent chapter. In addition to beam and bottle approaches, the neutron lifetime has also been inferred from space-based measurements of neutron fluxes above planetary surfaces exposed to cosmic radiation [Wil+20]. The most precise result to date was obtained using a lunar orbiter. Although current uncertainties exceed 10 s, future improvements may provide an independent cross-check relevant to resolving the neutron lifetime puzzle [Wil+21].

Chapter 3

τSPECT

3.1 Experimental setup

τSPECT is a three-dimensional magnetic field gradient trap for UCNs. After its commissioning at the research reactor TRIGA Mainz, it is now located at beam port west 1 (W1) of the UCN source at the Paul Scherrer Institute (PSI) in Switzerland, which is operated in pulsed mode (see Section 3.4). A schematic cross section of the current experimental setup is shown in Fig. 3.1. At the core of the setup is the trap region, in which UCNs can be magnetically confined. When neutrons enter a magnetic field, their magnetic moment either aligns or anti-aligns with the field lines (see Section 2.1.2). In case of antiparallel alignment neutrons are low-field-seeking (LFS) and are decelerated by an increasing magnetic field. It is therefore possible to trap LFS neutrons in a region in which the absolute value of the magnetic field exhibits a minimum. The confining magnetic fields are produced by a solenoidal superconducting coil system and a permanent magnet Halbach octupole array [Hal80]. The superconducting coils are housed in a cylindrical cryostat and provide longitudinal confinement along the x-axis (see Fig. 3.1) by generating the field profile indicated by the red line in Fig. 3.1. Both the cryostat and the superconducting coil system were repurposed from the aSPECT experiment [Bec+20; Bae+08]. Radial confinement is provided by the magnetic field produced by the Halbach array, which is positioned such that it overlaps with the minimum of the longitudinal field. The magnitude of the superposition of the octupole and the longitudinal field in the xy-plane of the trap is shown in Fig. 3.2. A more detailed description of the magnetic field configuration and the experimental setup in general can be found in [Aul+24; Aul+25].

In the following, the working principle of the experiment is summarized by tracing the path of a neutron from the beam port into the trap. During this discussion, minor modifications relative to the setup described in [Aul+24] are highlighted, and references to more detailed descriptions of specific subsystems are provided. To guide neutrons from the beam port to the entrance of the experiment, a set of highly polished[1] stainless steel guides ($V_{\mathrm{F}} = 184\,\mathrm{neV}$, see Appendix A) with an inner diameter of 66 mm is used. In the first section, the neutrons are guided horizontally towards a vertical guide. At the bottom and top of the vertical guide ^{58}NiMo (85:15, $V_{\mathrm{F}} = 313\,\mathrm{neV}$, see Appendix A) coated mirrors are installed. This section guides the neutrons up to the height of the cryostat. Downstream of the top mirror, a monitor UCN detector is positioned. The neutron sensitive detector surface can be reached through a 6 mm diameter pinhole in the neutron guide wall and counts a small fraction of the

[1] According to the manufacturer (NEUMO GmbH + Co. KG), the internal surface finish complies with DIN 11866 (HE5o finish) and is additionally electropolished, resulting in an average surface roughness below 0.25 µm.

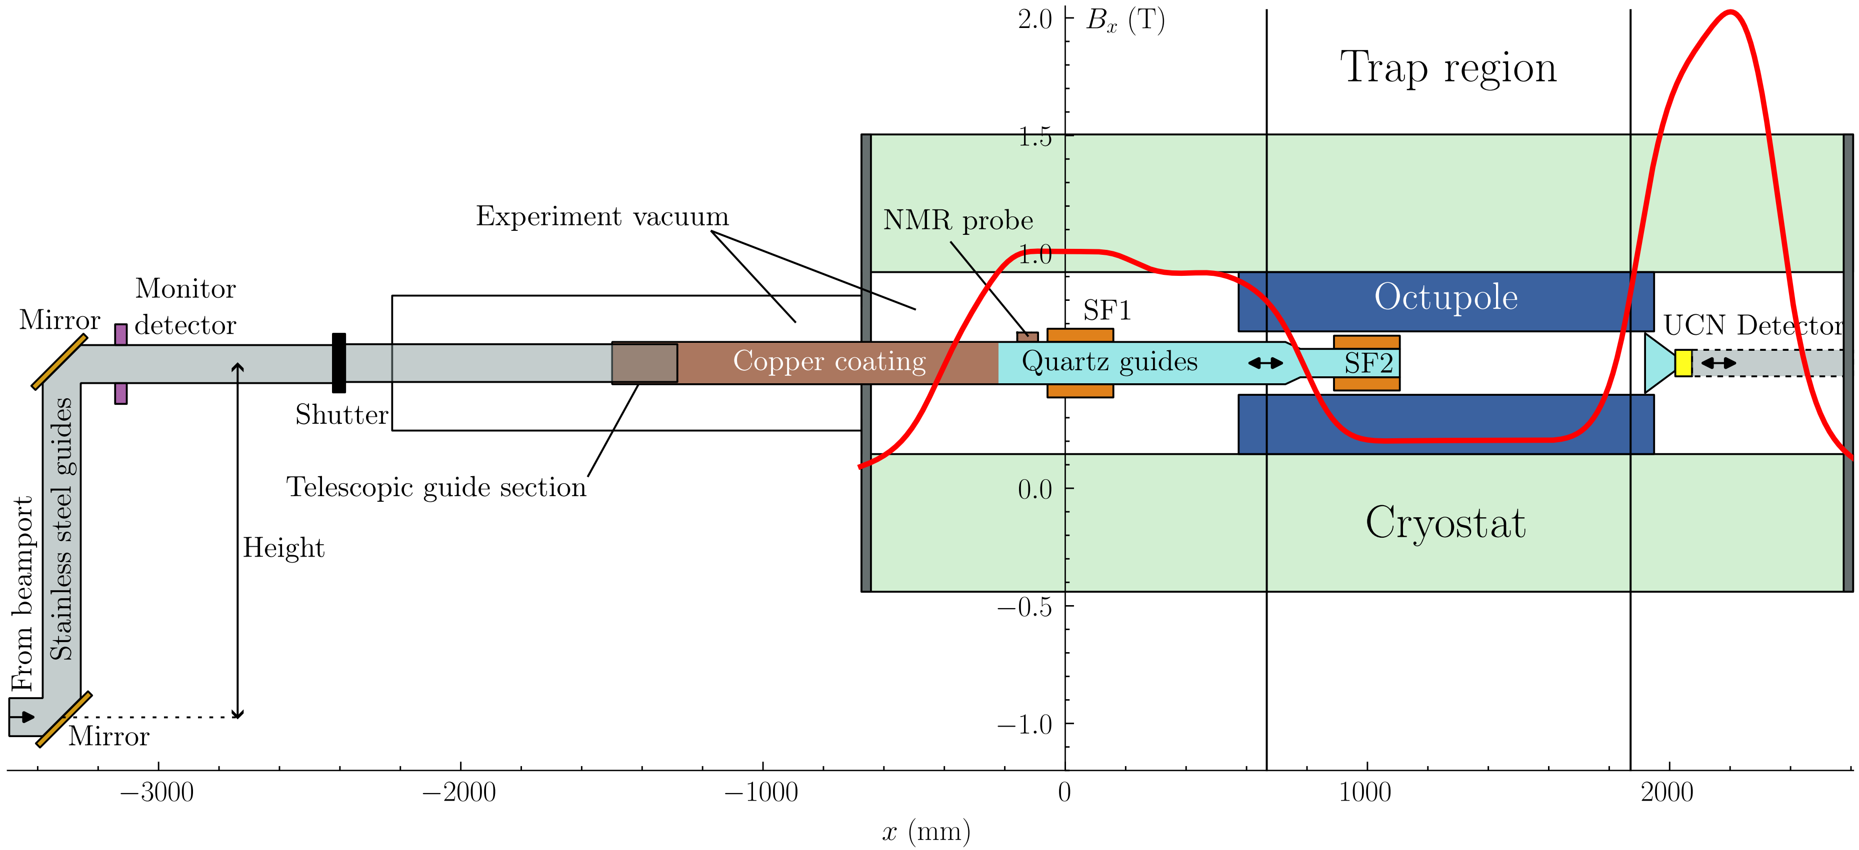

FIGURE 3.1: Schematic cross section of the τSPECT setup at PSI in 2024. UCNs from beam port W1 are guided into the experiment via stainless steel guides. Two ^{58}NiMo (85:15) coated mirrors reflect UCNs at the top and bottom of the vertical section. A monitoring detector is positioned downstream of the top mirror. A shutter is used to separate experiment and neutron guide vacuum and to allow UCNs into the experiment. Inside the experiment, a telescopic guide section with an inner stainless steel guide and an outer quartz glass guide directs UCNs into the trap. The first 1280 mm of the quartz guide are copper coated. Two spin-flipping units (SF1 and SF2) are mounted onto the quartz guides. The trap is formed by a radial field created by a permanent magnet octupole array and a longitudinal field (red line) created by superconducting coils inside the cryostat. A UCN detector is located at the opposite end of the cryostat and counts the trapped neutrons. Both detector and quartz guides are moved in and out of the trap by stepper motor driven translation stages.

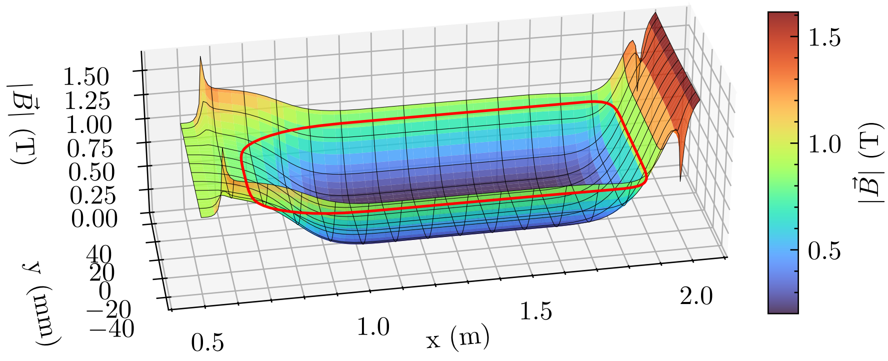

FIGURE 3.2: Magnitude of the combined longitudinal and radial magnetic field of τSPECT in the xy-plane in the trap region. The field was simulated using the τSPECT simulation framework described in [Aul+25]. The y-axis is defined perpendicular to the cross section plane shown in Fig. 3.1. The red line marks the edge of the trap (48 neV) as defined in [Aul+25], which also considers gravity.

traversing neutrons in this way. Its purpose is to monitor the UCN source performance, enabling the identification of abnormal neutron pulses and potential issues with the experimental setup. Furthermore, it can be used to account for variations in the UCN yield between different measurement runs. A detailed description of the monitor detector is provided in [Eng25]. The majority of neutrons bypass the monitor detector and subsequently reach the UCN shutter at the entrance of the experiment, which separates the experiment vacuum from the neutron guide vacuum and controls whether UCNs can reach the trap region. The shutter is opened only while neutrons are filled into the trap. After passing the UCN shutter, neutrons enter a telescopic guide section consisting of an inner stainless steel guide and an outer quartz guide with inner diameters of 66 mm and 73 mm, respectively. This guide section directs neutrons into the trap region and must be retracted after filling to achieve a material free trap.

To successfully trap a neutron, it must be in the LFS state once it reaches the trap region. However, since the confining field is constant over time, no LFS neutrons in the trappable energy range can pass the initial longitudinal magnetic field peak. Therefore, the neutron spin and thus the magnetic moment of high-field-seeking (HFS) neutrons reaching the trap region must be flipped. This is achieved using an adiabatic fast passage (AFP) spin flipping unit (SF2) mounted at the end of the quartz guide. A detailed description of the AFP spin flip technique, which was first proposed for UCN experiments in [LT84], is provided in [Hol+12]. The spin flipping unit irradiates a magnetic field oscillating at the Larmor frequency of the neutrons at its position, which is about 6 MHz for SF2. To maintain neutron polarization, the longitudinal field is kept at about 0.21 T in the trap region [Aul+25]. As a consequence, a HFS neutron entering the trap gains about 12.5 neV of energy when spin-flipped only by SF2. To exploit the full trap depth, an additional spin flipping unit (SF1) is installed near the maximum of the magnetic field peak towards the source and operates at a frequency of around 29 MHz. In this configuration, LFS neutrons entering the trap lose kinetic energy on both flanks of the magnetic field peak at the entrance and no longer gain energy when entering the trap. A more detailed discussion of the spin-flipping units, the energy acceptance of the trap and the filling process can be found in [Aul+24; Aul+25; Roß21]. In addition to the superconducting coils generating the

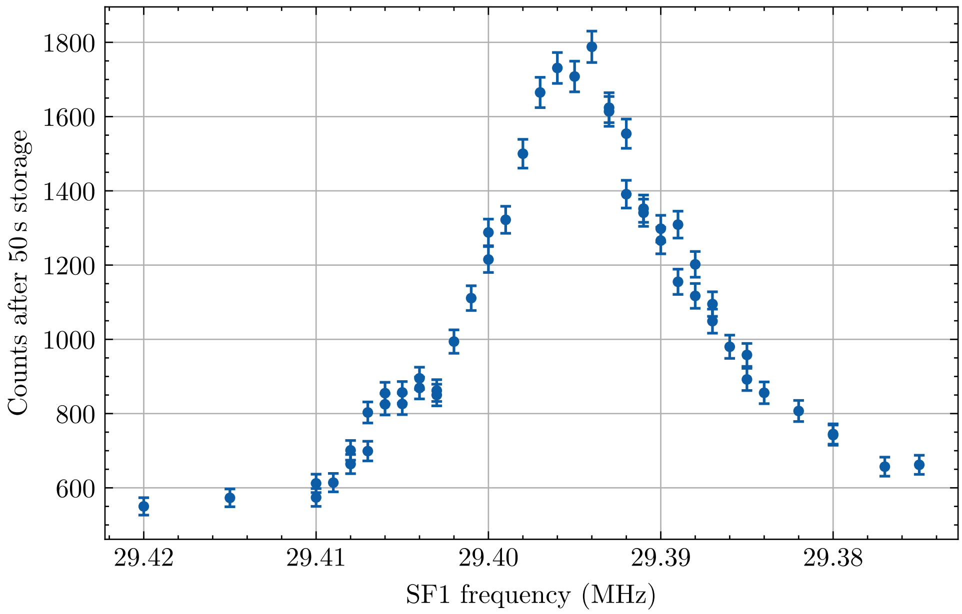

Figure 3.3: Number of counted neutrons after a storage time of 50 s as a function of the SF1 RF frequency. Dataset recorded in June 2025.

main longitudinal field, gradient coils, referred to as C3 and C5, are used to adjust the gradient around the position of SF1. While the main coils exhibit an inductance of 212 H, the combined inductance of the gradient coils is only 2.4 m H at similar cross section and length at the same order of magnitude [Cry]. Because the field produced by the main coils is very flat around the position of SF1 (see Fig. 3.1), energizing the gradient coils is sufficient to modify the local gradient at the position of SF1 without significantly altering the magnetic field magnitude. Since the efficiency of the spin flip depends on the gradient of the magnetic field at the position of the spin flipping unit [Hol+12], the gradient coils were used to optimize the performance of SF1. Experimentally, it was found that currents of $I_{C3} = 12.0\,\text{A}$ and $I_{C5} = 7.0\,\text{A}$ at $I_{\text{main}} = 32.5\,\text{A}$ yielded the best SF1 performance [Eng25]. This configuration was therefore used for all results presented in this thesis. Because the resulting gradient is relatively small, the frequency of SF1 must be adjusted at the kHz level to match the resonance frequency. This is highlighted by Fig. 3.3, which shows the number of neutrons stored after a storage time of 50 s as a function of the SF1 frequency.

Other members of the collaboration measured that the magnetic field of the superconducting coils decreases by about 26 µT/d using the nuclear magnetic resonance (NMR) probe indicated in Fig. 3.1, which is mounted close to the position of SF1. The development of the NMR probe is described in [Erm22] and the probe has since been installed in the experiment. To accommodate the environmental constraints inside the cryostat, the water sample was replaced by an ethanol-acetone mixture and a new copper housing was produced in the mechanical workshop [Eng25]. The observed drift corresponds to a shift in the neutron Larmor frequency of approximately 0.75 kHz/d. To compensate for this effect, the SF1 frequency is typically adjusted every two to three days. In case of SF2, the gradient is significantly larger and thus the resonance is significantly broader, such that no such measures are required. Additionally, no additional gradient coils are available at the position of SF2 and therefore no optimization of the gradient could be performed. The magnetic field drift had previously been observed when the superconducting magnet system was used in the aSPECT experiment and was attributed to power dissipation in the superconductors [Gua11]. Measures to correct and compensate for this drift are currently under investigation by other collaboration members. To avoid disturbances of the generated radio frequency (RF) field, the neutron guides at the position of the spin-flipping units must be nonmagnetic and electrically insulating. For this reason, quartz guides ($V_F = 91\,\text{neV}$, see Appendix A) are used inside the experiment and remain uncoated at the positions of SF1 and SF2. However, as the trap accepts neutrons with initial kinetic energies up to 144 neV at the experiment entrance when both spin-flipping units are in operation [Aul+25], the first 1280 mm of the quartz guides are coated with copper ($V_F = 171\,\text{neV}$, see Appendix A) to increase transmission. Since a spin flip at the position of SF2 requires compensation of the octupole field of the trap [Aul+24], an additional Halbach octupole array is installed around SF2. The permanent magnets of this array are arranged to cancel the field of the trap octupole. Due to the additional space required by this compensation octupole, the inner diameter of the quartz guide is reduced to 50 mm at the position of SF2.

Using the telescopic guide section, the quartz guides and the spin-flipping units can be moved in and out of the trap by a translation stage driven by a stepper motor. Typically, this is done at a speed of 24 mm/s. The position of the spin-flipping units is defined by the position of the center of SF2 along the τSPECT longitudinal axis shown in Fig. 3.1. At the opposite end of the cryostat, the movable in situ UCN detector is located. Similar to the spin-flipping units, it is moved in and out of the

FIGURE 3.4: Photograph of the τSPECT experiment at PSI in 2024. The four heavy duty lifting columns are used to adjust the height of the experiment relative to the beam port. The cryostat is mounted inside an iron yoke that shields the environment from the stray field of the superconducting coils inside the cryostat. The neutron source is located behind the concrete shielding on the left, and neutrons enter the experiment from the left.

trap via a translation stage and a stepper motor, typically at a speed of 10 mm/s. The detector is inserted into the trap to count the remaining neutrons at the end of each measurement run. Its position is defined by the location of its neutron sensitive surface. A detailed description of the detector can be found in [Aul+24; Kah20].

To gravitationally adjust the UCN energy spectrum at the experiment entrance, the height of τSPECT can be varied relative to the beam port using four heavy-duty lifting columns on which the experiment rests (see Fig. 3.4). When changing the height, only the vertical guide connecting τSPECT to the beam port must be swapped for another guide whose length fits the new height. The height is measured from the center of the horizontal guide at beam port height to the center of the horizontal guides inside the cryostat. For all measurements presented in this thesis, the height was 59 cm. As the trap region is located within the cryostat cooling the superconducting coils, the material surfaces surrounding the trap region cool to a steady-state temperature of approximately 170 K. The inner walls of the cryostat reach temperatures as low as 50 K. As a result, many gases freeze out on the inner walls, contributing to good vacuum conditions inside the experiment. In combination with four turbomolecular pumps, a steady state pressure of around 2×10^{-7} mbar is achieved in τSPECT, which is measured at a position farther away from the pumps than the trap region and therefore represents an upper estimate [Aul+24].

3.2 Marginally trapped neutrons

An important systematic uncertainty that must be well controlled in order to measure the neutron lifetime is caused by so-called marginally trapped neutrons. These are neutrons whose energy exceeds the trap potential but that nevertheless remain confined for a considerable amount of time because their trajectories prevent them from leaving the trap. As a result, this population of neutrons has an additional loss channel besides the β-decay. If this effect is not accounted for, it leads to a systematic underestimation of the neutron lifetime. To mitigate this effect, a so-called spectral cleaning procedure is employed in τSPECT. The purpose of the cleaning procedure is to remove marginally trapped neutrons from the trap before the actual storage period is started. To achieve this, the detector is moved to a position inside the trap, that lies within the rising flank of the detector-end peak of the longitudinal magnetic field for a defined period of time. Because marginally trapped neutrons exhibit higher energies than storable neutrons, they are expected to traverse this region more frequently and are therefore more likely to be absorbed by the detector. Afterwards, the detector is moved to its storage position, and the storage phase begins.

3.3 Measurement phases

A τSPECT measurement run can be divided into four phases: filling, cleaning, storage, and counting. The different measurement phases are illustrated in Fig. 3.5, which shows the detector count rate throughout a typical τSPECT measurement run. During the filling phase, the center of SF2 is located at 980 mm and the detector is in the cleaning position, typically around 1800 mm. The start of the filling phase is marked by the source trigger, which signals that the source will be pulsed after 1 s. The experiment shutter is opened and both SF1 and SF2 are switched on throughout the filling phase, which typically lasts around 25 s depending on the experiment height. For each height, the duration of the filling phase is optimized experimentally. At

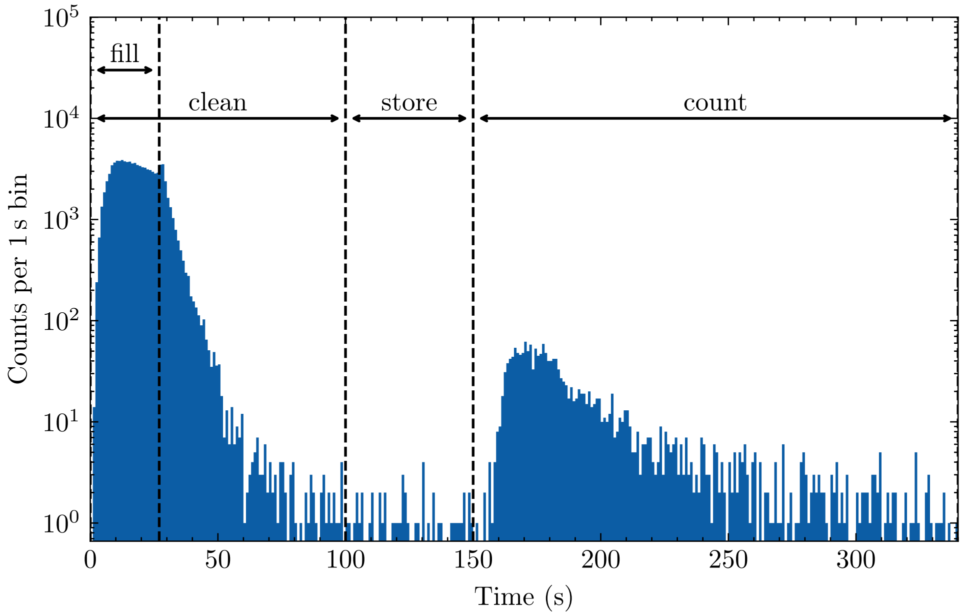

FIGURE 3.5: Count rate recorded by the τSPECT detector during a single exemplary measurement run. The four phases, filling, cleaning, storage, and counting, are indicated by dashed vertical lines and horizontal arrows. In the measurement shown, the filling phase lasted 27 s, the cleaning phase 100 s, the storage phase 50 s, and the counting phase 190 s. The measurement ID is 20241001_112758.

the end of the filling phase, SF2 is moved to its storage position at 400 mm. The detector remains in its cleaning position until the end of the cleaning phase, which starts simultaneously with the filling phase and typically lasts between 50 s and 100 s. At the end of the cleaning phase, the detector is moved to its storage position at 1920 mm. The storage phase begins when the detector starts moving towards its storage position. During this phase, neutrons remain trapped magnetically and do not interact with any material surfaces. At the end of the storage phase, the detector is moved to the counting position at 1605 mm and counts the remaining neutrons. Based on experimental data, a counting phase duration of 190 s was selected. Afterwards, the experiment is prepared for the next measurement run. The τSPECT UCN detector records data throughout the entire measurement run. In addition, the monitoring detector located upstream of the experiment records data during the first 240 s of each run. In contrast to earlier measurements discussed in [Yaz24; Aul+24; Roß21], the filling and cleaning phases were slightly adjusted. Instead of keeping the detector in its storage position during the filling phase and moving it to the cleaning position afterwards, it is already in its cleaning position at the start of the filling phase. That is, the filling and cleaning phases overlap for the duration of the filling phase. This approach was adopted because simulations using the τSPECT simulation framework [Aul+25] predict performance comparable to or better than the previous procedure while reducing the total duration of a single measurement run.

3.4 PSI UCN source

This section summarizes the working principle of the UCN source at PSI as described in [Bis+20; Bis+22]. The UCN source at PSI is driven by PSI's High Intensity Proton Accelerator (HIPA), which generates a 590 MeV proton beam with a beam current of up to 2.4 mA. A simplified schematic cross section of the UCN source is depicted in

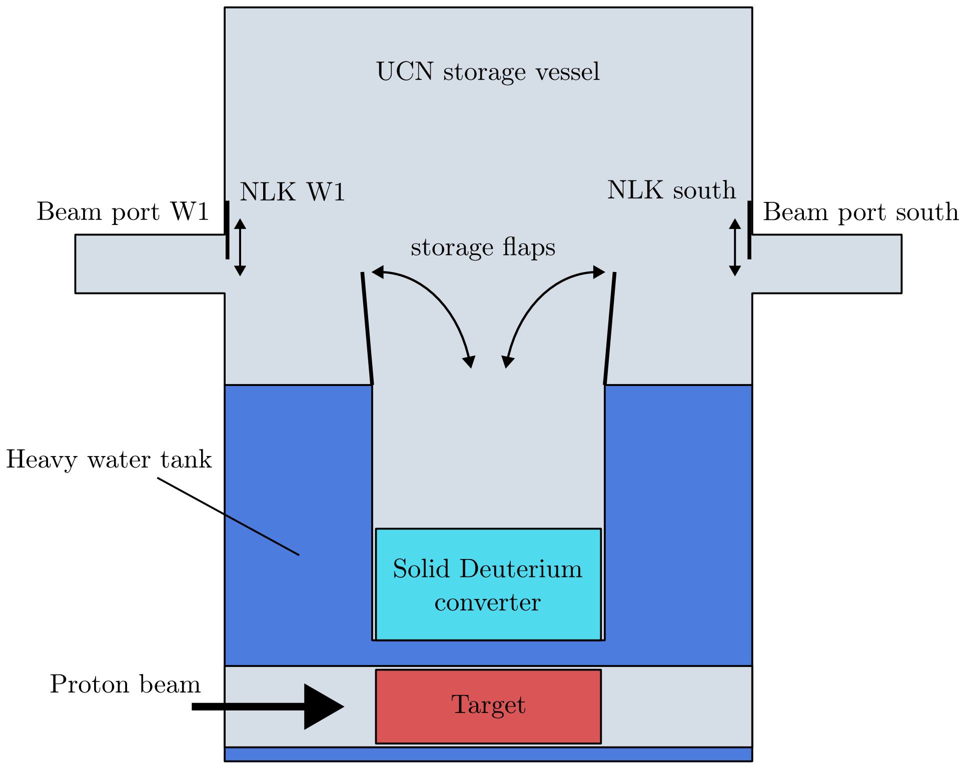

FIGURE 3.6: Simplified schematic cross section of the PSI UCN source (not to scale). Neutrons are produced by a 590 MeV proton beam generating spallation neutrons in the target. The neutrons thermalize in the heavy water tank and are subsequently converted to the UCN regime in the solid deuterium converter. While the proton beam is on target, the storage vessel flaps remain open, allowing UCNs to fill the storage vessel. After the proton pulse ends, the storage flaps close. UCNs are delivered to experiments through beam ports west 1 (W1) and south. Flaps NLK W1 and south can be used to block the respective beam port.

Fig. 3.6. The source is operated in a pulsed mode. For each pulse, the proton beam is guided onto the spallation target for up to 8 s by a kicker magnet. While the beam is on target, fast neutrons with kinetic energies in the MeV range are produced by spallation. In a first step, the fast neutrons undergo thermalization in the heavy water tank of the source, where they lose kinetic energy by elastic collisions with deuterium nuclei. As a result, the neutrons reach the thermal neutron regime corresponding to kinetic energies on the order of 10 meV. In a second step, thermal neutrons are converted into cold and ultracold neutrons in a solid deuterium converter, which is cooled to approximately 5 K and forms a crystal. First, thermal neutrons are converted to cold neutrons by elastic scattering from the cold deuterium atoms. In a subsequent step known as superthermal cooling [GP77], cold neutrons can excite phonons in the solid [GB83; YMG86], thereby losing almost all their kinetic energy and reaching the ultracold neutron regime. After exiting the solid deuterium converter, the neutrons enter the UCN storage vessel. Once the proton beam is no longer on target, flaps at the bottom of the storage vessel are closed, preventing the stored neutrons from re-entering the converter. From the storage vessel, UCNs are delivered to experiments through beam ports west 1 (W1) and south. Both beam ports are equipped with an additional flap referred to as the NLK ("NeutronenLeiterKlappe"), which allows the respective beam port to be closed, thereby increasing the UCN yield at the other port. The source can be pulsed with a minimum period of 300 s. Currently, the n2EDM experiment is connected to beam port south, while τSPECT is operating at beam

port W1. To accommodate this configuration, three different modes of cooperation have been explored:

1. Exclusive use: Only one NLK is open.
2. Shared use: Both NLKs are open.
3. Delayed use: The NLK status is switched after the experiment with higher priority has been filled.

While the first and second modes have both been employed during the 2024 beam time, the third mode has only been tested towards the end of the year and still requires upgrades to the beam port flap control electronics for use in standard operation. In addition to the cooperation mode, it must also be considered that the yield of the solid deuterium UCN source is known to slowly decrease over time [Ang+18]. This behavior has been attributed to surface changes of the deuterium crystal that accumulate over many pulses. To counteract this effect, a so-called conditioning of the deuterium crystal is performed in regular intervals. During conditioning, the deuterium crystal is slightly heated by reducing the cooling power over a small number of accelerator pulses. The heat deposited during these pulses removes a rough layer of solid deuterium from the crystal surface and subsequent freezing restores the crystal to its initial state prior to beam operation [Doo25]. However, while conditioning is ongoing, the UCN yield decreases substantially. For all measurements presented in this thesis, the source was operated with a 24 h conditioning period. This has previously been found to provide a good trade-off between the aforementioned effects [Doo25].

3.5 Data acquisition software

To perform a measurement with τSPECT, a large set of parameters across various subsystems must be configured. In 2023, the τSPECT experiment was moved from the research reactor TRIGA Mainz to the UCN source at PSI. In Mainz, most parameter changes required physical user intervention, with little to no automation available. Typically, the control software for each device consisted of a user interface that displayed the current set of parameters and allowed the user to set new values. However, no communication between different subsystems was possible. Parameters used during measurement runs were typically reconstructed from subsystem log files and written to a common database via an auxiliary script. All user interfaces and device control software were implemented in Python 3. The TRIGA reactor in Mainz is typically operated only during regular working hours. Thus, an experimenter was usually present on site throughout operation. In contrast, at PSI, the source is operated continuously, 24 h a day and seven days a week for several months. Additionally, the number of available on-site experimenters is significantly reduced due to the distance between PSI and Mainz. As a result, the level of automation, monitoring capabilities and remote operability needed to be increased substantially.

3.5.1 Custom Python package

To meet the requirements listed above, a Python package was developed that could be integrated into the existing control software. The primary purpose of the package is to enable subsystem control software to send status messages, updated sensor measurements or parameter values to other Python scripts running within the experiment's local network. This was achieved by employing a Python 3 remote procedure call (RPC)

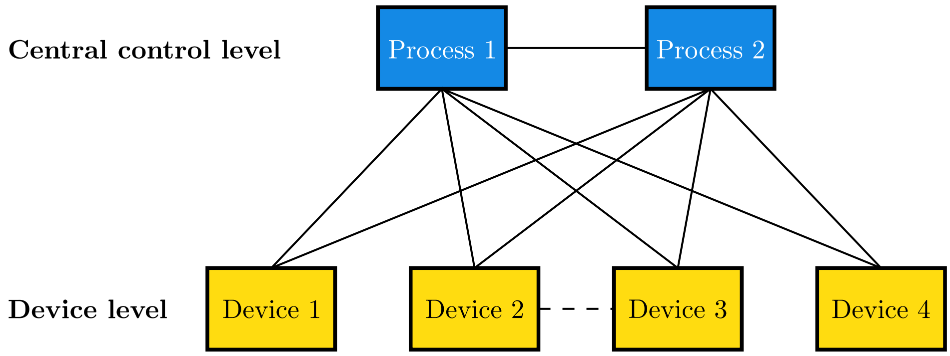

FIGURE 3.7: Basic structure of the control software network. For each device a corresponding class is implemented in Python and an object of that class is executed on one of the experiment's computers. These objects maintain the connection to the respective device and are able to query and set device parameters. By exposing object methods to the control software network, other objects on the network can set device parameters via network. Objects on the network can subscribe to updates of selected device parameters. Through central control processes, the experiment runs can be monitored and automated. Device objects can also access other device objects through the network. However, to reduce network complexity, this is only done if strictly necessary.

library[2]. After initialization, this RPC framework allows previously defined methods of objects in other Python processes on the local network to be invoked as if they were part of the same process. This capability was used to set up a subscription-based communication framework. Within this framework, each object in the control software can designate a set of parameters for subscription. Whenever one of these parameters changes, the new value is automatically distributed to all subscribed processes. To receive updates, objects on the network can register themselves as subscribers. Based on these features, the τSPECT control software network was structured as illustrated in Fig. 3.7. To integrate a device into the network, a corresponding device class is implemented in Python and an object of that class is executed on one of the experiment's computers. Each of these objects maintains a connection to the respective device, tracks all relevant device parameters and provides methods to access or modify them. Using the custom Python package these methods can be exposed to the network. This way, other objects can query and set device parameters simply by calling the respective method. Moreover, the subscription mechanism described above allows objects to receive parameter or status updates from other objects without periodically querying them. In the same way, virtual components such as databases can also be integrated into the network. The framework, however, is not designed for transferring large amounts of data. Therefore, only metadata such as measurement parameters should be distributed, while large data sets such as detector raw data should be transferred through other channels. Through central control processes that aggregate information from all devices, the experiment can be monitored and automated. These processes are described in detail in the following section. In principle, information can also be exchanged directly between devices, but to reduce network complexity this is only done when strictly necessary. In summary, a dedicated, RPC-based Python package was developed to facilitate communication between the data acquisition (DAQ) software and the control software of the experiment subsystems. The package makes it possible to maintain a modular control software structure, which is easily expandable and enables a high degree of automation of experiment operation.

[2]Python package: Pyro5 (v. 5.14)

3.5.2 Central control software

To monitor and automate τSPECT measurement runs, two central control systems were developed: the scheduling and the monitoring system. The scheduling system functions as a queue in which measurement configurations can be stored and subsequently executed in sequential order. Each measurement configuration specifies the parameters for:

- the two frequency generators feeding the amplifiers of SF1 and SF2,
- the two motors operating the translation stages of UCN detector and telescopic guide section,
- the timing control box.

For safety reasons, certain parameters such as the height of the experiment or the current of the superconducting coils still require user intervention and cannot be modified by the scheduling system. All subsystems of the experiment that are required to perform a time-controlled action during a measurement run can be either edge- or level-triggered by 5 V logic signals. In the edge-triggered case, the action is initiated on the rising edge of the signal, whereas in the level-triggered case, the action is performed for the duration of the high signal level. The timing of these signals is controlled by the timing control box, which provides a large number of 5 V logic output channels and is described in detail in Section 3.6. Upon receiving the trigger from the UCN source, the timing control box executes a predefined sequence of logic signals. Execution of a measurement configuration from the scheduling system queue proceeds as follows:

1. Load configuration from queue.
2. Verify that all devices are online.
3. Set measurement parameters on devices.
4. Instruct motors to move to their initial position.
5. Wait until motors reach their initial positions.
6. Enable timing control box to accept trigger from UCN source.
7. Wait for notification from timing control box that trigger has been received.
8. Disable trigger acceptance in timing control box.
9. Wait for notification from timing control box that the measurement run has been completed.
10. Send parameters of completed measurement to measurement database.

Throughout all steps, the scheduling system heavily relies on the control software network described in the previous section.

The scheduling control system is complemented by the monitoring system, whose primary function is to continuously monitor the experiment during measurement runs and to detect irregular conditions. Runs identified as irregular can be flagged accordingly and can subsequently be excluded from data analysis. A dedicated user interface serves as the central access point for experimenters to monitor the progress of ongoing runs and to review logs of status updates as well as warning and error messages. The monitoring system is subscribed to virtually all available parameters on the control software network. This enables cross-checking between different subsystem configurations to identify potential issues, such as a motor failing to move at its designated time. Such problems have occurred in the past and often went unnoticed for hours, or required lengthy debugging, even when an experimenter was on site during operation. Moreover, the system tracks whether a measurement is currently running and issues a warning whenever parameters are modified during these periods. Automating the detection of these problems significantly reduces the time required to diagnose and resolve issues with the experiment, thereby improving the data quality and increasing the amount of usable data.

The structure of the control software network further facilitates the integration of new interfaces, such as live plots of sensor values, which can be implemented simply by subscribing to the corresponding parameter updates. In addition, the central control software enables the experiment to operate for extended periods without requiring hourly manual intervention by experimenters. Through secure shell (SSH) access to the local network, the experiment can be operated fully remotely from Mainz during regular operation, significantly reducing the demand for on-site experimenters at PSI.

3.6 Timing control box upgrade

As explained in the previous section, all subsystems of the experiment that are required to perform time-controlled actions are triggered by 5 V logic signals provided by the so-called timing control box. The previous version of this box was designed for the early stage development of the experiment and was used until December 2023. It was based on a microcontroller[3] and provided a total of 24 output channels and 4 trigger inputs. During operation in 2022 and 2023, a few issues and limitations of this design became apparent. The most significant problem was related to the USB connection used to interface the box with the experiment control system. Over time, this connection proved to be unreliable, leading to intermittent connection losses that repeatedly interrupted data taking. Such failures typically occurred once or twice per week, which is too frequent for an experiment that is intended to operate continuously, 24 h a day and seven days a week. Since the root cause of the connection losses could not be conclusively identified, it was decided to upgrade the timing control box, with the primary change being a transition from USB to Ethernet communication. In addition, all input and output channels were galvanically isolated for the upgraded design in order to eliminate ground loops formed by the large number of coaxial cables connected to the box. Such ground loops had been found to introduce noise into the system, occasionally preventing subsystems from being triggered as intended. The upgraded hardware was designed and assembled by an electronics technician at the Institute of Physics at JGU and other members of the collaboration and provides a total of 22 output channels and 4 trigger inputs. To enable Ethernet communication, the microcontroller was replaced by a Raspberry Pi single board computer[4]. While a dedicated real-time solution would provide superior timing performance, the precision of the implemented system was evaluated experimentally and found to be sufficient for the current precision goals of the experiment (see subsequent sections). In the long term, the Raspberry Pi is intended to be replaced by a field programmable gate

[3] Arduino MEGA2560
[4] Raspberry Pi 4 model B

array (FPGA) system synchronized to a GPS-referenced clock[5]. A first attempt to implement the required functionality on an FPGA has already been made [Tsv22], but has so far not been successful. Therefore, the solution described in this section was adopted for the current experimental phase.

3.6.1 Timing implementation

One advantage of the Raspberry Pi is that it runs a full-fledged operating system (OS), which simplifies the communication with the timing control box considerably. Moreover, the OS time is synchronized to a Network Time Protocol (NTP) server provided by a GPS-referenced clock[6] on the local network with approximately millisecond precision. At the same time, however, the OS prevents the implementation of a true real-time system, as task scheduling and execution order are ultimately controlled by the OS. To nevertheless achieve precise timing, several system-level adjustments were implemented. First, one of the central processing unit (CPU) cores is isolated from the scheduler, ensuring that no OS processes are assigned to this core unless explicitly specified. In addition, the CPU clock frequency is fixed to its minimum value to avoid timing jitter caused by dynamic frequency scaling. Precise timing is achieved using the system timer registers. These include a 32-bit counter that increments at a frequency of 1 MHz, corresponding to a theoretical resolution of one microsecond. The timing process uses this counter directly as its time reference. Because the counter is only 32 bits wide and increments every microsecond, a counter rollover occurs approximately every 72 min. To handle this rollover safely, the counter must therefore be read at least once within each 72 min interval. In principle, precise timing could be achieved using a busy-wait scheme, in which the process continuously queries the counter until the desired wait time has elapsed. Since this method is limited only by CPU speed, it provides reliable timing precision. However, it also places a significant computational burden on the CPU. An alternative is to rely on blocking delay functions, which instruct the OS to suspend execution of the process for a specified duration. However, this method is unsuitable for precise timing, as the process is only guaranteed to be paused for *at least* the specified time, resulting in variable and potentially long delays. To combine the advantages of both strategies while avoiding their drawbacks, a hybrid approach was adopted. For the majority of the waiting period, blocking delay functions are used, followed by a busy-wait phase during the final 100 µs. This busy-wait interval serves as a buffer for cases in which the blocking delay duration is longer than the specified time. If the total wait time is shorter than 100 µs, the initial blocking phase is skipped entirely. The timing code was implemented in C++, using the usleep function for the blocking delay phase. This function accepts delay times specified in microseconds and is therefore well suited for this application. However, according to the specification [The04], it is only guaranteed to be precise for delays shorter than 1 s. For this reason, the initial sleep phase was subdivided into 1 s intervals. After each sleep interval, the remaining wait time is recalculated to prevent accumulation of timing inaccuracies. This strategy also ensures that the system timer counter is queried at least once every 72 min. In summary, the timing control process maintains precise timing by running on a CPU core isolated from the OS scheduler and using a combination of blocking delay functions and a busy-wait approach.

[5]Meinberg IMS Lantime M1000
[6]Meinberg IMS Lantime M1000

3.6.2 Trigger handling and performance evaluation

The logic input and output channels of the timing control box are provided by the general purpose input/output (GPIO) pins of the Raspberry Pi. Interfacing with these pins requires a dedicated software library. While the output channels only need to be switched between high and low logic levels, the input channels require more sophisticated handling. In particular, precise timing requires that the rising edges on input pins are detected with minimal latency. To achieve this, the library must provide a mechanism that allows the OS to suspend the execution of the calling process until a rising edge is detected, rather than relying on active polling. Although all tested libraries provided a function to wait for rising edges, their performance differed substantially. This discrepancy is likely due to differences in implementation, where some libraries repeatedly poll the pin state, whereas others request the OS to notify the process when an edge is detected. For this reason, the *gpiod* library was eventually selected. To evaluate the performance of the trigger handling, a simple test program was implemented. The program waits for a rising edge on an input pin and immediately sets an output pin to high once the trigger is detected. An oscilloscope[7] was used to measure the delay between the rising edges of the trigger signal and the corresponding output signal. The test was conducted before the upgraded timing control box was assembled and therefore characterizes only the performance of the Raspberry Pi itself. A total of 3471 samples were recorded and the resulting delay distribution is shown in Fig. 3.8. The measured delays ranged from 44 µs to 150 µs, corresponding to a mean delay of 64 µs and indicating significant jitter. The distribution is strongly asymmetric, exhibiting a sharp cutoff around 40 µs towards shorter delays and a long tail towards longer delays. In addition, three distinct peaks are visible around 60 µs. As the peak structure does not have a large influence on relative trigger timing between measurements, its origin was not investigated further. The long tail towards larger delays could, in principle, introduce relative timing shifts between individual

[7] Siglent SDS5104X

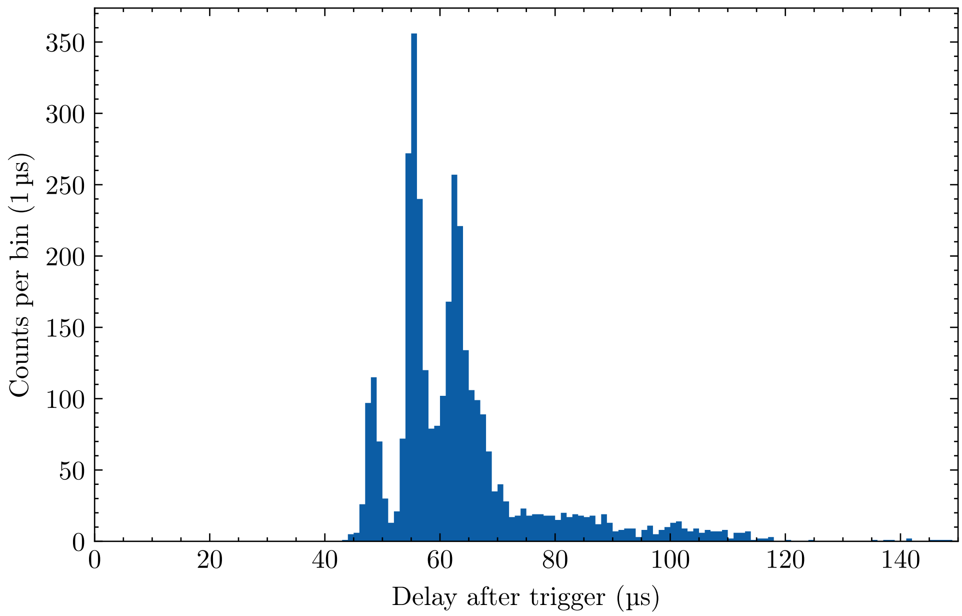

Figure 3.8: Distribution of the measured delay between rising edges of the trigger signal and the corresponding output signal generated by the timing control Raspberry Pi.

measurement runs. However, to assess its relevance, the magnitude of the observed jitter must be compared to the target precision of the experiment. Ultimately, the τSPECT experiment aims to measure the neutron lifetime with a precision of a few hundred milliseconds. Based on the recorded samples, the maximum delay is expected to remain below 1 ms. Timing effects of this magnitude, which are also likely to occur only sporadically, are negligible at the targeted level of precision. It was therefore concluded that the trigger performance satisfies the requirements.

3.6.3 Absolute timing accuracy

In addition to trigger performance, the absolute timing precision of the timing control box was investigated. At the time of this test, the upgraded timing control box was already in operation at the experiment. Therefore, a second Raspberry Pi of the same model[8] was configured identically to the one mounted in the timing box. Subsequently, the pulse width of rectangular pulses generated by one of the Raspberry Pi's GPIO pins was measured using an oscilloscope[9]. For this measurement, pulse widths between 1.8 s and 900 s were tested, with approximately 200 samples recorded for each setting. The oscilloscope has a memory depth of 250 million samples [Sig], corresponding to a sampling time of 4 µs or better for acquisition times of up to 1000 s. To ensure accurate timing of the recorded waveforms, the oscilloscope was phase-locked to a rubidium frequency standard[10] throughout the measurements. The results of the test are shown in Fig. 3.9. The upper panel displays the average offset between programmed and measured pulse width, while the lower panel shows the corresponding jitter, quantified by the sample standard deviation. As expected, the average offset increases with the pulse width. The longest neutron holding time typically used in the τSPECT experiment is about 2000 s. Extrapolating the observed trend to this timescale yields

[8] Raspberry Pi 4 model B
[9] Siglent SDS5104X
[10] Stanford Research Systems FS725

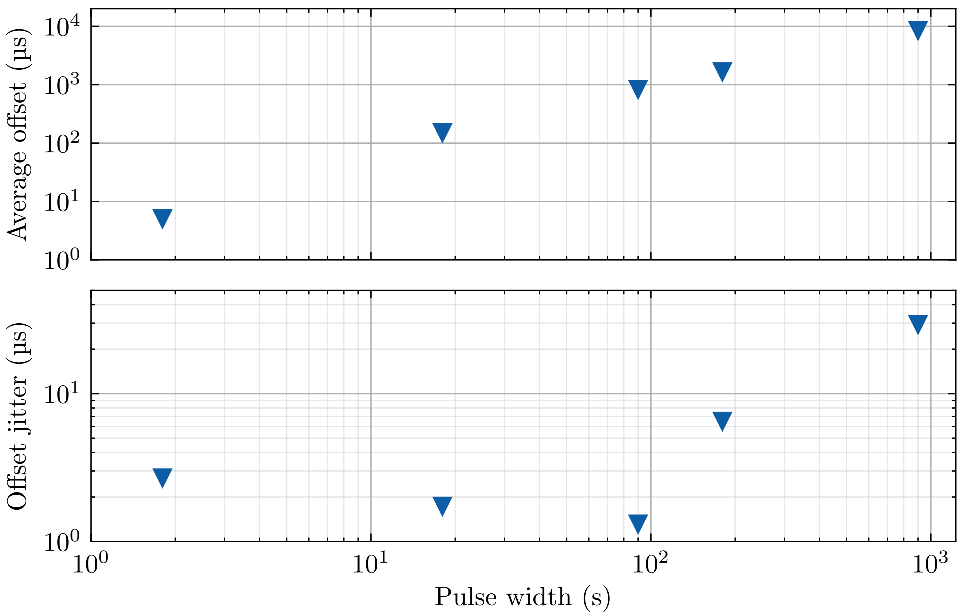

FIGURE 3.9: Top: average offset between the specified and measured pulse width of the timing control box. Bottom: corresponding offset jitter, quantified by the sample standard deviation.

an expected average offset of about 20 ms. Given the ultimate precision goal of a few hundred milliseconds, such an offset begins to constitute a relevant systematic effect. However, during the initial phase of the experiment, the goal is to obtain a first result with a precision of 1 s, for which this effect remains negligible. The measurements further indicate that the fluctuation of the offset is small. For pulse widths below 100 s, the observed jitter is likely dominated by the microsecond resolution of the timing counter itself. In principle, the offset could be corrected by characterizing the timing behavior of the installed timing control box using a procedure similar to that described here. However, it should be noted that the presented measurements were recorded over the course of a single week, limiting their relevance for assessing long-term clock drifts of the Raspberry Pi. Consequently, additional long-term studies would be required to implement a reliable offset correction.

3.6.4 Control software integration

To integrate the timing control box into the experiment control software network, a dedicated communication layer written in Python was added. Communication between the timing process and the communication layer is realized using shared memory and semaphores[11]. This design allows predefined signals and status information, encoded as integers, to be exchanged between the two processes. Whenever new information is written to the shared memory, the receiving process is notified via a semaphore shared between both processes. Timing sequences for upcoming measurement runs are transferred by first writing a list of channels and corresponding times to a file. Subsequently, the timing process is instructed to read and apply this configuration via the communication interface. Both the timing process and the communication layer employ a dedicated thread to handle the inter-process communication. For the timing process, this thread is explicitly assigned to a separate CPU core to avoid any impact on timing performance. Through this interface, status messages such as "trigger from UCN source registered" and "measurement run finished" can be forwarded to the control software network. Furthermore, the timing control box can be instructed not to accept any triggers in this way. This functionality is e.g. used to prevent triggers from the UCN source from being accepted while measurement parameters for the next run are still being transmitted to the subsystems.

3.7 Spectral cleaning study

As explained in Section 3.2, a so-called spectral cleaning procedure is employed in τSPECT to remove marginally trapped neutrons from the trapped neutron population. The performance of this procedure depends mainly on two parameters: the detector position during the cleaning phase and the duration of the cleaning phase. To evaluate the performance of different configurations, a dedicated series of measurements was conducted in fall of 2024, during which both cleaning position and cleaning duration were systematically varied. For each configuration a storage time constant τ_{st} was determined. As increasingly rigorous cleaning, i.e. positioning the detector deeper inside the trap or extending the cleaning duration, is expected to remove a larger fraction of the marginally trapped neutron population, the extracted storage time constant is expected to increase with increasing cleaning depth and duration. In the limit of very deep cleaning, virtually all marginally trapped neutrons are expected to be removed. At the same time, however, a very rigorous cleaning also removes a

[11]The posix_ipc (v. 1.1.1) package was used in the Python communication layer.

sizable fraction of storable neutrons, thereby increasing the number of measurement runs required to achieve a given statistical precision in a neutron lifetime measurement. Furthermore, it is possible that a fraction of the marginally trapped neutrons remains on trajectories that do not intersect the detector surface during the cleaning period. In this case, deeper cleaning may preferentially remove storable neutrons while leaving part of the marginally trapped population unaffected, thereby decreasing the ratio of storable to marginally trapped neutrons. As a result, the extracted storage time constant could decrease again. Depending on whether the cleaning duration is sufficiently long, the extracted storage time constant is therefore expected to either plateau or to exhibit a maximum as the cleaning depth is increased. Based on the results of such a study, suitable cleaning configurations, defined by the cleaning position and duration, can be selected for neutron lifetime measurement runs. Moreover, the measurements provide valuable input for fine-tuning and validating the τSPECT simulation framework.

To extract a storage time constant for a given configuration, it is in principle sufficient to measure two different storage times. In practice, however, this is generally not done, because this does not allow deviations of the decay curve from a single exponential decay behavior to be identified. Instead, storage curves consisting of approximately 5 to 10 different storage times are typically recorded. For the purpose of comparing different cleaning configurations, however, measuring only two storage times is sufficient, since systematic effects associated with marginally trapped neutrons can be identified by comparison with measurements obtained using other cleaning configurations. The primary advantage of this approach is that it minimizes the time required to reach a given level of precision, maximizing the number of cleaning configurations that can be tested within a fixed amount of available beam time.

Initially, five different cleaning positions between 1850 mm and 1760 mm were selected and storage time pair measurements were performed with a cleaning duration of 100 s. Additionally, measurements with shorter cleaning durations of 50 s and 30 s as well as a longer cleaning duration of 500 s, were conducted. For all measurements the short and long storage times were fixed to $t_\mathrm{s} = 50\,\mathrm{s}$ and $t_\mathrm{l} = 1250\,\mathrm{s}$, respectively (see Section 3.7.3).

3.7.1 Statistical model and considerations

To extract the storage time constant from measuring a pair of storage times, a constant trap escape rate $\lambda_\mathrm{st} = 1/\tau_\mathrm{st}$ is assumed, where τ_st is referred to as storage time constant. Under this assumption, the number of neutrons within the trap at time t is given by

$$N(t) = \hat{N}_0 \exp\left(-\frac{t}{\tau_\mathrm{st}}\right), \tag{3.1}$$

where $\hat{N}_0$ denotes the number of neutrons in the trap at the start of the storage period. From a statistical perspective, the storage process can be modeled as $\hat{N}_0$ independent Bernoulli trials, where the probability for a neutron to survive a storage period of duration t_s is given by

$$p_\mathrm{s} = \exp\left(-\frac{t_\mathrm{s}}{\tau_\mathrm{st}}\right). \tag{3.2}$$

Consequently, the probability to register k counts after the storage phase is described by the binomial distribution

$$P(N = k \mid \hat{N}_0 = n) = \binom{n}{k} p_{\mathrm{s}}^k (1 - p_s)^{n-k}. \tag{3.3}$$

This implicitly assumes that the initial number of neutrons $\hat{N}_0$ is known and constant. In the experiment, however, this condition is not fulfilled, as both UCN production in the source and the filling process exhibit intrinsic statistical randomness. During each pulse, a large number of neutrons traverse the deuterium crystal of the source, while only a small fraction being converted into UCNs. The population of the available UCN phase space can be modeled as an inhomogeneous Poisson point process under the following conditions [DV98]:

1. The number of UCNs in any finite phase space segment A is Poisson distributed.
2. The number of UCNs in disjoint phase space segments are independent random variables.

UCNs are produced at phase space coordinates ξ and time t according to a distribution $\Lambda(\xi, t)$. For a sufficiently small phase space time segment A, $\Lambda(\xi, t)$ can be approximated as constant, resulting in a constant local count rate. Under this condition, the number of UCNs in A is Poisson distributed, satisfying the first condition. Moreover, it is reasonable to assume that the production of a UCN in one segment A_i of the crystal does not influence production in other, disjoint segments A_j, such that the second criterion is also fulfilled. By definition of the Poisson point process, the total number N_{source} follows a Poisson distribution with mean

$$\mu_{\mathrm{source}} = \int \Lambda(\xi, t) d\xi dt. \tag{3.4}$$

A probability $p(\xi, t)$ to survive the filling process can be assigned to each phase space coordinate and point in time. The trap acceptance distribution is then given by

$$\Lambda'(\xi, t) = \Lambda(\xi, t) p(\xi, t). \tag{3.5}$$

The filling process therefore modifies the initial phase space distribution, without altering the underlying Poisson nature of the process. Thus, the combined production and filling process remains a Poisson process and the probability to initially store n neutrons is

$$P(\hat{N}_0 = n) = e^{-N_0} \frac{N_0^n}{n!}, \tag{3.6}$$

where

$$N_0 = \int p(\xi, t) \Lambda(\xi, t) d\xi dt \tag{3.7}$$

denotes the mean number of initially stored neutrons. The probability to count k neutrons after the storage phase is then given by

$$P(N = k) = \sum_n P(N = k \mid \hat{N}_0 = n)\, P(\hat{N}_0 = n). \tag{3.8}$$

It can be shown (see Appendix B.1) that this expression eventually simplifies to

$$P(N = k) = e^{-p_{\mathrm{s}} N_0} \frac{(p_{\mathrm{s}} N_0)^k}{k!}, \tag{3.9}$$

which is a Poisson distribution with mean $p_s N_0$. A similar argument can be made for the neutron counts recorded by the monitor detector. Accordingly, these are also assumed to follow a Poisson distribution.

In this study, pairs of short storage times t_s and long storage times t_l were measured. To apply the statistical model derived above, a constant N_0 must be assumed. However, as described in Section 3.4, the yield of the UCN source is known to slowly drift over time. Furthermore, pulse-to-pulse fluctuations, for example caused by fluctuations in the proton beam current, introduce further fluctuations in the neutron yield and consequently in N_0. To mitigate these effects, two complementary strategies are employed. First, measurements of each storage time pair are performed as close in time as possible, ideally using successive pulses, minimizing the impact of slow source yield drifts. Second, the measured neutron counts are normalized using the counts recorded by the monitor detector, which compensates for both residual drifts and pulse-to-pulse fluctuations. In combination, these measures substantially reduce source-related systematic effects. The magnitude of the residual uncertainty in N_0 after applying these corrections is subject of ongoing investigations and is further discussed in [Eng25]. In the following analysis, it is assumed to be negligible.

For short and long storage times t_s and t_l, the expected number of neutrons remaining in the trap is given by

$$N_s = N_0 \exp\left(-\frac{t_s}{\tau_{st}}\right) \tag{3.10}$$

and

$$N_l = N_0 \exp\left(-\frac{t_l}{\tau_{st}}\right) \tag{3.11}$$

respectively. In addition to stored neutrons, the detector also records background counts during the counting phase. The total number of counts is therefore

$$N_{meas} = N_{st} + N_{bg}, \tag{3.12}$$

where N_{st} denotes the stored counts and N_{bg} the background counts recorded during the counting phase. Assuming the background count rate to be constant over the period required to record one pair of short and long storage time measurements, N_{bg} is expected to follow a Poisson distribution, yielding

$$\Delta N_{meas} = \sqrt{N_{st} + N_{bg}} = \sqrt{N_{meas}} \tag{3.13}$$

for the uncertainty of the measured counts. That is, the uncertainty of the measured counts accounts for both the uncertainty of the stored counts and the uncertainty of the background counts. To estimate the number of stored neutrons N_{st}, the expected number of background counts

$$\bar{N}_{bg} = R_{bg}\, t_{count} \tag{3.14}$$

is subtracted from N_{meas}. Here, R_{bg} is the background rate determined from dedicated background measurements and t_{count} is the duration of the counting phase. The uncertainty of $\bar{N}_{bg}$ is then given by

$$\Delta\bar{N}_{bg} = t_{count}\Delta R_{bg}, \tag{3.15}$$

since t_{count} is known with sufficient precision that its uncertainty contribution is negligible. Accordingly, the estimated number of remaining neutrons after short and long storage times are

$$N_\mathrm{s} = N_\mathrm{meas,s} - \bar{N}_\mathrm{bg} \tag{3.16}$$

and

$$N_\mathrm{l} = N_\mathrm{meas,l} - \bar{N}_\mathrm{bg} \tag{3.17}$$

with uncertainties

$$\Delta N_\mathrm{s} = \sqrt{N_\mathrm{meas,s} + \Delta\bar{N}_\mathrm{bg}^2} \tag{3.18}$$

and

$$\Delta N_\mathrm{l} = \sqrt{N_\mathrm{meas,l} + \Delta\bar{N}_\mathrm{bg}^2}, \tag{3.19}$$

respectively. To correct for source-related fluctuations, the number of remaining neutrons after short storage, N_s, is scaled by the ratio of monitor detector counts for long and short storage times, $N_\mathrm{m,l}$ and $N_\mathrm{m,s}$ for each pair of measurements. The normalized short-storage counts are thus given by

$$\hat{N}_\mathrm{s} = N_\mathrm{s}\frac{N_\mathrm{m,l}}{N_\mathrm{m,s}}. \tag{3.20}$$

Typically, counts recorded by both the monitor detector and the τSPECT detector are well above 100. Consequently, the Poisson distribution is well approximated by a Gaussian distribution (see Appendix F) and the uncertainties can be computed via Gaussian propagation of uncertainties. The resulting uncertainties are therefore given by

$$\Delta\hat{N}_\mathrm{s} = \sqrt{\left(\frac{N_\mathrm{m,l}}{N_\mathrm{m,s}}\right)^2 \Delta N_\mathrm{s}^2 + \left(\frac{N_\mathrm{s}}{N_\mathrm{m,s}}\right)^2 N_\mathrm{m,l} + \left(\frac{N_\mathrm{s}N_\mathrm{m,l}}{N_\mathrm{m,s}^2}\right)^2 N_\mathrm{m,s}} \tag{3.21}$$

and ΔN_l for short and long storage time measurements, respectively.

3.7.2 Estimator evaluation

To extract a storage time constant from the measured data, three different estimators were evaluated. In the following, these estimators are referred to as the p-, λ- and Σ-estimators, respectively. For the p- and λ-estimators, a bias was initially found by other members of the collaboration, which was then investigated in greater detail and compared to a third estimator as part of this work. The p-estimator is the most straightforward estimator and can be derived by combining Eq. (3.10) and Eq. (3.11) to eliminate N_0 and solving for τ_st. This yields a storage time constant

$$\tau_\mathrm{p} = \frac{t_\mathrm{l} - t_\mathrm{s}}{\ln\left(\frac{N_\mathrm{s}}{N_\mathrm{l}}\right)}. \tag{3.22}$$

for each pair of long and short storage times. Subsequently, the mean $\bar{\tau}_\mathrm{p}$ of all computed τ_p values is calculated and used as an estimate of the storage time constant. In order to assess whether the estimators yield satisfactory results, a simple Monte Carlo model of the process was employed. The model assumes the storage time constant to be $\tau_\mathrm{MC} = 880\,\mathrm{s}$, which is of the same order as previous measurements of the neutron lifetime (see Section 2.3). This choice ensures that the magnitude of potential estimator biases is representative of what can be expected for the experimental data. To model the number of counts after long and short storage times, the expected number of counts

$$N(t_\mathrm{st}) = N_0 \exp\left(-\frac{t_\mathrm{st}}{\tau_\mathrm{MC}}\right) \tag{3.23}$$

was computed. In order to account for the statistical nature of the process, the measured number of counts was assumed to follow a normal distribution with mean $\mu = N(t_{\mathrm{st}})$ and standard deviation $\sigma = \sqrt{N(t_{\mathrm{st}}) + \bar{N}_{\mathrm{bg}}}$. For every run, a random sample was drawn from this distribution. Here, $\bar{N}_{\mathrm{bg}} = 169$ was assumed. This value was estimated via Eq. (3.14) by assuming a counting phase duration of 190 s, which was the default during the measurement campaign, and a background count rate of $R_{\mathrm{bg}} = 0.89\,\mathrm{s}^{-1}$, which was determined from a small number of preparatory measurements. Furthermore, to account for the slow drift of the UCN yield of the source, N_0 was assumed to be uniformly distributed in the interval $[1500, 1750]$. However, N_0 was assumed to be equal for each pair of short and long storage times and no normalization based on simulated monitoring detector counts was performed for this simulation. In the simulation runs, a fixed short storage time of $t_{\mathrm{s}} = 50\,\mathrm{s}$ was chosen and the long storage time t_{l} was varied. For each value of t_{l}, the storage time constant was estimated 100 times by sampling 150 000 pairs of counts after long and short storage times as described above. Subsequently, the storage time constant was computed according to Eq. (3.22) for each pair of samples and the mean calculated. The statistical uncertainty was estimated by computing the standard deviation of the 100 repetitions. For the p-estimator, a positive bias was observed, whose magnitude decreases as t_{l} increases. This behavior is shown in Fig. 3.10. The bias arises because the ratio $N_{\mathrm{s}}/N_{\mathrm{l}}$ is formed from two normally distributed random variables. Since division is a nonlinear operation, the resulting distribution of the ratio is asymmetric and computing the mean introduces a bias.

In an attempt to reduce the estimator bias, the λ-estimator was tested. Instead of computing τ_{p} for each pair of short and long storage times,

$$\lambda_{\mathrm{p}} = \frac{\ln\left(\frac{N_{\mathrm{s}}}{N_{\mathrm{l}}}\right)}{t_{\mathrm{l}} - t_{\mathrm{s}}} \tag{3.24}$$

is computed, which corresponds to the inverse of τ_{p}. Subsequently, the mean $\bar{\lambda}_{\mathrm{p}}$ of all computed λ_{p} values was calculated. The estimate for the storage time is then given by

$$\tau_{\lambda} = \frac{1}{\bar{\lambda}_{\mathrm{p}}}. \tag{3.25}$$

As shown in Fig. 3.10, this method exhibits a negative bias. In contrast to the p-estimator, the magnitude of the bias increases with increasing t_{l}, as expected from the fact that λ_{p} is the inverse of τ_{p}. Nevertheless, over most of the probed range, the λ-estimator performs better than the p-estimator. In order to identify an estimator that performs robustly across the entire range of typically probed storage times, a third estimator, referred to as Σ-estimator, was evaluated. Unlike the p- and λ-estimators, the Σ-estimator follows a different approach. Instead of computing an individual result for each pair, the counts measured after short and long storage times are summed separately. Subsequently, the ratio of the sums is formed. The storage time constant is thus given by

$$\tau_{\Sigma} = \frac{t_{\mathrm{l}} - t_{\mathrm{s}}}{\ln\left(\frac{\sum_{i=1}^{n} \hat{N}_{\mathrm{s},i}}{\sum_{i=1}^{n} N_{\mathrm{l},i}}\right)}. \tag{3.26}$$

This approach relies on the assumption that N_0 is identical for each pair of short and long storage time measurements. The estimator can be derived by profiling the

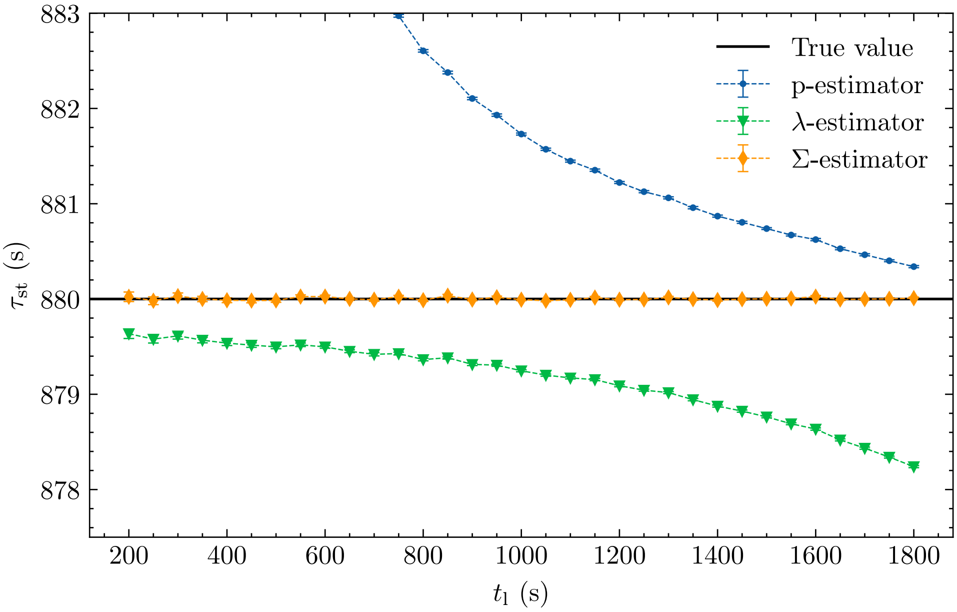

FIGURE 3.10: Comparison of different storage time constant estimators. A Monte Carlo simulation assuming a neutron lifetime of 880 s was used to generate counts after a fixed short storage time $t_s = 50$ s and a varying long storage time t_l. The different estimators were applied to the simulated data to extract the storage time constant.

likelihood function of this problem, which is shown in Appendix B.2. As for the p- and λ-estimators, the Σ-estimator was tested using the Monte Carlo simulation. The results are shown in Fig. 3.10 and indicate that the Σ-estimator does not exhibit a measurable bias at the 0.1 s level of precision probed by the simulation. For this reason, the Σ-estimator was used to determine the storage time constant throughout the following analysis.

3.7.3 Choice of storage times

In addition to selecting an appropriate estimator for the storage time constant, the pair of storage times used for this study had also to be determined. For this purpose, a second Monte Carlo simulation was performed with the goal of identifying the storage time combination that yields the smallest statistical uncertainty for a fixed amount of beam time. In this optimization, constraints imposed by the UCN source pulse period and the downtime of the experiment between subsequent measurements were considered. Throughout this study, the UCN source was operated with a pulse period of 400 s. A measurement run consisting of a 100 s cleaning period, followed by the storage period and a 190 s counting phase was assumed. Between consecutive measurements, approximately 24 s are required for the motors to move the UCN detector and spin flipping unit back to their respective initial positions. In addition, a few seconds are needed for the DAQ software to transmit and apply the parameter set for the subsequent run. Consequently, the experiment downtime was conservatively estimated to be 50 s. For the purpose of the simulation, a fixed available beam time window of 48 h was assumed. The short storage time was fixed at $t_s = 50$ s, as this had been the standard storage time during most of the commissioning measurements. This allowed for direct comparison of the configurations probed during this study with previously acquired data. For the long storage time t_l, values ranging from 400 s to

2400 s were probed in steps of 10 s. For each value, the number of measurement pairs that fit into the assumed available beam time window was calculated and counts were simulated in the same manner as for the estimator evaluation (see Section 3.7.2). In this case, however, a constant value of $N_0 = 1500$ was assumed. The storage time constant was subsequently estimated using the Σ-estimator and the corresponding statistical uncertainty determined via Gaussian uncertainty propagation. This yields

$$\Delta\tau_{\mathrm{st}} = \sqrt{\frac{(t_\mathrm{l} - t_\mathrm{s})^2 (N_{\mathrm{l,tot}}^2 \Delta N_{\mathrm{s,tot}}^2 + N_{\mathrm{s,tot}}^2 \Delta N_{\mathrm{l,tot}}^2)}{N_{\mathrm{l,tot}}^2 N_{\mathrm{s,tot}}^2 \ln\left(\frac{N_{\mathrm{s,tot}}}{N_{\mathrm{l,tot}}}\right)^4}}, \tag{3.27}$$

where $N_{\mathrm{s,tot}}$ and $N_{\mathrm{l,tot}}$ are given by

$$N_{\mathrm{s,tot}} = \sum_{i=1}^{n} \hat{N}_{\mathrm{s,i}} \tag{3.28}$$

and

$$N_{\mathrm{l,tot}} = \sum_{i=1}^{n} N_{\mathrm{l,i}} \tag{3.29}$$

respectively. Accordingly, the uncertainties of the total counts are given by

$$\Delta N_{\mathrm{s,tot}} = \sqrt{\sum_{i=1}^{n} \Delta \hat{N}_{\mathrm{s,i}}^2} \tag{3.30}$$

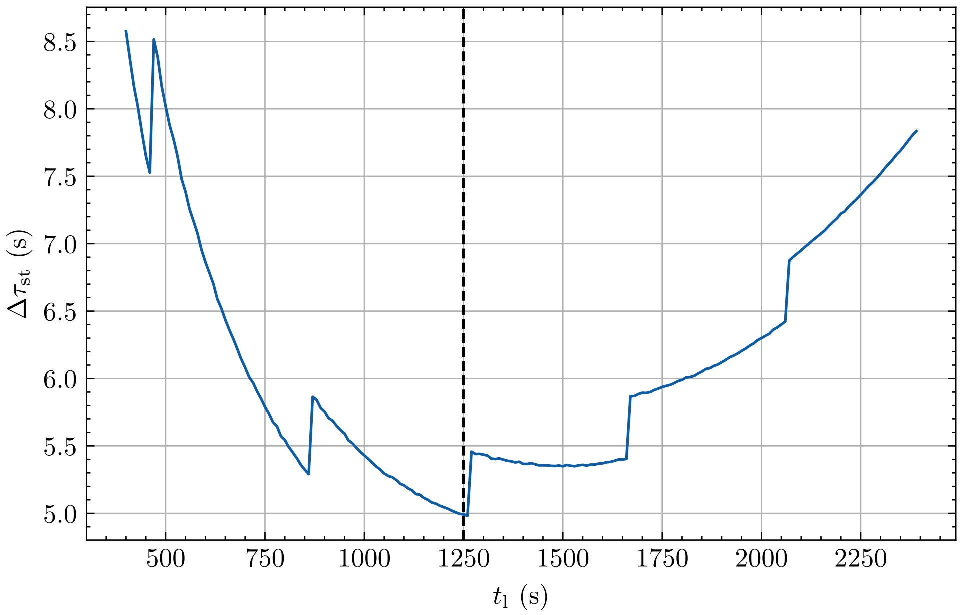

FIGURE 3.11: Storage time constant uncertainty $\Delta\tau_{\mathrm{st}}$ as a function of the long storage time t_l. A short storage time $t_\mathrm{s} = 50$ s is assumed. The uncertainty is calculated for a fixed available beam time of 48 h, assuming a pulse repetition period of 400 s. The dashed black line indicates the selected long storage time of 1250 s.

and

$$\Delta N_{\mathrm{l,tot}} = \sqrt{\sum_{i=1}^{n} \Delta N_{\mathrm{l,i}}^2}. \tag{3.31}$$

To reduce statistical fluctuations, the estimation was repeated 100 times for each pair of storage times and the average uncertainty $\Delta\tau_{\mathrm{st}}$ was calculated. The results are shown in Fig. 3.11. Based on these results, a long storage time of $t_{\mathrm{l}} = 1250\,\mathrm{s}$ was selected. The step-like features visible in the figure occur whenever an additional UCN pulse is missed during the long storage time measurement. This reduces the number of measurement pairs that can be completed within a fixed beam time window and therefore leads to an increased statistical uncertainty of the storage time constant.

3.7.4 Outlier identification and measurement pairing

After completion of the measurements, the dataset was filtered to exclude runs recorded under irregular conditions and valid pairs of short and long storage times were formed. Irregular measurement conditions most commonly occurred during UCN source conditioning (see Section 3.4) and when there were issues with the proton accelerator that initiates the spallation process in the UCN source. During source conditioning, the UCN yield of the source is significantly reduced. For this reason, τSPECT measurements performed during conditioning periods were excluded from the analysis. Accelerator related issues can likewise interrupt proton pulses or reduce the beam current, leading to reduced UCN yield. The corresponding measurements were therefore also excluded. The source performance was monitored using the monitor detector mounted upstream of τSPECT (see Section 3.1). Irregular pulses were identified by defining a threshold of monitor detector counts that a measurement had to exceed in order to be included in the analysis. This threshold was determined by plotting the temporal evolution of the monitor detector counts for each cleaning configuration. Subsequently, a count value that clearly separated measurements with normal source performance from those recorded under degraded source conditions was selected. Fig. 3.12 shows such a plot for the dataset recorded with a cleaning position of 1820 mm and a cleaning duration of 100 s. For some cleaning configurations, two different thresholds had to be defined because the source operation mode was changed from shared use to exclusive use for τSPECT partway through the measurement campaign. When the source feeds only a single beam port, the yield at that port increases significantly, necessitating an adjustment of the monitor detector threshold. As evident in Fig. 3.12, the applied threshold reliably identifies large outliers. However, measurements with slightly reduced monitor detector counts are consistently observed at the beginning of source conditioning periods. These cases are difficult to identify using a simple threshold cut. Because the number of such measurements is small and the normalization procedure is expected to compensate for most of the reduction, no further attempts were made to exclude them in this analysis. To enable improved identification of source status changes in future measurements, a new source beam port flap control system is currently under development. In addition to reporting the current position of the beam port flaps, it will provide explicit information on whether a source conditioning is ongoing. Furthermore, an improved outlier identification procedure is presently being developed within the collaboration.

In addition to source-related effects, the drift of the magnetic field produced by the main coils also had to be considered. As explained in Sec. 3.1, this required adjustments of the SF1 frequency every two to three days. Between frequency adjustments,

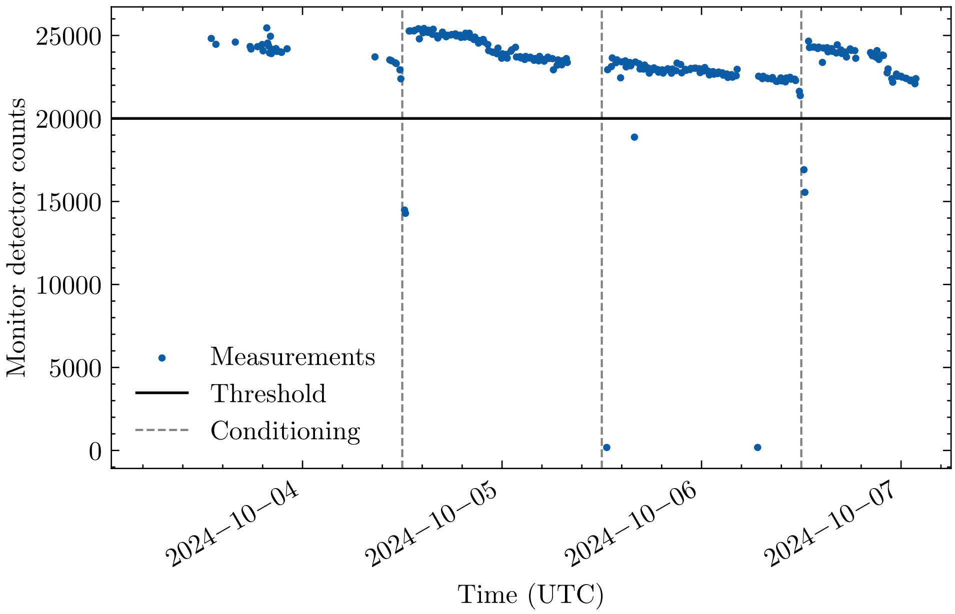

FIGURE 3.12: Temporal evolution of the monitor detector counts for the cleaning configuration with a position of 1820 mm and a duration of 100 s. The black horizontal line indicates the threshold used to determine which measurement runs were included in the analysis, while the vertical dashed lines indicate when a source conditioning was performed.

the number of stored neutrons gradually decreased as the optimum SF frequency slowly drifted away from the applied value. Although this effect is negligible for successive measurements, measurements recorded at different SF1 frequencies were not paired for the analysis. This avoids artificial jumps in the initial neutron number that occur whenever the frequency is adjusted. Additionally, a temporal filter was applied requiring that the two measurements forming a pair were recorded no more than 2000 s apart. This is necessary because the UCN source yield slowly decreases with each pulse between conditionings. During long storage time measurements, multiple source pulses can occur that are rejected by the experiment (i.e. the UCN shutter remains closed) but still decrease the UCN yield. For this reason, the measurement order within each pair was required to be (short, long) with respect to the recording time. This ensures that only measurements recorded during successive pulses are paired, minimizing the difference in UCN source yield.

3.7.5 Extraction of detector counts

After the measurement pairing procedure was completed, the detector counts recorded during the respective measurements were extracted. This was accomplished by determining the number of events in the detector raw data using a peak detection algorithm described in [Aul+24]. For this work, the slightly updated version 24 of the algorithm was used. This version was developed by other members of the collaboration and is optimized to suppress noise introduced by γ-radiation. The algorithm produces an output file containing a list of timestamps corresponding to registered events, referenced to the start of the measurement. Thus, the number of counts in a given time interval can be obtained by simply counting the timestamps within that interval. For all measurements, a time window of 190 s starting immediately after the end of the storage period was selected.

3.7.6 Background estimation

After extracting the raw detector counts, the expected number of background counts must be subtracted. To estimate the background count rate, the storage periods of the long storage time measurements were analyzed. During these periods, no neutrons enter the experiment from the source and the detector is located in a high magnetic field region. As a result, neutrons stored in the trap cannot reach the detector and only background is expected to be recorded. However, several short intervals within the storage period had to be excluded because of elevated γ-background levels that occur when the source is pulsed and when activated cooling water from the source passes through the cooling plant beneath the experiment after the pulse. This was first observed in detector data by other members of the collaboration. Because the source is pulsed with a fixed period rather than on demand, multiple pulses not associated with any τSPECT measurement can occur during the storage phase of long measurements. Owing to the fixed period of the pulses, the corresponding intervals of elevated background can be readily identified and excluded from the background rate estimation. Fig. 3.13 shows the stacked detector counts that occurred during the storage period of all long storage measurements with a cleaning duration of 100 s. The source is pulsed with a period of 400 s for a duration of 8 s and the activated cooling water passes beneath the experiment approximately 60 s after each pulse. To account for small timing jitter, a 30 s interval starting 10 s before the expected start of each elevated background period was excluded from the background rate estimation. Additionally, the first 10 s of the storage period had to be excluded. During this time, the detector is still moving from cleaning to storage position. Especially for deep cleaning positions, a small number of neutron counts can still be registered during this transition. Measures to reduce the elevated background are subject of ongoing investigations and will be discussed in a future publication. To account for potential long term drifts of the background rate, each pair of long and short storage times was

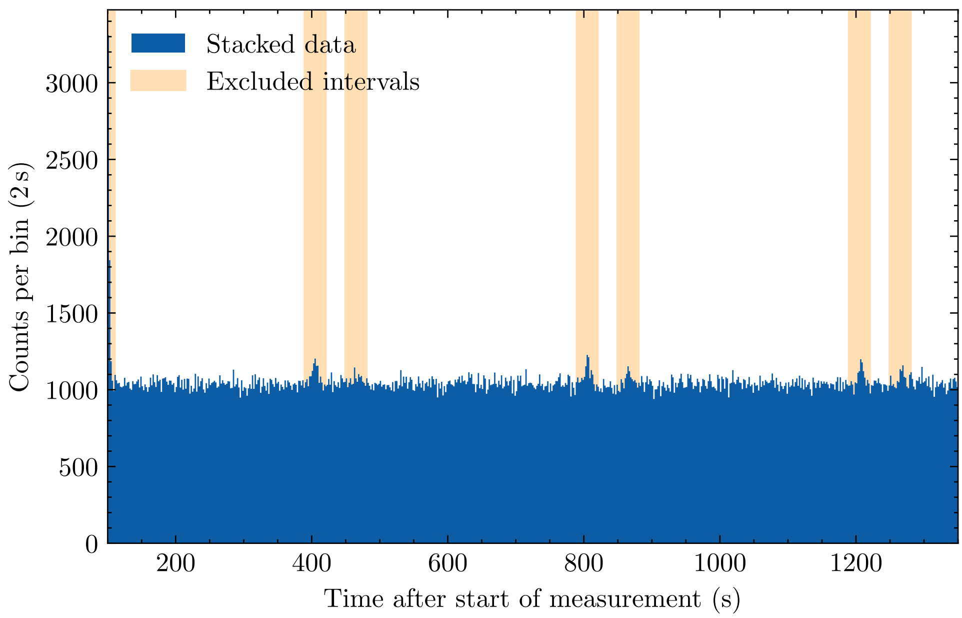

FIGURE 3.13: Stacked histograms of the storage periods of all 595 long storage time measurements recorded with a cleaning phase duration of 100 s as part of this study. Intervals highlighted in orange correspond to periods of elevated background and were excluded from the background count rate estimation.

assigned an individual background count rate determined from the storage period of the long storage measurement. This implicitly assumes that the background count rate remains constant over a period of approximately 2000 s.

3.7.7 Phase space corrections

In addition to the background correction, the temporal distribution of neutron events within the counting phase, in the following referred to as counting distribution, must also be taken into account. Up to this point, it was implicitly assumed that the time required to count the neutrons remaining in the trap after the storage period is negligible compared to the storage time constant. However, with a counting phase duration of 190 s and an expected storage time constant of approximately 880 s, this assumption is not justified. As a consequence, a neutron counted towards the end of the counting phase has a higher probability of having decayed than one counted at the beginning. This effect can only be neglected if the counting distribution shape is approximately identical for short and long storage times.

Construction of counting distributions

To assess whether the counting distributions after short and long storage are sufficiently similar, all measurements with short storage times and all measurements with long storage times were combined separately for each cleaning configuration and binned with a bin width of 1 s. The bin counts were assumed to follow a Poisson distribution. The bin width was chosen as a compromise between temporal resolution and statistical precision, ensuring a minimum bin content of 50. For bin counts above this threshold, a Poisson distribution can be reasonably approximated by a Gaussian distribution (see Appendix F), allowing bin count uncertainties to be treated as Gaussian in the subsequent analysis. The bins were then individually corrected for background and the neutron counts after the short storage time were rescaled by the ratio of the monitor detector counts as described above. Accordingly, the number of events counted in a given bin for short and long storage is defined as

$$\hat{N}_{\mathrm{s,B}} = \sum_{i=1}^{n} (N_{\mathrm{meas,s,b,i}} - \bar{N}_{\mathrm{bg,b,i}}) \frac{N_{\mathrm{m,l,i}}}{N_{\mathrm{m,s,i}}} \tag{3.32}$$

and

$$N_{\mathrm{l,B}} = \sum_{i=1}^{n} \left(N_{\mathrm{meas,l,b,i}} - \bar{N}_{\mathrm{bg,b,i}}\right), \tag{3.33}$$

respectively. Here n denotes the number of measurement pairs, $N_{\mathrm{meas,b,i}}$ is the number of events recorded in the respective bin for measurement i and $\bar{N}_{\mathrm{bg,b,i}} = R_{\mathrm{bg,i}} t_{\mathrm{bw}}$ represents the expected number of background counts within a bin of width t_{bw} for each pair of measurements. The uncertainty of $N_{\mathrm{l,B}}$ is given by the combined Poisson uncertainty

$$\Delta N_{\mathrm{l,B}} = \sqrt{N_{\mathrm{l,B}}}. \tag{3.34}$$

For the uncertainty of $N_{\mathrm{s,B}}$, however, the situation is more complex, since $N_{\mathrm{s,B}}$ is not simply a sum of Poisson distributed counts. Instead, each bin is corrected by a factor $c_i = N_{\mathrm{m,l,i}}/N_{\mathrm{m,s,i}}$ prior to summation. Moreover, $N_{\mathrm{meas,s,b,i}}$ can become very small towards the end of the counting phase, such that some bins contain only one or even zero counts. In this regime, the Poisson distribution describing the bin counts can no longer be reasonably approximated by a Gaussian distribution. Consequently, the uncertainty of $N_{\mathrm{s,B}}$ cannot, in general, be estimated using Gaussian propagation of uncertainty. To account for the Poisson nature of the underlying distribution, the variance of the product of two random variables was considered. For an independent normally distributed variable c_i and a Poisson distributed variable $N_{\mathrm{b,i}} = N_{\mathrm{meas,s,b,i}} - \bar{N}_{\mathrm{bg,b}}$ this yields

$$\Delta\hat{N}_{\mathrm{B,s}} = \sqrt{\sum_i \left(N_{\mathrm{meas,s,b,i}} c_i^2 + \Delta c_i^2 (N_{\mathrm{b,i}}^2 + N_{\mathrm{meas,s,b,i}})\right)} \tag{3.35}$$

for the uncertainty of the total bin counts (see Appendix B.3 for a derivation). The uncertainty of c_i can be obtained using Gaussian propagation of uncertainties, since the monitor detector typically records more than 20 000 counts per measurement run. At this level, the underlying Poisson distribution is very well approximated by a Gaussian distribution. This yields

$$\Delta c_i = \sqrt{\frac{N_{\mathrm{m,l,i}}}{N_{\mathrm{m,s,i}}^2} + \frac{N_{\mathrm{m,l,i}}^2}{N_{\mathrm{m,s,i}}^3}}. \tag{3.36}$$

Observed counting distribution shape differences

The resulting distributions for three selected datasets recorded with a cleaning duration of 100 s are shown in Fig. 3.14. The figure shows that the distribution of counts during the counting phase differs between short and long storage time measurements. This effect becomes increasingly pronounced for deeper cleaning. Fig. 3.15 presents the same datasets grouped by storage times. For long storage times, the distributions primarily exhibit a reduction in peak amplitude as the cleaning depth increases,

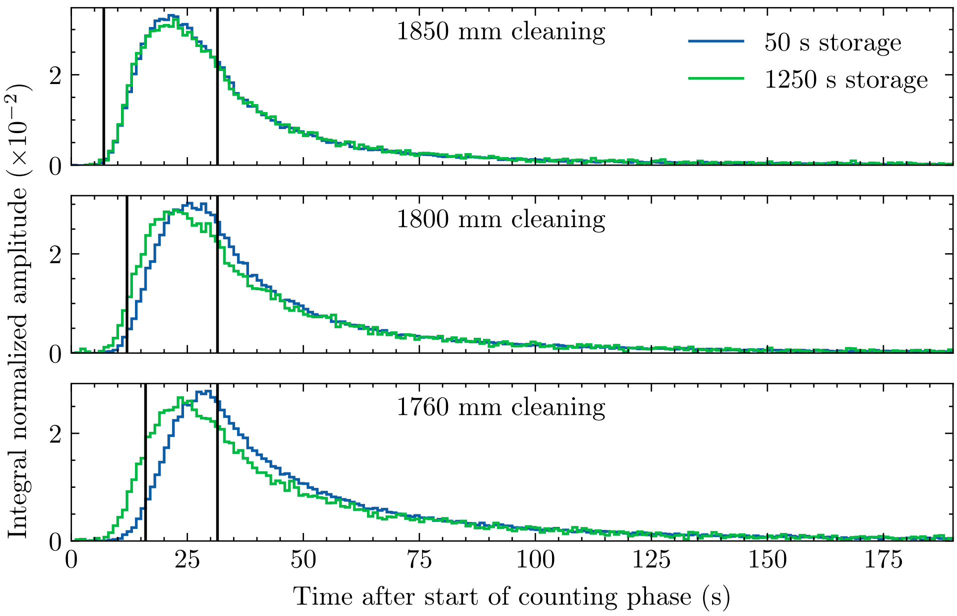

FIGURE 3.14: Counting distributions after short (50 s) and long (1250 s) storage for different cleaning positions. The cleaning duration is 100 s for all configurations shown. For each configuration, all corresponding measurements were combined to extract the distributions. Smaller cleaning position values correspond to deeper cleaning. The vertical black lines indicate the times at which the detector reaches the cleaning position and final counting position, respectively.

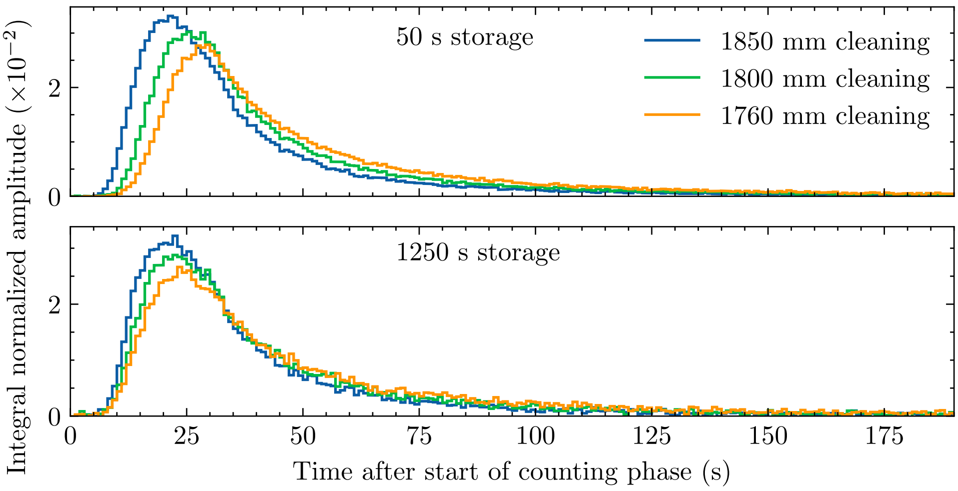

FIGURE 3.15: Comparison of the counting distribution for short (50 s) and long (1250 s) storage and different cleaning positions. The cleaning duration is 100 s for all configurations shown. Smaller cleaning position values correspond to deeper cleaning. All measurements for the respective configuration were combined to extract the histograms.

whereas an additional shift is observed for the short storage measurements. This indicates that the dominant contribution to the observed shape difference arises from the short storage measurements. An amplitude reduction for deeper cleaning is expected, since deeper cleaning preferentially removes high-energy UCNs, which are on average detected earlier during the counting phase. The observed temporal shift for the short storage times can likely be attributed to the removal of neutrons occupying the phase space volume probed while the detector sits in the cleaning position during the cleaning phase. While this volume has been repopulated after long storage times, the short storage time appears insufficient to achieve the same effect. It can therefore be concluded that a significant shape difference between the counting distributions after short and long storage times is present and must be corrected. Since neutrons are, on average, counted later during the counting phase after short storage, the correction is expected to reduce the extracted storage time constant.

To correct for this effect, the bin counts must be reweighted to account for the additional decay time of neutrons counted towards the end of the counting phase. The reweighted bin counts are given by

$$N_{\mathrm{w}} = N_{\mathrm{B}} \exp\left(\frac{t_{\mathrm{B}}}{\tau_{\mathrm{n,e}}}\right), \tag{3.37}$$

where N_{B} and t_{B} denote bin counts and bin center times, respectively and $\tau_{\mathrm{n,e}}$ is an estimate of the neutron lifetime. It should be noted that this correction requires assuming a value for the neutron lifetime, which is the quantity ultimately measured by the τSPECT experiment. It is therefore essential to ensure that the magnitude of the correction depends only weakly on the assumed neutron lifetime and remains well below the statistical uncertainty of the measurement. As established above, the dominant contribution to the difference between the counting distributions after short and long storage is a temporal offset, which can be corrected without assuming a neutron lifetime value. By separating the shape correction into an offset correction and a residual shape correction, the dependence on the assumed neutron lifetime can be reduced.

In addition to the aforementioned effects, the finite duration of the counting phase must be considered in context of the observed shape differences. However, a dedicated investigation of this effect would require measurements with an extended counting phase, which were not available for this study. Therefore, this effect was not quantified within the scope of this work and is not explicitly accounted for in the following analysis.

Offset correction

To determine the offset correction, the distributions after short (d_s) and long (d_l) storage were shifted relative to each other such that their overlap was maximized. For this purpose, smooth distributions were generated from the binned data using kernel density estimation[12] with a Gaussian kernel of width 1.5 s. The distributions were evaluated on a uniform grid of 1000 points t_i spanning the counting phase. The overlap between the two distributions was quantified using the sum of squared differences

$$\delta^2 = \sum_{i=1}^{1000} \left(d_\mathrm{s}(t_i) - d_\mathrm{l}(t_i)\right)^2 . \tag{3.38}$$

The optimal offset t_off was obtained by minimizing δ^2 using a numerical minimization package[13]. The uncertainty of t_off was estimated by randomly sampling the binned distributions by drawing values from Gaussian distributions with bin counts as mean and the corresponding bin uncertainties as standard deviations. For each sampled

[12]Python package: KDEpy.TreeKDE (v. 1.1.12)
[13]Python package: scipy.optimize.minimize (v. 1.15.2)

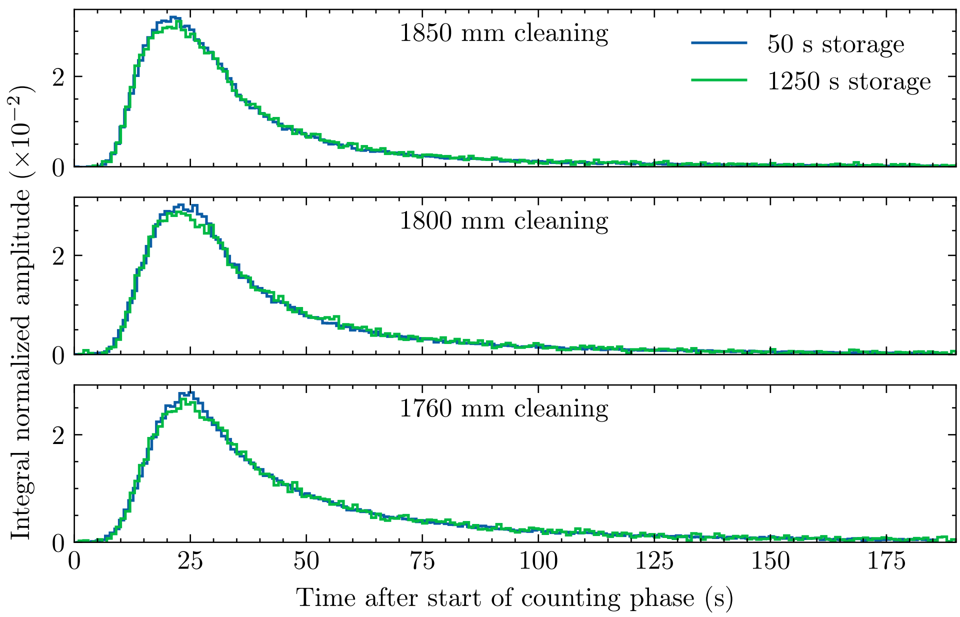

FIGURE 3.16: Offset corrected counting distributions after short (50 s) and long (1250 s) storage for different cleaning positions. The cleaning duration is 100 s for all configurations shown. For each configuration, all corresponding datasets were combined to extract the distributions. Smaller cleaning position values correspond to deeper cleaning.

distribution, the offset was recalculated as described above. This procedure was repeated 100 times and the uncertainty of t_{off} was estimated by computing the sample standard deviation of the resulting offsets. The resulting values are summarized in Tab. 3.1 and can be used to compute the offset corrected storage time constant, which is defined as

$$\tau_{\text{st,off}} = \frac{t_\text{l} - t_\text{s} + t_{\text{off}}}{\ln\left(\frac{\sum_{i=1}^{n} \hat{N}_{\text{s},i}}{\sum_{i=1}^{n} N_{\text{l},i}}\right)}. \tag{3.39}$$

Fig. 3.16 shows the counting distributions after applying the offset correction. After this correction, the distributions already exhibit a large overlap. However, the distributions after short storage still display a more pronounced peak, indicating a residual shape difference in addition to the corrected temporal shift.

Residual shape correction

To correct for the observed residual shape difference, the reweighting approach described above was applied to the offset corrected histograms. For this purpose, an estimated neutron lifetime of $\tau_{\text{n,e}} = 880\,\text{s}$ was used. After reweighting, the bin counts were summed to obtain the reweighted total counts $\hat{N}_{\text{s,w}}$ and $N_{\text{l,w}}$ for short and long storage times, respectively. Using these reweighted counts, the shape corrected storage time constant

$$\tau_{\text{st,sc}} = \frac{t_\text{l} - t_\text{s} + t_{\text{off}}}{\ln\left(\frac{\sum_{i=1}^{n} \hat{N}_{\text{s,w},i}}{\sum_{i=1}^{n} N_{\text{l,w},i}}\right)} \tag{3.40}$$

was computed. By comparing $\tau_{\text{st,sc}}$ to the offset corrected value $\tau_{\text{st,off}}$, the magnitude of the residual shape correction

$$t_{\text{rs}} = \tau_{\text{st,sc}} - \tau_{\text{st,off}} \tag{3.41}$$

Table 3.1: Parameters extracted from the cleaning position study. The parameters t_{off} and t_{rs} correct for differences in the shape of counting distributions after short and long storage times. $\bar{N}_\text{s}$ denotes the average number of counts after 50 s of storage for the respective cleaning configuration. All counts were normalized to 24 000 monitor detector counts before averaging and only measurements recorded with the UCN source operating in shared mode were included. The given uncertainty represents the sample standard deviation. $\tau_{\text{st,sc}}$ is the storage time constant corrected for shape differences between the counting distributions. Smaller cleaning position values correspond to deeper cleaning.

x_{clean} (mm)	t_{clean} (s)	t_{off} (s)	t_{rs} (s)	$\bar{N}_\text{s}$	$\tau_{\text{st,sc}}$ (s)
1850	100	−0.203(71)	0.7(2)	1870(140)	852.1(41)
1820	100	−1.501(69)	0.8(2)	1860(100)	865.5(39)
1800	50	−1.892(70)	1.2(2)	1830(30)	866.6(35)
1800	100	−2.519(87)	0.7(3)	1550(60)	866.6(45)
1800	500	−2.82(20)	1.0(6)	590(50)	882.4(98)
1780	100	−3.290(79)	1.0(2)	1140(90)	866.5(41)
1760	30	−1.739(74)	0.7(2)	1590(70)	857.2(34)
1760	50	−3.25(14)	1.4(4)	1250(40)	879.6(68)
1760	100	−4.313(85)	0.8(3)	940(30)	876.2(43)
1750	50	−3.493(83)	1.3(2)	1070(70)	877.8(41)

was determined. In order to estimate the uncertainty of the residual shape correction, the offset corrected histograms were sampled 1000 times, using the same procedure as for the uncertainty estimation of the offset correction. For each sampled dataset, the residual shape correction was recalculated and the sample standard deviation was computed as an estimate of the uncertainty of t_{rs}. The results are summarized in Tab. 3.1. Finally, the uncertainty of the shape corrected storage time constant

$$\Delta\tau_{\mathrm{st,sc}} = \sqrt{\Delta\tau_{\mathrm{st,off}}^2 + \Delta t_{\mathrm{rs}}^2} \tag{3.42}$$

was calculated. To assess the sensitivity of the residual shape correction to the assumed neutron lifetime, the correction procedure was repeated using $\tau_{\mathrm{n,e}} = 870\,\mathrm{s}$ and $\tau_{\mathrm{n,e}} = 890\,\mathrm{s}$. The resulting deviations were approximately one order of magnitude smaller than the estimated uncertainty of t_{rs}. This demonstrates that the residual correction is sufficiently insensitive to the choice of $\tau_{\mathrm{n,e}}$ for the present analysis.

Interpretation and implications for future measurements

In addition to their role in correcting the presented dataset, the shape corrections provide insight into both the underlying phase space dynamics and the requirements for future high-precision measurements of the neutron lifetime. Interestingly, the extracted residual shape corrections are largely consistent within their statistical uncertainties. This suggests that the cleaning procedure predominantly introduces a temporal offset rather than significantly altering the intrinsic shape of the counting distribution. Instead, an alternative explanation could be that certain regions of the phase space contributing to the long tail of the counting distribution are not populated during the filling phase and instead require a substantial fraction of the storage time to be populated by neutrons unaffected by the spectral cleaning. To test this hypothesis, the counting distribution should be studied for intermediate storage times not covered in this study. In particular, such a study could also inform a decision on whether choosing a longer short storage time could reduce the counting distribution shape differences to such an extend that the loss in statistical sensitivity is compensated for by the reduced magnitude of the required corrections. Moreover, it could also provide insight into the time constants associated with the phase space evolution within the trap.

In contrast to the statistical uncertainties of the shape corrections, the dependence of the residual shape correction on the choice of the assumed neutron lifetime $\tau_{\mathrm{n,e}}$ does not decrease with an increasing number of measurements. This underscores the necessity to reduce the magnitude of this correction for future high-precision neutron lifetime measurements to below the $0.1\,\mathrm{s}$ level for a next generation experiment. Achieving this goal requires faster counting or enhanced momentum randomization within the trap. In a τSPECT like geometry, a shorter counting phase could be realized by employing a detector capable of moving further into the trap region. This would increase the interaction rate of the remaining neutrons with the detector surface, reducing the required duration of the counting phase. Enhanced momentum randomization would allow the phase space volume depleted during cleaning to be repopulated more rapidly, reducing shape differences between the counting distributions after long and short storage. This could e.g. be achieved by optimizing the shape of the magnetic field in this regard in a future experiment. To assess this, detailed simulations of the phase space evolution for different field geometries are required.

3.7.8 Results and discussion

The storage time constants extracted for all investigated cleaning configurations are shown in Fig. 3.17. These include the corrections for the effects discussed above, with the exception of possible effects related to the finite counting phase duration, which as described above could not be quantified with the available dataset. The following discussion therefore assumes that contributions from this effect are subdominant compared to the statistical uncertainty of the extracted storage time constants. As expected, deeper and longer cleaning generally results in larger storage time constants, consistent with a more effective removal of marginally trapped neutrons. At the achieved level of statistical precision, no significant difference is observed between cleaning durations of 50 s and 100 s. Since a shorter cleaning period leaves a larger number of neutrons in the trap at the beginning of the storage phase, this favors the shorter cleaning period from a statistics accumulation perspective. However, more precise comparison measurements are required to draw a definitive conclusion. In contrast, a very short cleaning of 30 s, which effectively only covers the filling phase of a measurement run, yields a significantly shorter storage time constant. Similarly, a very shallow cleaning position at 1850 mm also results in a markedly reduced storage time constant. These observations indicate that, for both configurations, a substantial fraction of marginally trapped neutrons remains in the trap after the cleaning phase, rendering them unsuitable for future measurements. In addition to an active cleaning procedure, spectral cleaning can in principle also be achieved by allowing sufficient time for marginally trapped neutrons to leave the trap by themselves. To investigate this possibility, a long cleaning period of 500 s at a position of 1800 mm was tested. However, the number of neutrons counted after storage was significantly smaller than for any other configuration (see Tab. 3.1). Under these conditions, a substantially larger number of measurement runs would be required to reach the desired statistical precision. Consequently, the measurements for this configuration were

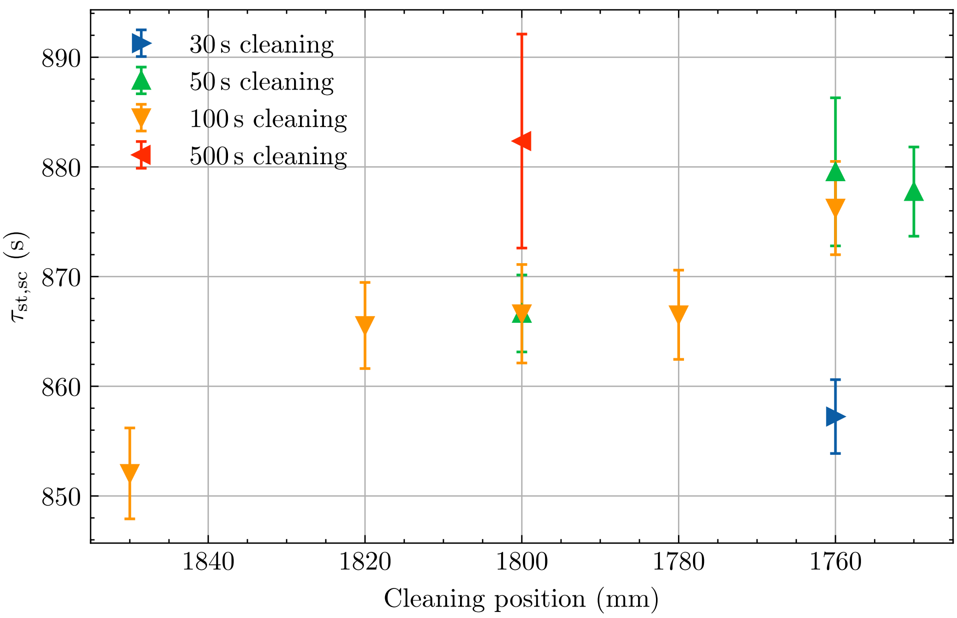

FIGURE 3.17: Storage time constants for different cleaning configurations. Smaller valued cleaning positions (to the right) correspond to deeper cleaning. The corrections for the the shape difference between the counting distributions after short and long storage derived in this section have been applied to all data points.

terminated before a statistical precision comparable to the other configurations could be achieved. Although the extracted storage time constant appears larger than those obtained with shorter cleaning durations at the same position, the large statistical uncertainty resulting from the reduced number of neutron counts renders this comparison inconclusive.

The most notable result of the study is the trend of the 50 s and 100 s cleaning data points in the range between cleaning positions 1820 mm and 1750 mm. While the trend between 1820 mm and 1780 mm suggests that a plateau in the storage time constant may have been reached, an additional increase is observed between 1780 mm and 1760 mm. However, at the present level of precision, this apparent increase could still be attributed to statistical fluctuation. These results therefore call for a dedicated follow-up measurement campaign with improved statistical precision, focusing on the cleaning position range between 1800 mm and 1750 mm. Moreover, complementary measurements employing storage times significantly different from 50 s and 1250 s for a selected subset of cleaning configurations are required. Such measurements would provide an important consistency check by testing whether the extracted storage time constants are independent of the chosen storage times and could thereby provide evidence for the effectiveness of the corresponding cleaning configuration. Furthermore, the same configurations can be investigated using the simulation framework described in [Aul+25] to assess whether the observed trend can be reproduced. This would provide an important cross-check of the phase space dynamics of marginally trapped neutrons modeled in the simulation.

3.7.9 Summary and outlook

As part of this study, different cleaning configurations in τSPECT were systematically investigated by measuring storage time constants for the first time. Although a cleaning configuration suitable for a high-precision determination of the neutron lifetime could not be identified conclusively, the study provided valuable guidance by narrowing the range of configurations that should be investigated further at higher precision in subsequent measurement campaigns. Additionally, differences in the counting distributions after short and long storage times were observed and strategies for correcting these effects were established. However, additional measurements with longer counting phase durations are required to allow for a dedicated study of potential effects arising from the finite counting phase duration in the context of the observed differences.

The observed counting distribution shape differences indicate that significant changes occur in the neutron phase space distribution over the probed storage times. Based on this finding, additional systematic effects unrelated to the counting distribution shapes may arise. For example, differences in the spatial distribution of neutron counts on the detector surface could lead to systematic distortions of the extracted neutron counts if spatial variations in detection efficiency are present. Investigating such effects requires a well characterized, spatially resolving detector. In combination with the magnetic field of τSPECT, such a detector would also enable the extraction of information about the stored UCN energy spectrum and the neutron phase space distribution in the trap, since the potential varies across the detector surface depending on the shape of the magnetic field. This capability could e.g. be used to monitor the evolution of the spatial distribution of neutron events on the surface and thus the phase space distribution over many neutron pulses. In this way, potential long-term changes in the energy spectrum of the UCN source during extended beam times, which might become relevant for τSPECT, could be identified. Moreover, spatial resolution would allow neutron counts on detector regions located at a trap potential level above the nominal trap depth to be rejected. Neutrons with sufficient kinetic energy to reach these positions can be classified as marginally trapped neutrons, providing an additional handle for suppressing or studying this systematic effect. The required level of spatial resolution to achieve these objectives is estimated in Section 3.8.

3.8 Spatial resolution requirements

To assess the level of spatial resolution required to extract useful information with a spatially resolving detector, the trap potential, the experimental setup and the obtainable number of neutron counts were considered. Given that the trap depth in the τSPECT experiment is approximately 36 neV [Aul+25], an energy resolution at the level of approximately 1 neV is a reasonable target. This would provide sufficient resolution across the stored energy spectrum without imposing overly stringent requirements on the number of counts per energy spectrum segment. The spatial resolution required to achieve this precision depends on the trap potential gradient across the detector surface. In τSPECT, the maximum gradient of the radially confining magnetic field is approximately 40 mT/mm according to magnetic field simulations performed with the τSPECT simulation framework described in [Aul+25], which corresponds to a potential gradient of 2.4 neV/mm. However, to make practical use of high spatial resolution, a sufficient number of neutrons per surface area element associated with that resolution is required. In τSPECT, the detector has a circular cross section with a diameter of approximately 100 mm, corresponding to an area of approximately $A = 7850\,\mathrm{mm}^2$. The average number of counts per spatial bin on the surface is then given by

$$\bar{n}_\mathrm{b} = \frac{N}{A} x_\mathrm{b}^2, \tag{3.43}$$

where N is the total number of counted neutrons and x_b is the bin size. Solving for the bin size yields

$$x_\mathrm{b} = \sqrt{\frac{\bar{n}_\mathrm{b} A}{N}}. \tag{3.44}$$

Thus, identifying a meaningful bin size requires choosing a reasonable value for $\bar{n}_\mathrm{b}$, which depends on the desired statistical sensitivity. Assuming the bin counts to be Poisson distributed, the relative uncertainty is given by $\sqrt{\bar{n}_\mathrm{b}}/\bar{n}_\mathrm{b}$. For a set of common choices for $\bar{n}_\mathrm{b}$, corresponding to relative uncertainties between 10 % and 30 %, the statistically meaningful bin size is shown in Fig. 3.18 as a function of the total number of detected neutrons. In a typical measurement with τSPECT, approximately $N = 1500$ neutrons are counted after a short storage time of 50 s, whereas only about one quarter of this number is counted after a long storage time of 1250 s. Therefore, multiple measurements must be stacked to extract meaningful information about the stored neutron ensemble in a given measurement configuration. In practice, the number of measurements per day conducted with the same experimental configuration can vary significantly and depends strongly on how many different cleaning configurations and storage times are measured during the respective period. To maintain sufficient temporal resolution for identifying potential long-term changes in the phase space properties and the energy spectrum of the trapped ensemble, it is expected that, at most, a number of measurements equivalent to a few days will be stacked. In the statistically most favorable case of measuring pairs of short and long storage times,

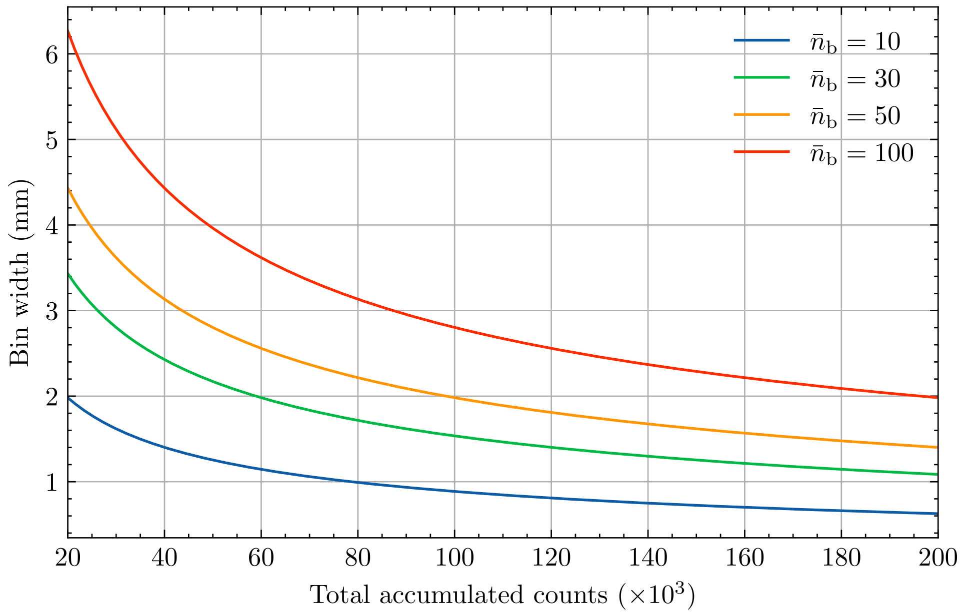

FIGURE 3.18: Required bin width for achieving a given number of counts per bin as a function of the total number of accumulated counts, shown for a set of common choices for the typical number of counts per bin $\bar{n}_\mathrm{b}$. The selected counts per bin correspond to relative uncertainties per bin between 10 % ($\bar{n}_\mathrm{b} = 100$) and 30 % ($\bar{n}_\mathrm{b} = 10$).

as done during the cleaning study presented in the previous sections, about 40 measurements per day can be expected for each configuration. Using data accumulated over three days would therefore yield 120 000 neutron counts. Based on Fig. 3.18, a bin size between 1 mm and 2 mm can be expected to yield statistically meaningful insights. This would correspond to an energy resolution between 2.4 neV and 4.8 neV in the regions of steepest gradient. To achieve this, the detector spatial resolution must be smaller than the expected bin size. It should be noted, however, that this estimation assumes a uniform distribution of counts across the detector surface, which is not expected for physical energy distributions in the trap. Furthermore, the estimation only considers the statistically ideal case of short storage time measurements when short and long storage times are measured pairwise. Therefore, this estimate should be treated as an upper limit estimate for the required spatial resolution. In conclusion, while a slightly higher spatial resolution would be desirable from the perspective of energy resolution, the actually achievable energy resolution and therefore the required spatial resolution is currently limited by the obtainable neutron counts per measurement. The development and characterization of such a suitable spatially resolving detector is subject of the subsequent chapters.

Chapter 4

Detector concept and design

4.1 UCN detection

Because neutrons are electrically neutral and interact only weakly with atomic electrons, they generally do not directly ionize matter. Since detector systems typically rely on converting particles into electrical signals, neutron detectors generally operate by detecting ionization produced by secondary particles generated in neutron-matter interactions [Kra88, p. 451]. In the following, the general principles of neutron detection are briefly outlined and their implications for detector design are discussed in the context of the requirements motivated by the τSPECT experiment. While fast neutrons with kinetic energies in the MeV range produce recoil particles in elastic collisions with nuclei that can subsequently be detected, the kinetic energy of slow neutrons in the thermal, cold or ultracold regimes is insufficient for this mechanism. Instead, detectors for slow neutrons rely on radiative capture reactions, nuclear reactions emitting charged particles or nuclear fission [Pie+20]. For most applications, nuclear reactions that generate charged particles are preferred, as the resulting ionization signals can be detected with high efficiency and allow for good discrimination against γ-backgrounds, which is not necessarily the case for radiative capture reactions. Fissile materials, on the other hand, are typically avoided due to their natural radioactivity, toxicity and regulatory restrictions [Pie+20]. Suitable materials are characterized by a large reaction cross section, a high energy release and practical availability. Based on these criteria ^{3}He, ^{6}Li and ^{10}B are the most common choices [Pie+20] and frequently used as conversion materials in UCN detectors (see e.g. [Wan+15; Ban+16; And+17]). In addition to the choice of conversion material, the neutron optical properties of material surfaces must be considered in the context of UCN detection. This is particularly important for gas-based detectors, which require a separation window to isolate the detector gas from the experiment vacuum. A common mitigation strategy is to insert a vertical UCN guide section between the experiment and the detector in order to gravitationally accelerate the neutrons towards the detector. However, such guide sections also introduce additional transport losses. To reduce these losses and avoid separation windows altogether, many UCN detectors employ solid conversion layers based on ^{6}Li and ^{10}B, which provide high absorption cross sections and can be directly exposed to vacuum. In the case of ^{10}B, such layers can also provide favorable neutron optical properties (see e.g. [Wan+15]). Solid conversion layers do, however, introduce additional efficiency losses, due to the finite range of the nuclear reaction products within the conversion material. To minimize these losses, the layer thickness must be matched to the expected UCN energy spectrum. For the case of ^{10}B in the context of the detector prototype presented in this chapter, this topic is discussed in detail in Section 4.4.1. An additional advantage of detectors employing solid conversion layers is that they can be inserted directly into UCN storage volumes for in-situ detection, thereby eliminating transport losses entirely.

4.2 Spatially resolved UCN detection

In the context of the detector development discussed in this thesis, it is instructive to consider the current landscape of spatially resolving UCN detectors and the trade-offs that characterize their performance. Spatially resolving UCN detectors can broadly be grouped into two categories:

1. detectors with high spatial resolution on the order of 1 µm, typically characterized by limited active area and comparatively low count rate capability and
2. large-area detectors with high count rate capability and spatial resolution on the order of 1 mm to 1 cm.

High spatial resolution in UCN detectors is primarily required in experiments that measure quantum states of neutrons in Earth's gravitational field (see [Nes+02]). The highest precision is typically achieved using passive detectors, which require post exposure analysis in a laboratory to extract the recorded data. Traditionally, nuclear track detectors have been employed in experiments investigating gravitational quantum states of UCNs [Nes+05] and continue to be used in more recent iterations of these experiments [Jen+13]. These detectors typically provide spatial resolution on the order of 1 µm [Jen+13]. More recently, even higher spatial resolution on the order of 100 nm has been achieved using passive nuclear emulsion detectors [Nag+18; Mut+22]. In an effort to develop online detectors suitable for this class of experiments, multiple detector prototypes based on boron coated, pixelated semiconductor devices such as charge-coupled devices (CCDs) or complementary metal-oxide-semiconductor (CMOS) devices have been developed. These devices typically achieve spatial resolutions better than 10 µm [Kaw+10; Kuk+21; Clé+22; Lin+23], but are generally limited in active area (typically less than $10\,\mathrm{cm}^2$) and in temporal resolution due to their frame-based readout, which typically operates at frame rates on the order of $30\,\mathrm{s}^{-1}$. These constraints also limit the achievable event rate, as event overlap within individual frames becomes significant at high rates. While semiconductor detectors with event based readout have been implemented in UCN detection and can, in principle, handle higher event rates, their size remains limited [Jak+09].

A more general-purpose UCN detector, the commercially available CASCADE detector[1], provides moderate spatial resolution by assigning UCN events to pixels of size $12.5 \times 12.5\ \mathrm{mm}^2$ [CDT18]. This resolution is, in principle, sufficient to perform coarse studies of UCN phase space and beam profile in UCN guides. The CASCADE detector is also capable of handling count rates of up to $40 \times 10^6\,\mathrm{s}^{-1}$ and covers areas of up to $200 \times 200\ \mathrm{mm}^2$ [CDT18]. However, as a gas-based detector, it requires a separation window, which limits its sensitivity to low-energy UCNs and complicates in-situ deployment. In the context of the τSPECT experiment, a detector of the second category is required due to the active area and rate requirements and because spatial resolution of approximately 1 mm to 2 mm is sufficient (see Section 3.8). Moreover, the experiment requires a detector capable of performing in-situ detection of low energy UCNs in a cryogenic, strong magnetic field environment at rates on the order of up to $1 \times 10^4\,\mathrm{s}^{-1}$. For these reasons, a solid conversion layer is required. The parameter space defined by a large detection area on the order of $100 \times 100\ \mathrm{mm}^2$, online readout capability, spatial resolution on the 1 mm scale, moderate global count rate capability and compatibility with in-situ deployment has, however, remained underexplored in

[1] CDT CASCADE Detector Technologies GmbH, CASCADE-U

the UCN detector landscape. In addition to its use case for a τSPECT-like experiment, these detector features are also relevant for more general applications, such as phase space diagnostics in UCN beam guides or studies of UCN transport properties. The following sections describe the setup and concept of a detector prototype designed to explore this parameter space.

4.3 Detector prototype setup overview

Fig. 4.1 illustrates the operating principle of the detector prototype design. From top to bottom, the detector consists of two thin conversion layers, a light guide and a grid of silicon photomultipliers (SiPMs). The conversion layers convert an incoming UCN into a scintillation light pulse. The light guide then distributes this light onto multiple SiPMs, which convert the the light into an electrical signal. By analyzing the relative SiPM signal amplitudes, the event position is reconstructed. As it employs solid conversion layers, SiPMs and a solid, non magnetic light guide, the design is fully compatible with vacuum environments and strong magnetic fields. The individual components from conversion layers to SiPMs are described in detail in the following sections.

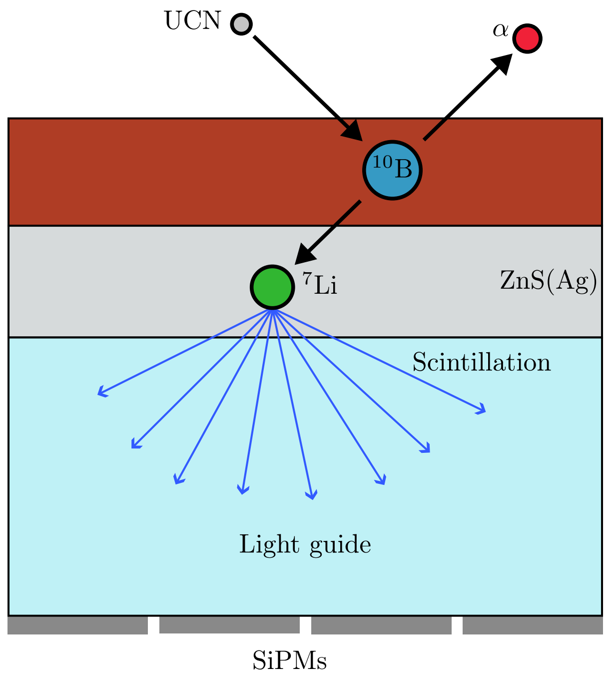

Figure 4.1: Schematic cross section of the detector setup illustrating its operating principle. Incoming UCNs are converted into a ^{7}Li nucleus and an α-particle via the ^{10}B neutron capture reaction. These ions generate a light pulse in the ZnS(Ag) scintillation layer, which is guided onto a grid of SiPMs by a light guide. Spatial resolution is achieved by evaluating the relative SiPM signal amplitudes.

4.4 Conversion layers

To convert incoming UCNs into light signals, an 80 nm ^{10}B layer stacked with an 8 µm ZnS(Ag) scintillation layer is employed. The scintillation layer was adhesively coated onto a 0.25 mm thick polyester sheet by the vendor and is commercially available[2]. The boron layer was deposited on top of the scintillation layer by a second company[3].

4.4.1 Boron layer

The boron layer has a calculated Fermi potential of −0.4 neV and a large neutron absorption cross section, marked by an imaginary Fermi potential component of 39 neV (see Appendix A). These properties make it well suited for the detection of UCNs down to the lowest energies. The relevant interaction of ^{10}B nuclei with UCNs is the (n, α) process

$$\begin{aligned} {}^{10}\mathrm{B} + \mathrm{n} &\longrightarrow \alpha + {}^{7}\mathrm{Li} + 2.79\,\mathrm{MeV}\ (6\,\%) \\ &\longrightarrow \alpha + {}^{7}\mathrm{Li} + \gamma(0.48\,\mathrm{MeV}) + 2.31\,\mathrm{MeV}\ (94\,\%), \end{aligned}$$

which converts neutrons into an α-particle and a ^{7}Li nucleus. The reaction exhibits two branches, one in which a ground state ^{7}Li nucleus is produced and one in which the ^{7}Li nucleus remains in an excited state that subsequently decays via γ-emission. The nuclear data for the ^{10}B neutron interaction were obtained from the ENDF database [Bro+18a]. It should be noted that cross sections are tabulated only down to neutron energies of 10 µeV. However, the cross section data show no resonance structure and follow the characteristic $1/v$ law for low-energy neutron cross sections [Kra88, pp. 456-459]. Since nuclear excitation energies are typically in the keV to MeV range, no additional resonances are expected in the UCN energy regime. The branching ratios and cross sections were therefore extrapolated to UCN energies.

The released energy is determined by the mass difference between initial and final state and is defined by

$$Q = (m_\mathrm{n} + m_\mathrm{B})\, c^2 - (m_\alpha + m_\mathrm{Li})\, c^2, \tag{4.1}$$

where m denotes the mass of the respective particle and c is the speed of light in vacuum. The particle masses used for these calculations are listed in Tab. 4.1. For the excited state branch, Q is reduced by the excited state energy. The kinetic energy of the reaction products follow from energy and momentum conservation. Because the released energy is small compared to the energy equivalent of the rest masses of the particles involved, a non-relativistic treatment is sufficient. The kinetic energies are therefore given by

$$E_\alpha = Q\, \frac{m_\mathrm{Li}}{m_\alpha + m_\mathrm{Li}} \tag{4.2}$$

and

$$E_\mathrm{Li} = Q\, \frac{m_\alpha}{m_\alpha + m_\mathrm{Li}}. \tag{4.3}$$

In the excited state branch, the momentum and kinetic energy of the ^{7}Li nucleus are modified when it decays to the ground state by emitting a γ-photon. To estimate the magnitude of these effects, two limiting cases are considered. The maximum deflection occurs when the γ-photon is emitted perpendicular to the momentum vector of the

[2]Eljen technology - EJ-440
[3]CDT CASCADE Detector Technologies GmbH, `https://n-cdt.com/`

TABLE 4.1: Mass values used for kinematic computations of the ^{10}B (n, α) reaction. Atomic mass values marked by (*) were reported as atomic masses in the cited reference and were converted to nuclear masses by subtracting the masses of the electrons of the corresponding neutral atoms. In all cases the quoted uncertainties are smaller than the numerical precision given here.

Particle	mass (u)	Reference
e	0.000 549	[Moh+25]
n	1.008 665	[Moh+25]
α	4.001 506	[Moh+25]
^{10}B	10.010 192*	[Wan+21]
^{7}Li	7.014 356*	[Wan+21]

nucleus. The photon momentum is defined as

$$p_{\mathrm{ph}} = \frac{E_{\mathrm{ph}}}{c} \tag{4.4}$$

and the corresponding deflection angle θ is given by

$$\theta = \arctan\left(\frac{p_{\mathrm{ph}}}{p_{\mathrm{Li}}}\right), \tag{4.5}$$

where p denotes the particle momentum. The maximum kinetic energy shift occurs when the photon momentum is antiparallel to the nucleus momentum. In this case,

$$E'_{\mathrm{Li}} = \frac{1}{2m_{\mathrm{Li}}}\left(p_{\mathrm{Li}}^2 + p_{\mathrm{ph}}^2 + 2p_{\mathrm{ph}}p_{\mathrm{Li}}\right) \tag{4.6}$$

describes the kinetic energy of the ^{7}Li nucleus after photon emission, leading to

$$\Delta E_{\mathrm{max}} = E'_{\mathrm{Li}} - E_{\mathrm{Li}} = \frac{1}{2m_{\mathrm{Li}}}\left(p_{\mathrm{ph}}^2 + 2p_{\mathrm{ph}}p_{\mathrm{Li}}\right). \tag{4.7}$$

This yields a maximum deflection angle of $\theta = 0.24°$ and a maximum kinetic energy shift of $\Delta E_{\mathrm{max}} = 7.7\,\mathrm{keV}$. Both effect are negligible. The reaction can therefore be treated, to good approximation, as a two-body decay with an angle of 180° between the products and with the particle energies listed in Tab. 4.2. Consequently, one of the ions is generally emitted towards the scintillation layer.

Based on the reaction product energies, their average ranges in ^{10}B were estimated using SRIM simulations [ZZB10]. Because the SRIM database only provides material properties for natural boron, the atomic mass was set to the value for ^{10}B recommended in [Wan+21] and the density was scaled by the ratio $m_{^{10}\mathrm{B}}/m_{\mathrm{B}}$. Here, $m_{^{10}\mathrm{B}}$ and m_{B} denote the atomic masses of ^{10}B and natural boron recommended in [Wan+21] and [BW11], respectively. The results are listed in Tab. 4.2. In addition to ion ranges, the absorption length of UCNs in the ^{10}B layer must be considered when selecting the layer thickness. The absorption length increases with the UCN velocity, as the absorption cross section is proportional to the inverse of the velocity. The optimal layer thickness therefore represents a trade-off between maximizing UCN absorption probability and ensuring transmission of the reaction products to the scintillator. This optimum depends on the expected UCN energy spectrum. Surface reflection and

Table 4.2: Kinetic energies and estimated mean ranges of the ions generated in the ^{10}B (n, α) reaction (branching ratios in brackets). Ranges in ^{10}B and ZnS(Ag) were obtained from SRIM simulations [ZZB10] as described in the main text.

Particle	Kinetic energy (MeV)	Range in ^{10}B (µm)	Range in ZnS(Ag) (µm)
α	1.78 (6 %)	4.4	5.1
	1.47 (94 %)	3.5	4.2
^{7}Li	1.01 (6 %)	2.1	2.5
	0.84 (94 %)	1.8	2.3

possible neutron transmission must also be considered. If transmission becomes significant, which is generally the case for higher energies, the properties of the substrate surface have to be considered as well. Such optimization procedures have e.g. been demonstrated in [Cle+19]. For the design of the τSPECT detector, similar studies were performed and multiple layer thicknesses were tested [Kah20]. Ultimately, a ^{10}B layer thickness of 80 nm was selected. Spare material from the τSPECT detector was still available and used for this work.

4.4.2 Scintillation layer

In order to convert the ions produced in the ^{10}B (n, α) reaction into a light pulse, a ZnS(Ag) scintillation layer is employed. The specified surface density is 3.25 mg/cm^2 [ELJ21a]. With a material density of 4.09 g/cm^3 [ELJ21b], this corresponds to an average thickness of approximately 8 µm. Since the typical particle size is also about 8 µm [ELJ21b], the surface is likely uneven, with local variations on the micrometer scale. The light yield of the scintillator is quoted as 300 % of that of Anthracene [ELJ21a], which is approximately 20 000 ph/MeV (photons per MeV) [YWF15]. This implies a light yield of roughly 60 000 ph/MeV for ZnS(Ag). Since the ion energies are independent of the incident UCN energy, the absolute light yield does not encode additional information about the neutron. For position reconstruction, only the relative light distribution across the SiPM grid is therefore relevant. A comparison of the scintillator emission spectrum with the spectral photon detection efficiency of the SiPM model used for this detector is shown in Fig. 4.15 in Section 4.6.1. Using SRIM simulations [ZZB10], the ranges of the ^{10}B neutron capture reaction products in the scintillator were estimated. For these simulations, the scintillator was approximated as a ZnS(50:50) compound layer with a density of 4.09 g/cm^3. No additional compound corrections were applied. The results are listed in Tab. 4.2 and indicate that the scintillation layer is sufficiently thick to fully stop all ions produced in the ^{10}B neutron capture reaction.

In addition to its response to heavy ions, the response of the scintillator to β- and γ-radiation, which in practice often contribute to the observed background, must also be considered. Due to the small thickness of the scintillation layer, the sensitivity to β- and γ-radiation is strongly reduced. Heavy charged particles deposit their energy on micrometer scales in solids, whereas electrons and photons typically exhibit substantially larger penetration depths at comparable energies. Thin ZnS(Ag) sheets therefore provide intrinsic suppression of β- and γ-induced background signals.

Next to stopping power and light yield, the scintillation time constant represents a key parameter. The scintillation process is commonly approximated by a simple exponential decay, which is defined as

$$L(t) = A \exp\left(\frac{-t}{\tau_{\mathrm{sc}}}\right), \tag{4.8}$$

where A is the initial amplitude, t the time after emission onset and τ_{sc} the scintillation time constant [Leo94, p. 158]. This time constant influences the signal shape, achievable timing precision and maximum event rate. For ZnS(Ag), the vendor specifies a time constant of $\tau_{\mathrm{sc}} = 200\,\mathrm{ns}$ [ELJ21a]. However, the literature reports a wide spread of time constants ranging from below 100 ns to tens of microseconds [Mor+17; Man+23; YT21]. In many cases, multiple exponential components are required to describe the scintillation process adequately. Furthermore, the signal shape depends not only on intrinsic scintillation dynamics but also on the response of the detector electronics. A dedicated characterization of the effective time constant for the final detector configuration is therefore necessary.

For the final position reconstruction, the spatial light intensity distribution of the scintillation light is of central importance. While the light emission process itself is isotropic, the opaque, polycrystalline structure of the material affects the light transport within the layer. As a consequence of this structure, scintillation light is expected to undergo scattering during propagation. In addition, partial self-absorption may occur. Consequently, light emitted at large angles relative to the surface normal is expected to experience stronger attenuation than light emitted at small angles. Accurate modeling of light transport would require detailed knowledge of optical properties, grain sizes and surface roughness, which are typically not available and must therefore be treated as unknowns in the detector design process. Implications of these effects for the position reconstruction of neutron events are discussed in detail in Chapter 5.

4.5 Light guide

To achieve spatial resolution, the light generated in the scintillation layer must be distributed onto a grid of SiPMs. This is realized by means of a so-called light guide. Light guides are typically fabricated from glass, acrylic glass or similar transparent materials and rely on internal reflections at their surfaces to transport light. Consequently, the geometric shape of the light guide determines how the scintillation light is distributed across the SiPM grid. The design of the light guide therefore constitutes a central element of the overall detector concept. To identify a suitable geometry, the relevant constraints must be considered first. These constraints arise from

- the cross section to be covered in the experiment,
- the expected distribution of events on the detector surface,
- the number of available SiPM readout channels,
- the size of the SiPMs
- and manufacturability considerations.

The most critical parameter is the detector cross section required in the experiment. For the present design, the cross section of the trap covered by the detector currently employed in τSPECT was used as a reference. This cross section is circular with a radius of $R = 49.5\,\mathrm{mm}$. Since events in τSPECT are expected across the entire

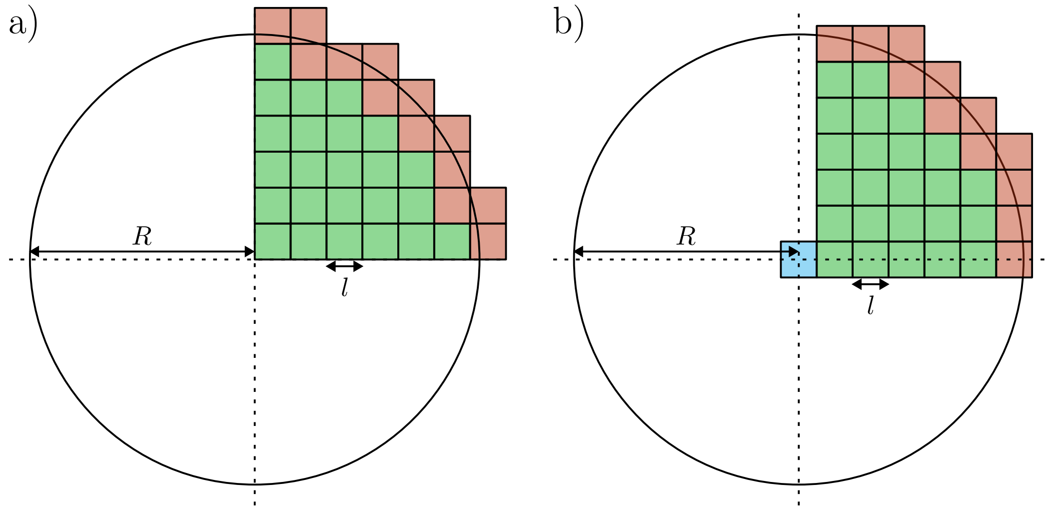

FIGURE 4.2: The two symmetric grid patterns for filling a circle of radius R with a grid of squares with edge length l. Due to the symmetry, only a quarter circle needs to be considered. Squares fully contained within the circle are shown in green, whereas squares intersected by the circle boundary are colored red. (a) Pattern (a): a square corner is placed in the circle center. (b) Pattern (b): a square center coincides with the circle center. In this case, the central square is counted only once and colored blue.

detector surface, the ideal configuration would cover the full circular area with a symmetric grid of SiPMs. However, the available installation space in the experiment is restricted to this radius by the octupole magnet that confines the UCNs radially. As a result, the SiPMs must be arranged such that they fit within the circular boundary. To determine the optimal packing, the number of SiPM sized squares that can be fully contained within a circle of radius R was calculated. The SiPMs feature an active area with edge length $s = 6\,\mathrm{mm}$ [Sem22], which corresponds to the largest commonly available size. Due to symmetry, it is sufficient to consider only a quarter circle when determining how many squares fit into a circle of radius R. The two symmetric packing patterns are shown in Fig. 4.2. In pattern (a), a corner of a square coincides with the center of the circle, whereas in pattern (b) the center of a square is located at the circle center. The pattern that accommodates the larger number of squares depends on the ratio

$$\eta = \frac{R}{l} \tag{4.9}$$

between the circle radius and the square edge length l. In the following, only the case $\eta \geq \sqrt{2}$ is considered, corresponding to at least one square contained in a quarter circle. For both configurations, the number of squares is obtained by determining how many squares fit into each row. For case (a),

$$N_\mathrm{a}(\eta) = 4 \sum_{i=1}^{\left\lfloor \sqrt{\eta^2-1} \right\rfloor} \left\lfloor \sqrt{\eta^2 - i^2} \right\rfloor \tag{4.10}$$

gives the total number of squares $N_\mathrm{a}(\eta)$. Here $\lfloor x \rfloor$ denotes the floor operation, which rounds x down to the nearest integer. The upper summation limit $\left\lfloor \sqrt{\eta^2 - 1} \right\rfloor$ corresponds to the number of rows, while the inner floor function determines how many square edge lengths fit between the circle center and the circle boundary at the height

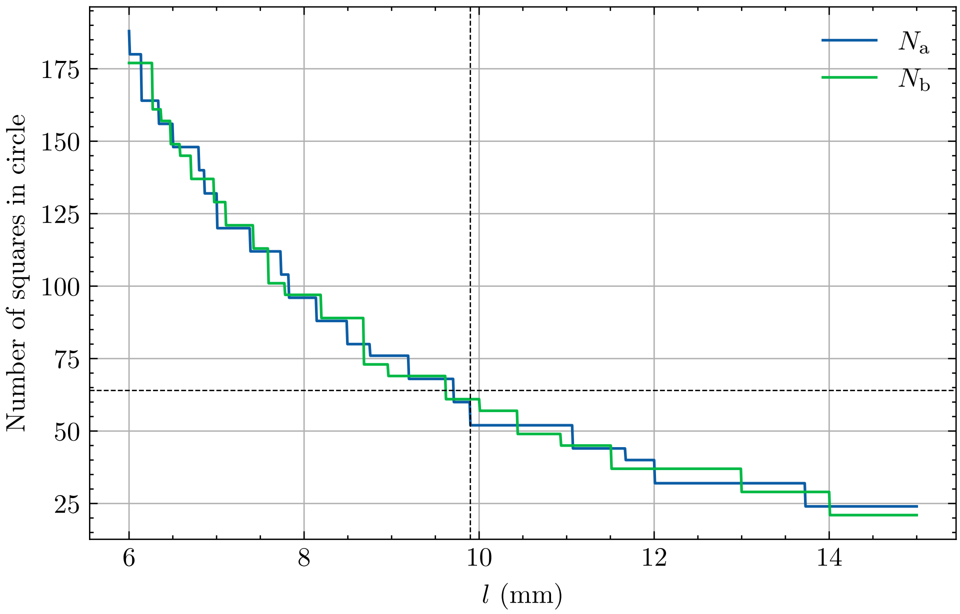

FIGURE 4.3: Maximum number of squares that can be fully contained within a circle of diameter 99 mm for the two grid patterns shown in Fig. 4.2, plotted as a function of the square edge length l. The horizontal dashed line at 64 indicates the maximum number of SiPM readout channels. The vertical dashed line indicates an edge length of 9.9 mm, which was ultimately selected.

of the upper edge of each row. The result is multiplied by 4 to account for symmetry. For case (b),

$$N_b(\eta) = 1 + 4 \sum_{i=0}^{\left\lfloor \sqrt{\eta^2 - \frac{1}{4}} - \frac{1}{2} \right\rfloor} \left\lfloor \sqrt{\eta^2 - (i + \frac{1}{2})^2} - \frac{1}{2} \right\rfloor \tag{4.11}$$

applies. The procedure is analogous to case (a), but all rows are shifted down and to the right by $l/2$ and the center square must be added explicitly. In Fig. 4.3 N_a and N_b are shown for a circle diameter of 99 mm and as a function of the square edge length l. The plot demonstrates that the optimal grid pattern depends on the chosen edge length. For $l = 6$ mm, a total of 188 SiPMs would be required to achieve dense coverage of the detector cross section. However, 64 channels typically represent the upper limit of a single SiPM readout system and were therefore adopted as a practical constraint for the first prototype. Consequently, the SiPMs must be arranged in a sparse grid. The plot can then be used to determine both the maximum feasible number of SiPMs and the corresponding grid spacing. The horizontal dashed line in Fig. 4.3 indicates the threshold of 64. From this, it follows that 61 SiPMs arranged according to pattern (b) with a grid spacing below 10 mm constitute the optimal solution under the given constraints. A spacing of 9.9 mm was ultimately selected in order to maximize coverage while retaining a small tolerance margin. Fig. 4.4 shows a schematic cross section of the resulting SiPM distribution within the circular cross section. For SiPMs with an active area of $6 \times 6\text{mm}^2$, this configuration results in a gap of $g = 3.9$ mm between adjacent SiPM edges that is not covered by a light sensitive surface. To mitigate light losses in these regions, the light guide geometry can be designed such that part of the light incident on these gaps is redirected onto neighboring SiPMs. Fig. 4.5 provides a schematic illustration of the corresponding vertical detector cross section and dimensions. As discussed above, shape parameters

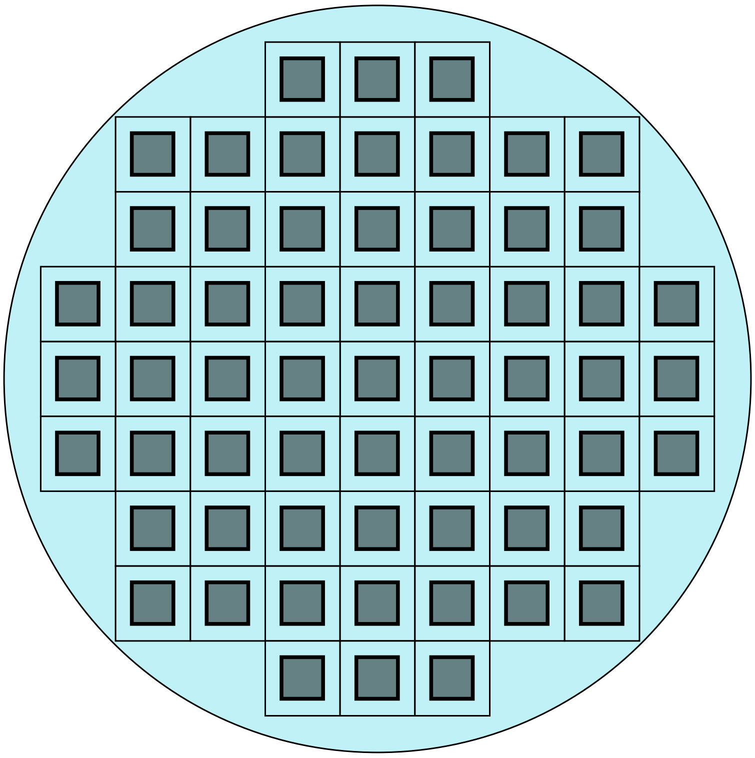

Figure 4.4: Optimized symmetric SiPM grid within a circular cross section of 99 mm diameter, corresponding to the size of the current τSPECT detector. The SiPMs are shown as grey squares. The black grid indicates the grid used to determine the optimal packing. The figure is drawn to scale.

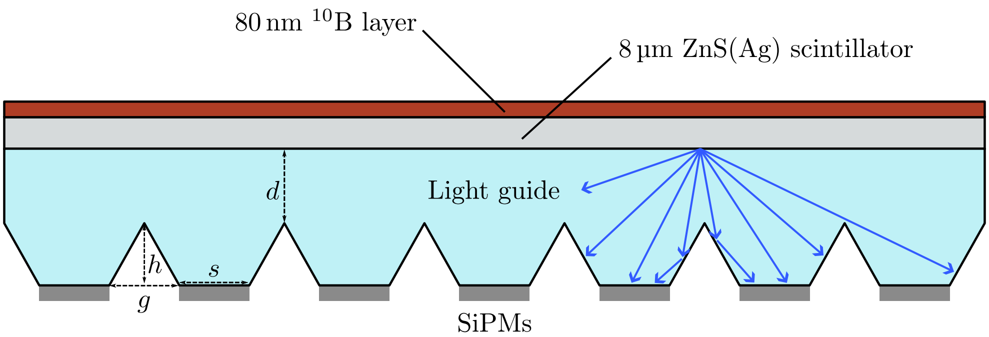

Figure 4.5: Schematic vertical cross section of the detector stack. UCNs are captured in the ^{10}B layer and converted into high energy ions via neutron capture. These ions produce scintillation light in the ZnS(Ag) layer, which is subsequently guided to the SiPM grid by the light guide. s, g, d and h denote the parameters defining the light guide geometry.

s and g are fixed by the SiPM dimensions and the required detector diameter. In order to optimize the remaining geometric parameters d and h, a dedicated Monte Carlo simulation tool was developed, which is discussed in detail in Section 4.5.1. The proposed circular arrangement exhibits uncovered regions near the perimeter (see Fig. 4.4), implying reduced detection efficiency in these areas. Additionally, spatial reconstruction near the edges may be degraded, as the regular grid ends and the boundary modifies the light propagation asymmetrically. In principle, the SiPM arrangement is also one of the factors determining the achievable spatial resolution of the detector. However, the statistical uncertainties associated with estimating the detected amount of light from the SiPM signals are difficult to predict a priori. As a consequence, an experimental prototype is required to obtain a reliable estimate of the achievable spatial resolution. Since the first detector iteration is intended primarily for characterization studies, including an evaluation of the spatial resolution, and is unlikely to be installed in the τSPECT experiment without further modifications, a square light guide geometry was selected for the initial prototype. The selected shape consists of an 8×8 grid of SiPMs with a spacing of 9.9 mm, thereby approximating the central region of the circular design in a geometrically simpler configuration.

4.5.1 Light guide simulation

The purpose of the simulation is to determine a light guide geometry that ensures a sufficient number of SiPMs receive a sufficient light signal for all possible event positions. To optimize the light guide geometry in this regard, the simulation estimates the light distribution on the SiPM array as a function of the event position and light guide shape. Photons emitted by the scintillator are treated as rays and their trajectory is computed using the Fresnel equations. Whenever a photon reaches an SiPM in the simulation, a counter associated with that SiPM is incremented. A detailed description of the implementation is provided in Appendix C. However, the simulation does not include the optical properties of the scintillator. As explained in Section 4.4.2, the scintillation layer consists of an opaque, polycrystalline powder adhesively bonded to a thin polyethylene sheet. Because of this complex structure, its scattering and absorption properties are not well characterized. As scintillation light is emitted isotropically, an isotropic point like source was adopted as an idealized approximation. In addition, the interface between scintillator and polyethylene sheet is not well defined due to the granular structure of the scintillator. Therefore, refraction at this boundary is neglected in the simulation as well. In reality, photons emitted at larger angles with respect to the surface normal are expected to experience stronger attenuation than predicted by the model. Consequently, the simulation likely overestimates the photon fraction detected by SiPMs located farther from the event position. To partially account for photons that are reflected back into the scintillator, an absorber is placed on top of the light guide in the simulation. This constitutes a conservative assumption that accounts for the limited knowledge of the scintillator's optical properties.

To identify a suitable light guide geometry, several constraints must be considered. Since an event can occur anywhere on the detector surface, the geometry cannot strongly favor specific event positions. Moreover, enhancing photon collection for certain SiPMs at specific positions inevitably reduces the photon flux received by others (see Fig. 4.5). Therefore, the set of SiPMs that is supposed to receive a significant photon signal must be defined as a function of the event position. To evaluate this, the geometrically extreme cases of an event position directly above the center of an

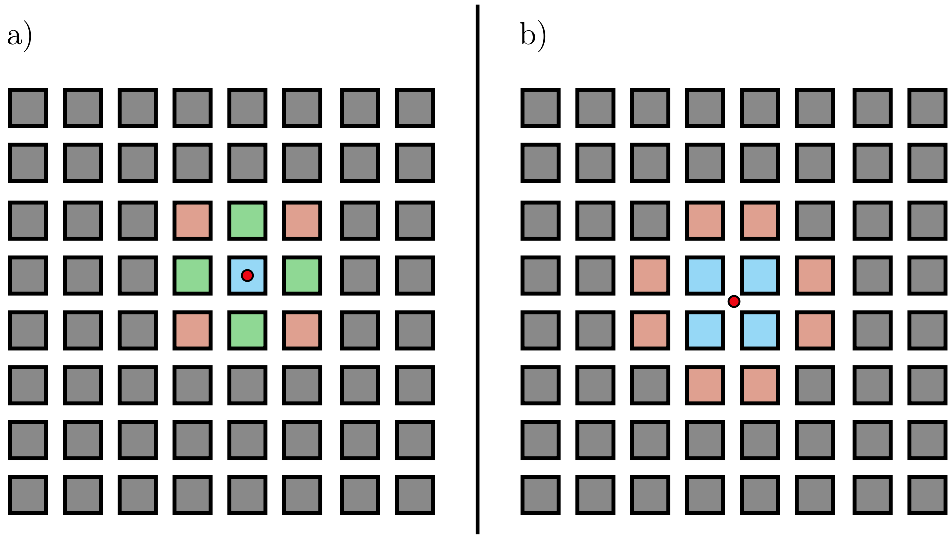

Figure 4.6: Geometrically extreme cases of (a) an event above the center of an SiPM and (b) an event above the barycenter of four diagonally adjacent SiPMs. The red dots indicate the simulated event positions, whereas the squares represent the SiPM grid. The colored SiPMs denote the distance based groups considered in the analysis.

SiPM (a) and above the barycenter of four diagonally adjacent SiPMs (b) were considered. In each case, the SiPMs can be grouped according to their geometric distance from the event position, forming discrete distance based groups. Fig. 4.6 illustrates these groups for the first three and first two distance groups in cases (a) and (b), respectively. The number of distance based SiPM groups included in the optimization must be defined a priori, as it determines the set of channels required to provide a usable signal. Restricting this to only the nearest groups would limit the number of contributing channels in the least favorable geometric configuration to a small set of SiPMs, reducing robustness against statistical fluctuations and deviations from the assumed light transport model. Furthermore, geometric truncation is expected to reduce the number of SiPMs receiving a significant photon flux for events occurring near the edge of the light guide. To ensure robust position reconstruction under these conditions and to provide a margin against photon count overestimation due to model idealizations, additional distance based SiPM groups were included. The selected SiPMs in the two geometric extreme cases are highlighted in Fig. 4.6. This choice results in nine contributing SiPMs in case (a) and twelve for case (b). For intermediate positions, the illumination pattern is expected to interpolate between these cases.

To determine the optimal SiPM array distance, the fraction of photons received by the different SiPM groups was simulated as a function of distance $D = d + h$. For this study, the light guide was removed and all refractive indices were set to 1. For each value of D, the trajectories of 200 000 photons were computed. The results for cases (a) and (b) are shown in Fig. 4.7 and Fig. 4.8, respectively. As expected, the group 2 SiPMs in case (b) exhibit the lowest photon fractions. The light guide design was therefore optimized to maximize the photon flux directed towards these channels. A total thickness of $D = 10.0\,\mathrm{mm}$ was selected, leaving h as the only free design parameter. Varying h in steps of 0.5 mm, an optimum was found at $h = 1.0\,\mathrm{mm}$, corresponding to a photon flux fraction of 0.54(2) %. The uncertainty was computed

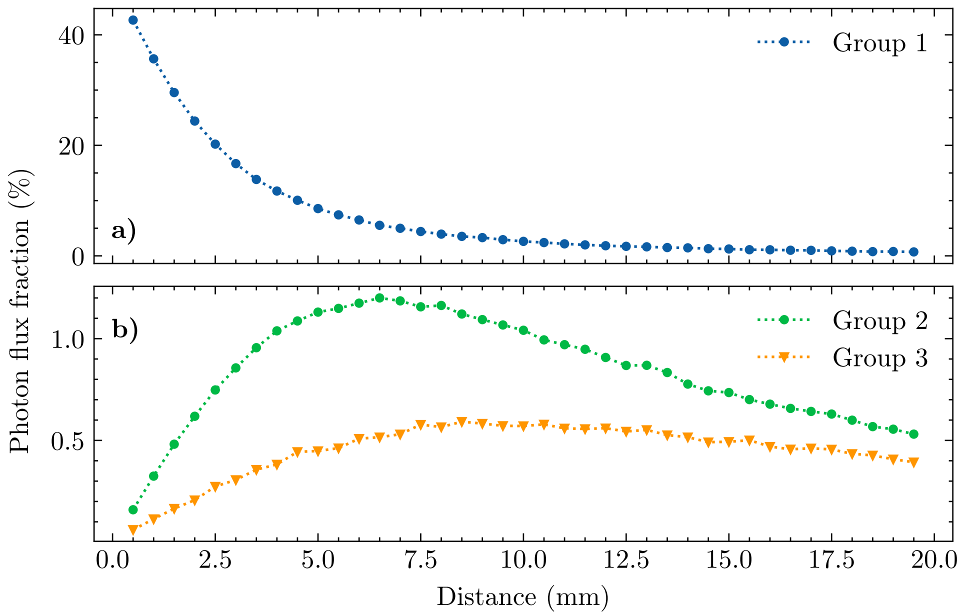

FIGURE 4.7: (a) Photon flux fraction received by the group 1 SiPM for an event above the center of one of the SiPMs (blue SiPM in Fig. 4.6 (a)). (b) Photon flux fraction received by group 2 and group 3 SiPMs corresponding to the green and red colored SiPMs in Fig. 4.6 (a).

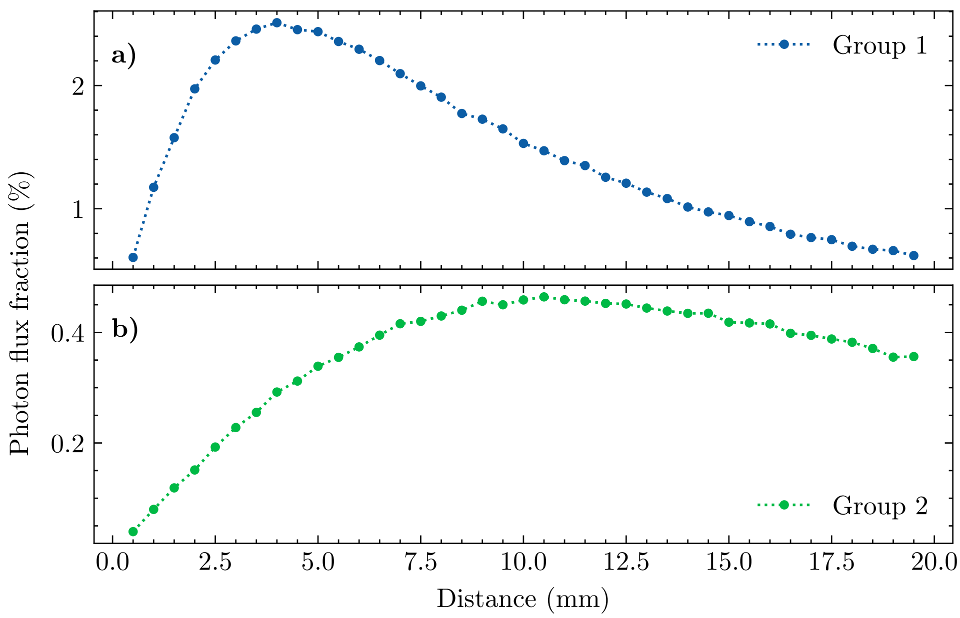

FIGURE 4.8: (a) Photon flux fraction received by group 1 SiPMs for an event between four SiPMs in the grid (blue SiPMs in Fig. 4.6 (b)). (b) Photon flux fraction received by group 2 SiPMs corresponding to the red colored SiPMs in Fig. 4.6 b).

assuming binomial statistics,

$$\sigma_{\mathrm{SA}} = \sqrt{N_{\mathrm{ph}} \frac{N_{\mathrm{c}}}{N_{\mathrm{ph}}} \left(1 - \frac{N_{\mathrm{c}}}{N_{\mathrm{ph}}}\right)}, \tag{4.12}$$

where N_{ph} is the number of simulated photons and N_{c} is the number of photons counted by one of the group 2 SiPMs in case (b), assuming 100 % detection efficiency. Compared to 0.46(2) % without a light guide, this corresponds to an improvement of approximately 17 %. For $h = 1.0$ mm, no direct line of sight between the event position and any of the considered SiPMs is obstructed by the guiding structures. Based on these results, a square 8×8 light guide with $s = 6.0$ mm, $g = 3.9$ mm, $d = 9.0$ mm and $h = 1.0$ mm was selected. To verify that a sufficient number of SiPMs receives a photon fraction comparable to the minimum observed in the extreme cases across the symmetry cell of the grid, a 100×100 grid of events was simulated within one quarter of the symmetry cell (see Fig. 4.9). Due to the square symmetry, simulating one quarter is sufficient. For each position, 200 000 photons were simulated. Since the minimum photon fraction observed by any SiPM in the extreme cases is approximately 0.5 %, this value was used as a threshold to define whether an SiPM receives a sufficient signal. Based on the light yield of ZnS(Ag) (see Section 4.4) and assuming no photon losses, this corresponds to approximately 250 photons for the minimum kinetic energy of the ions generated in the ^{10}B neutron capture reaction. This number is well above the expected detector threshold and provides a margin against losses and the SiPM photon detection efficiency. The results, shown in Fig. 4.10, indicate that at least 7 SiPMs exceed this threshold for all simulated positions. This was considered sufficient for constructing a first prototype. It should be noted, however, that this number is expected to decrease for events near the physical edge of the light guide. However, for the first prototype, the primary objective is to evaluate the performance in the central region of the detector. The light guide geometry at the edges will, in any case, require further adaptation for a design intended for use in the τSPECT experiment.

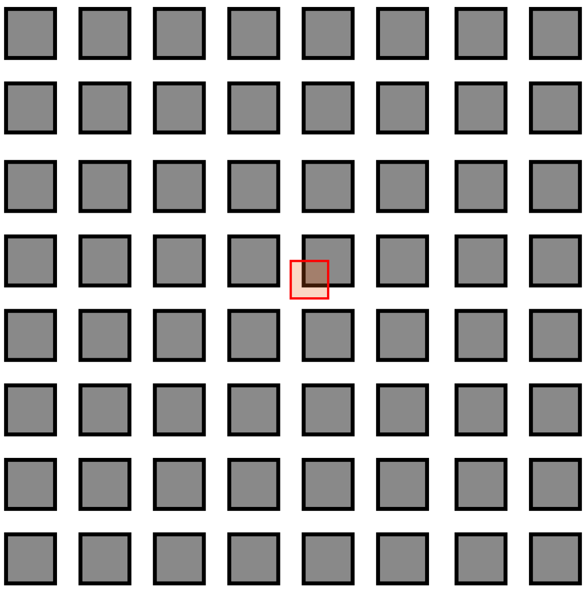

Figure 4.9: The grey squares represent the 8×8 SiPM grid used for the detector prototype. The red square marks a quarter of the symmetry cell of the grid, which is also the section of the grid shown in the simulation results in Fig. 4.10.

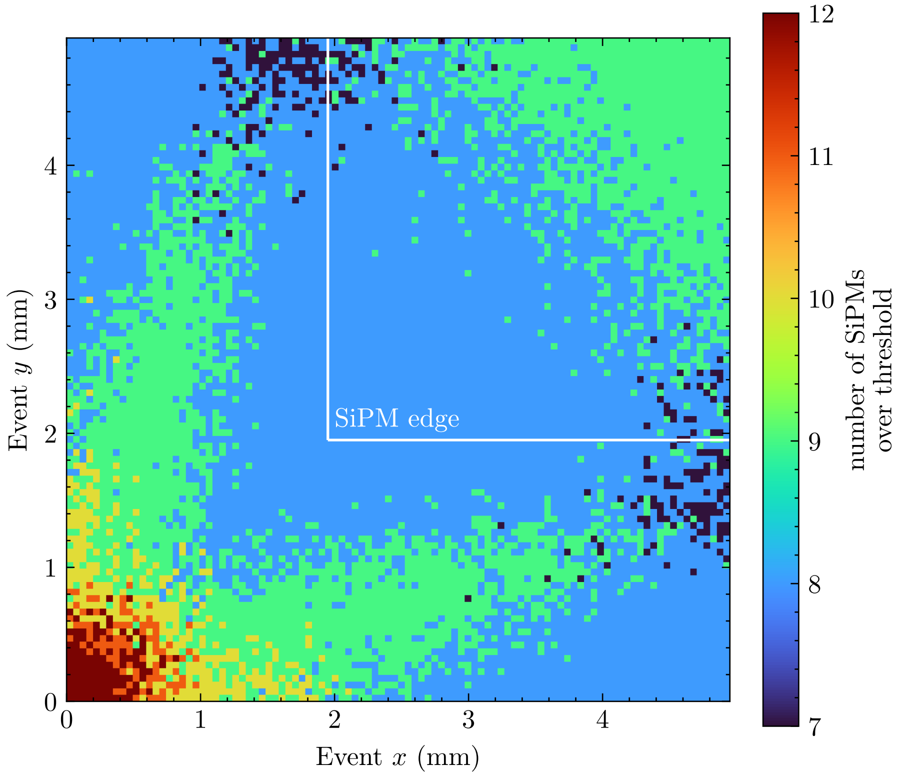

FIGURE 4.10: Each pixel corresponds to a simulated event position within a quarter symmetry cell (see Fig. 4.9). For each position, the color encodes the number of SiPMs receiving a photon flux fraction greater or equal to 0.5 %. A 100 × 100 grid of events was evaluated and the white lines indicate the boundary of the region covered by the SiPM.

4.5.2 Assessment of Light guide production methods

The selected light guide geometry is comparatively challenging to manufacture due to the narrow trenches between the SiPM contact surfaces. Conventional machining would therefore be time consuming and additionally constrained by the availability of suitable raw material. Recently, the detector laboratory of the PRISMA$^+$ cluster of excellence at JGU established a 3D-printing process for the production of light guides [Wei+23]. Although the printed light guides require subsequent mechanical post-processing to remove a non-transparent surface layer, this approach significantly reduces both production time and material cost. To assess whether a 3D-printed light guide satisfies the detector requirements, several small Plexiglass test samples of identical geometry were fabricated in the detector laboratory and the mechanical workshop. For the test measurements, a piece of ZnS(Ag) scintillator sheet and an SiPM (see Section 4.6) were attached to opposite ends of the light guides. The scintillator was irradiated with an ^{241}Am α-source and the resulting SiPM signals were recorded using an oscilloscope[4]. A cross sectional view of the measurement setup is shown in Fig. 4.11. The tested light guide shape consists of a cuboid section joined to a pyramidal frustum with dimensions as indicated in Fig. 4.11. The ^{241}Am α-source, with an activity of about 100 Bq, was custom produced for this kind of measurement by a researcher from another group at the research reactor TRIGA Mainz at JGU. It consists of a small droplet of ^{241}Am solution deposited onto a titanium foil. To ensure

[4]Siglent SDS5104X

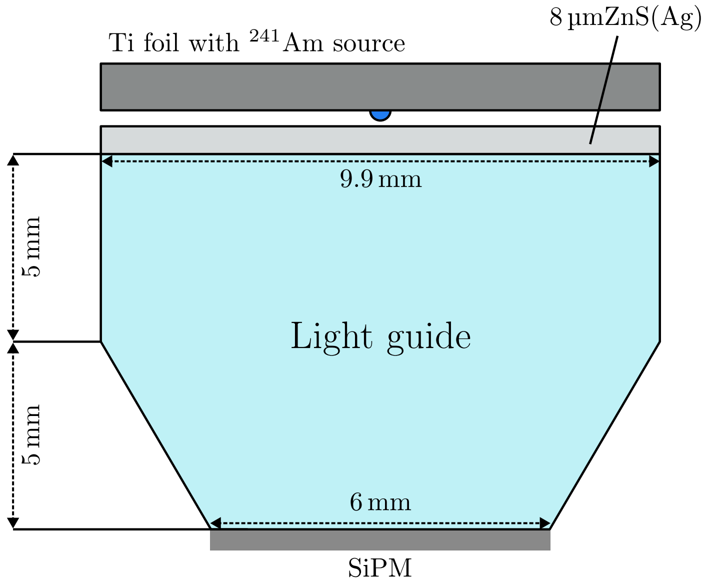

FIGURE 4.11: Cross sectional view of the setup used to compare light guides produced conventionally and by 3D-printing. α-particles emitted by the ^{241}Am source generate a light pulse in the scintillation layer which is recorded by the SiPM.

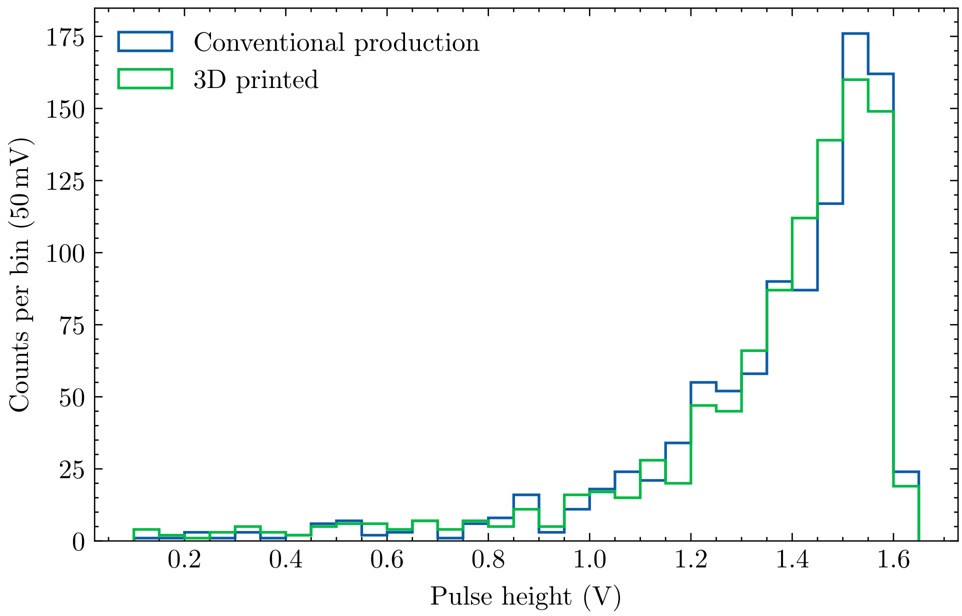

FIGURE 4.12: SiPM pulse height spectra measured for two geometrically identical light guides produced conventionally and by 3D-printing. The signals were generated by irradiating ZnS(Ag) with α-radiation from a ^{241}Am source.

consistent positioning of the source, the titanium foil was cut to match the light guide surface and was glued to a piece of ZnS(Ag) scintillator sheet of the same size. Care was taken not to cover the active spot with glue and the active spot was positioned close to the center. The resulting scintillation light source was then optically coupled to each light guide sample using optical grease[5], providing a non-permanent optical interface. The SiPM was coupled to the light guides in the same manner, allowing the samples to be exchanged without substantial modification of the setup. Since the SiPM signals require amplification for oscilloscope readout, an amplifier based on the design presented in [Cat+14] was employed. As part of this thesis, the design was implemented as a printed circuit board (PCB) and is currently also used in the monitor detector of the τSPECT experiment. For each light guide sample, a total of 1000 SiPM waveforms were recorded with the oscilloscope. The pulse amplitudes were extracted from the waveforms to generate pulse height spectra. By comparing the spectra obtained from 3D-printed and conventionally manufactured samples, the relative light guide transmission was evaluated. Representative pulse height spectra are depicted in Fig. 4.12. The results indicate that the samples produced conventionally exhibit slightly higher performance. However, the observed difference is marginal and comparable to the variation observed between samples from the same production method. Therefore, it was concluded that the advantages of the 3D-printing technique outweigh the minor reduction in performance and the 3D-printing production was selected for the production of the detector light guide.

4.6 Silicon Photomultipliers

A silicon photomultiplier (SiPM) is a semiconductor device designed for the detection of very low light levels. In contrast to more established low-light sensors such as photomultiplier tubes, SiPMs are insensitive to magnetic fields and are available in significantly smaller form factors. This makes them particularly well suited for use in the τSPECT experiment, which requires an in situ detector operating in magnetic fields of up to several Tesla. The working principle of an SiPM is e.g. described in [Sem23; GH20] and is summarized in this section. Depending on the specific device and its dimensions, SiPMs consist of hundreds to thousands of individual microcells per mm^2, arranged in a regular grid. Each microcell comprises a photodiode connected in series with a resistor and the microcells are connected in parallel (see Fig. 4.13). When a voltage is applied in reverse bias, the depletion zone of the pn-junction in the photodiode expands. Since no free charge carriers are present in this region, no current flows. However, if a photon is absorbed within the depletion zone, an electron-hole pair is generated. The applied electric field accelerates electrons and holes towards the cathode and anode, respectively, resulting in a current proportional to the light intensity. If the bias voltage is sufficiently high, charge carriers gain enough kinetic energy between collisions to generate secondary electron-hole pairs through impact ionization. Under these conditions, the absorption of a single photon can trigger an avalanche, leading to a breakdown of the pn-junction. A photodiode designed to operate in this regime is referred to as single photon avalanche diode (SPAD). The minimum voltage required to initiate breakdown is the breakdown voltage V_{bd}, a key characteristic parameter of a SPAD and thereby the SiPM as a whole. To prevent permanent damage to the diode, a so-called quench resistor is added in series within each microcell. This resistor limits the current and reduces the voltage across the diode below V_{bd} shortly after the avalanche is triggered, thereby quenching it. The

[5] Eljen technology EJ-550

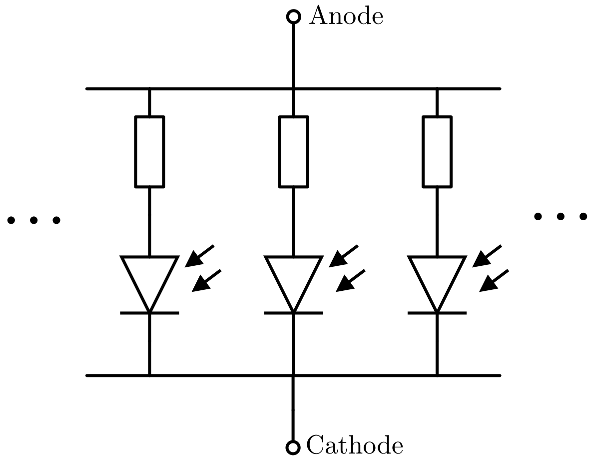

FIGURE 4.13: Effective circuit of an SiPM. Each photodiode, together with a series resistor, forms an individual microcell.

diode subsequently recharges with a time constant

$$\tau_{\mathrm{rc}} = R_{\mathrm{q}} C_{\mathrm{d}}, \tag{4.13}$$

where R_{q} is the quench resistance and C_{d} is the capacitance of a microcell. The sequence of avalanche breakdown and recharge produces the characteristic SiPM signal shape shown in Fig. 4.14. In real devices, the waveform tail typically exhibits an additional, faster time constant beyond τ_{rc}. This behavior is commonly attributed to parasitic capacitances associated with the quench resistor and the signal traces [Jha+13; KBB22]. Unlike a conventional photodiode operated well below V_{bd}, the output signal of a SPAD biased above V_{bd} is no longer proportional to the incident light intensity. Instead, each detected photon, down to the single photon level, produces a pulse with reproducible amplitude. Due to this digital, "click"-like behavior of the

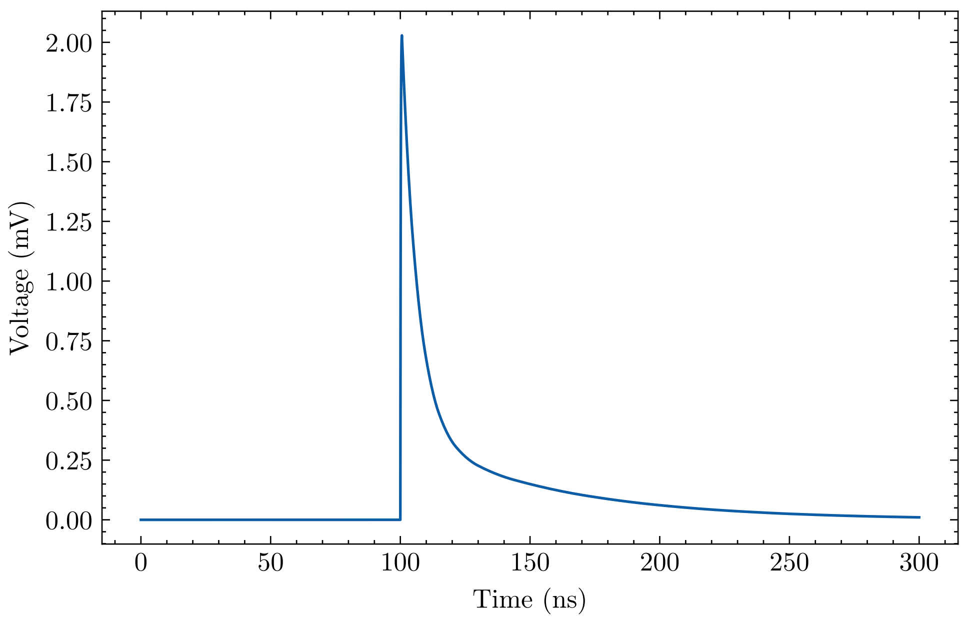

FIGURE 4.14: SPICE simulation of the signal generated by a single SiPM microcell. Circuit configuration and parameters as described in [Mar+13].

SPAD output signal, operation above V_{bd} is referred to as Geiger mode, by analogy with a Geiger-Müller tube. Because the response of a single SPAD is effectively digital, an SiPM integrates a large number of SPAD microcells connected in parallel. At light levels where single photon sensitivity is relevant, the probability that two photons hit the same microcell within one recharge cycle is small. Thus, the sum of the individual microcell currents is approximately proportional to the incident light intensity, resulting in a quasi-analog output signal.

4.6.1 Properties

SiPMs exhibit a wide range of performance characteristics, which are e.g. described in detail in [Sem23; GH20] that are summarized in this section. The following discussion focuses on the parameters most relevant for the presented detector design, which are photon detection efficiency over the emission spectrum of the scintillator, gain and breakdown voltage. The photon detection efficiency depends on three factors:

1. The quantum efficiency $\eta(\lambda)$, which is the probability that a photon of wavelength λ generates an electron-hole pair in a microcell.
2. The avalanche initiation probability $\epsilon(V)$, which depends on the bias voltage and describes the probability that a generated electron-hole pair triggers an avalanche. This bias voltage dependence can be understood as a consequence of the stronger electric field in the depletion zone at higher bias voltages, which enhances impact ionization.
3. The fill factor F, defined as the fraction of the SiPM surface area covered by light sensitive microcells.

So overall the photon detection efficiency is given by

$$\text{PDE} = \eta(\lambda)\,\epsilon(V)\,F. \tag{4.14}$$

This expression implicitly assumes uniform illumination of the SiPM surface due to the inclusion of the fill factor and is therefore not valid for strongly focused light. Additionally, both the quantum efficiency and avalanche probability exhibit a temperature dependence [Col+11]. The quantum efficiency decreases with decreasing temperature because the semiconductor band gap widens. In contrast, the avalanche probability increases at lower temperatures, as the mean free path of charge carriers increases. In general, the combined effect of these competing mechanisms is nontrivial, motivating extensive experimental studies of the photon detection efficiency of SiPMs at different temperatures [Col+11; Zha+22; Han+25]. For the presented detector design, it is essential that the photon detection efficiency spectrum of the selected SiPM model overlaps well with the emission spectrum of the scintillator. Fig. 4.15 shows both the scintillator emission spectrum and the photon detection efficiency spectrum of the SiPM model used in the detector prototype[7]. The substantial overlap indicates good spectral matching between scintillator and SiPM.

The detector concept relies on comparing the relative signal amplitude of multiple SiPMs to achieve spatial resolution. In this context, the gain is the most relevant SiPM parameter. The gain G is defined as the charge generated by a single photon

[6] `https://plotdigitizer.com/app`
[7] onsemi C-series MICROFC-60035-SMT

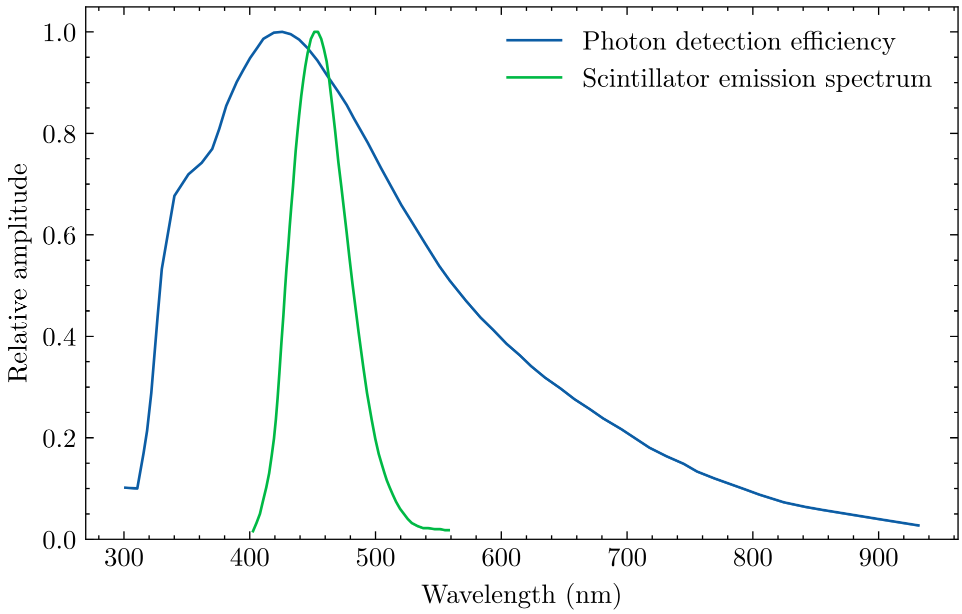

FIGURE 4.15: Emission spectrum of ZnS(Ag) compared with the photon detection efficiency curve of an onsemi C-Series SiPM. Both curves were manually extracted from the respective datasheets [ELJ21a; Sem22] using a plot digitizing tool[6].

and is proportional to the integral of a single photon signal. It is given by

$$G = \frac{(V_\mathrm{b} - V_\mathrm{bd})\, C_\mathrm{d}}{e}, \tag{4.15}$$

where V_b is the bias voltage, V_bd is the breakdown voltage, C_d is the microcell capacitance and e is the elementary charge. Since the gain is proportional to the overvoltage $V_\mathrm{OV} = V_\mathrm{b} - V_\mathrm{bd}$, it is directly linked to the breakdown voltage. According to the datasheet [Sem22], the breakdown voltage varies between individual SiPM samples of the same model. The datasheet of the selected SiPM model specifies a breakdown voltage range from 24.2 V to 24.7 V at a temperature of 21 °C and a gain of 3×10^6 at an overvoltage of 2.5 V [Sem22]. In the extreme case, two neighboring SiPMs operated at the same bias voltage could therefore exhibit a relative gain difference of 18 % if this variation were not accounted for. Since additional parameters, such as the photon detection efficiency, also depend on the overvoltage and thus on the breakdown voltage, the resulting mismatch becomes even more significant for signals containing larger numbers of photons. Moreover, the breakdown voltage exhibits a pronounced temperature dependence. With increasing temperature, enhanced charge carrier scattering reduces the mean free path, thereby requiring a higher voltage to reach the ionization threshold necessary for triggering an avalanche. Consequently, the gain of an SiPM changes with temperature, if the bias voltage remains constant. Therefore, the temperature of an SiPM based detector must be carefully monitored and, if necessary, stabilized or the change in breakdown voltage compensated to ensure constant gain. This is of particular importance for a detector intended for operation in a cryogenic environment as is e.g. the case in the τSPECT experiment.

4.6.2 SiPM breakdown voltage measurements

As established above, achieving uniform gain across all SiPMs used in the detector is important for good spatial resolution. Because the gain depends on the breakdown voltage, the breakdown voltage of each individual SiPM must be measured accordingly. Although the SiPMs are intended to operate at nearly identical overvoltages, the absolute value of the overvoltage is less critical, since position reconstruction depends primarily on the relative signal amplitudes of the individual sensors. Furthermore, the readout system used to bias and read out the SiPMs (see Section 4.7) applies bias voltages with a precision of 10 mV. At this level, gain variations arising from bias voltage fluctuations are expected to remain below 1 %. Consequently, a measurement setup must achieve a precision of approximately 10 mV and calibration to an absolute voltage reference is not required. Because the breakdown voltage is temperature dependent, the temperature of the measurement setup must be stabilized to ensure reproducible results.

In general, two widely used methods exist to determine the breakdown voltage of an SiPM [GH20]. The first method is based on gain measurements obtained by integrating SiPM signals produced by weak light pulses at different bias voltages. The breakdown voltage is then inferred via Eq. (4.15). The second method relies on recording the reverse bias current-voltage curve (IV-curve) of the SiPM by measuring the current as a function of applied voltage. At the breakdown voltage, a sharp increase of the current is observed. Since the first method requires both a weak pulsed light source and a precise bias voltage supply, recording an IV-curve is generally more straightforward. For this reason, the IV-curve method was selected for this study. Approaches for extracting the breakdown voltage from an IV-curve are discussed in the following sections.

Measurement setup

At voltages below the breakdown voltage, only small leakage currents on the order of a few nanoamperes are observed. Therefore, a device capable of measuring such low currents is required. For this study, a sourcemeter[8] was used. To record the IV-curve, the SiPM anode and cathode were connected to the negative (ground) and positive terminals of the sourcemeter, respectively (see Fig. 4.13). In addition to precise voltage and current measurement capabilities, the setup also had to meet the following requirements:

1. Temperature stabilization to prevent temperature induced drifts of the breakdown voltage as discussed above.
2. Operation under dark conditions to avoid ambient light exposure of the SiPMs.
3. A reliable electrical contact method that does not require soldering or any permanent connection, since the SiPMs are intended for subsequent use in other setups.

To measure the temperature of the setup and to establish electrical contact to the SiPM, a dedicated printed circuit board (PCB) was designed. Two temperature sensors[9] were soldered onto the PCB and read out using a Raspberry Pi[10] in combination with an analog to digital converter (ADC) module[11]. The implementation of the temperature readout was carried out as part of a Bachelor thesis [Hol23] and additional details on the temperature sensor readout are provided in Appendix D. To stabilize

[8] Keithley 2450
[9] Texas Instruments TMP61
[10] Raspberry Pi 4 model B
[11] DFRobot I2C ADS1115 16-bit ADC module

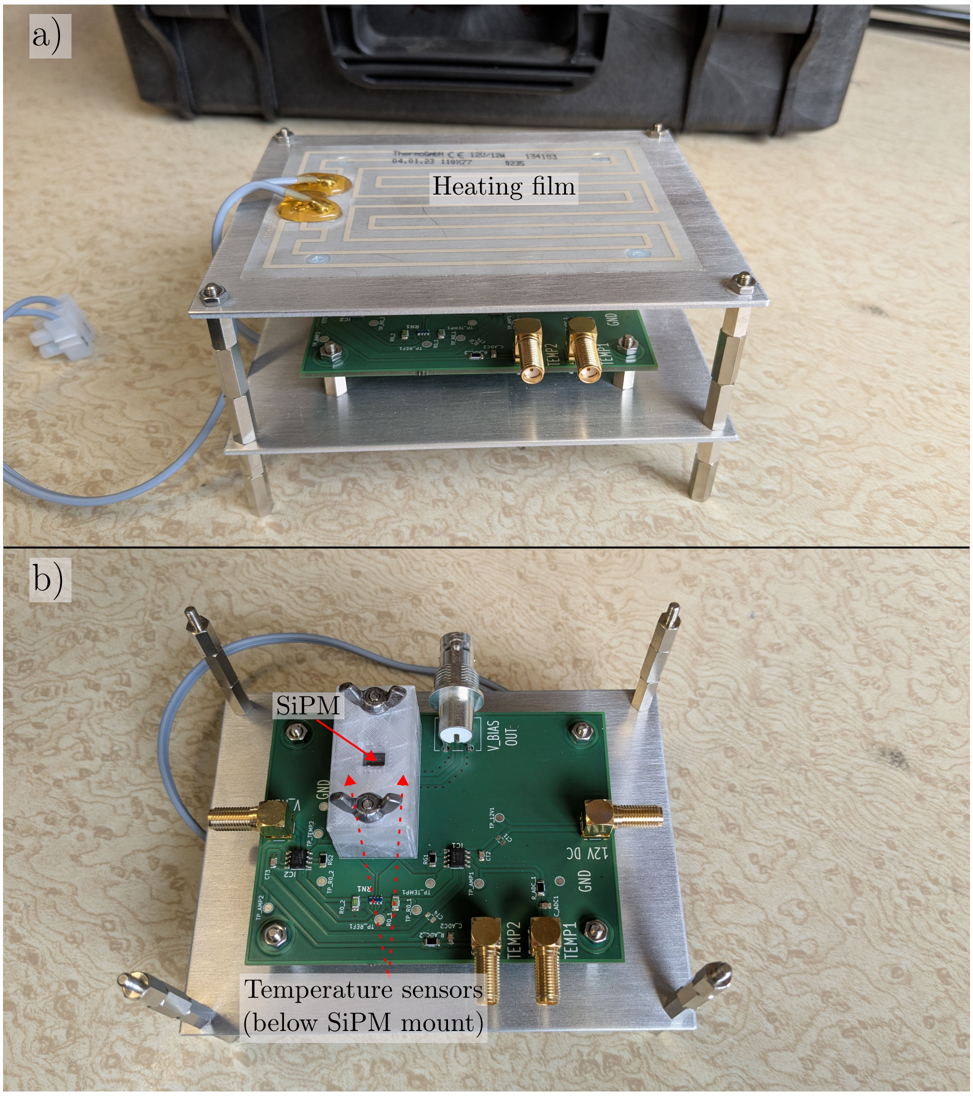

FIGURE 4.16: a) Photograph of the fully assembled SiPM test setup. The test setup PCB is mounted between two 3 mm thick aluminum plates used for temperature stabilization. A heating film is attached to both plates. Only the upper film is visible in the image. b) Photograph of the setup with the top aluminum plate removed. The 3D-printed SiPM mounting structure and an installed SiPM are visible. Two temperature sensors are located on the PCB below the SiPM mount, indicated by dashed red arrows.

the temperature, the setup was heated to a temperature slightly above room temperature, which could be realized using a 12 V, 12 W electrical heating film[12]. Two such films were attached to 3 mm thick aluminum plates, between which the PCB containing the SiPM mounting structure and temperature sensors was mounted using spacers. Photographs of the setup are shown in Fig. 4.16. To regulate the heating power, the average supplied current was controlled using solid-state relays[13] driven by the pulse width modulation (PWM) pins of the Raspberry Pi. Using the mean of the temperature sensor readings as the control input and the PWM duty cycle as the output variable, a proportional-integral (PI) controller was implemented on the

[12]Thermo Flächenheizungssysteme 2115404
[13]JOY-IT COM-MOSFET IRF9540N development board

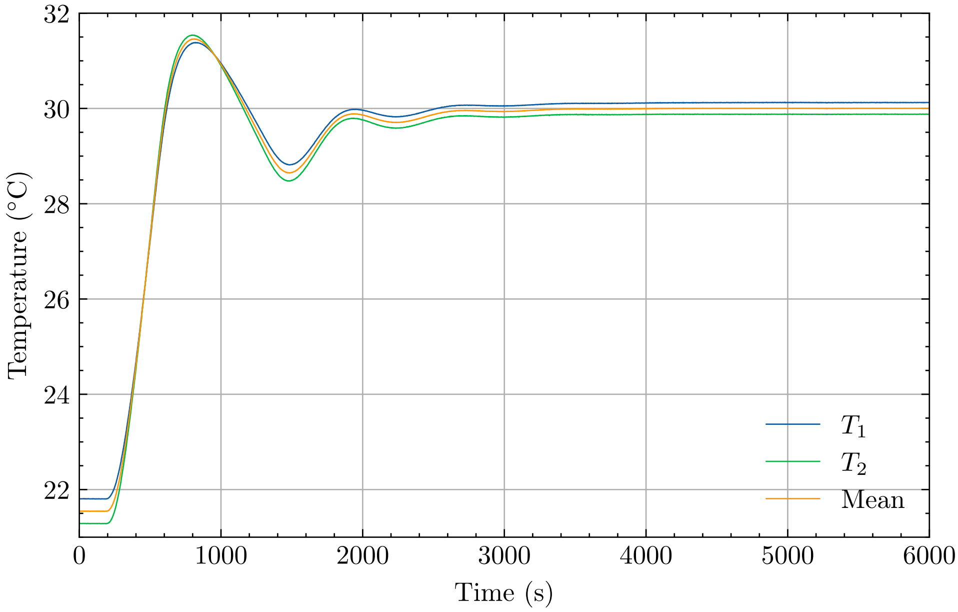

FIGURE 4.17: Temperature measured by sensors T_1 and T_2 of the SiPM test setup shortly after activation of the PI controller, which stabilizes the temperature to the mean of both sensor readings. The controller parameters were set to $K_\mathrm{p} = 35$ and $K_\mathrm{i} = 0.12$ (see Eq. (4.16)).

Raspberry Pi. The PI controller follows the control function

$$C(t) = K_\mathrm{p}\, e(t) + K_\mathrm{i} \int_0^t e(\tilde{t})\, \mathrm{d}\tilde{t}, \tag{4.16}$$

where t denotes the time and K_p and K_i are the proportional and integral constants, respectively. The error function is given by

$$e(t) = T(t) - T_\mathrm{target}, \tag{4.17}$$

where $T(t)$ is the current temperature and T_target is the target temperature. The PI controller was implemented as a Python script running on the Raspberry Pi and evaluated the control function every 2 s. If $C(t) \notin [0, 100]$, corresponding to the duty cycle range, the output was clipped to the nearest boundary value (0 or 100). Empirically, the parameters $K_\mathrm{p} = 35$ and $K_\mathrm{i} = 0.12$ were found to provide stable behavior. The initial temperature stabilization process is shown in Fig. 4.17. Typically, approximately 50 min were required to reach steady-state conditions. Fig. 4.18 shows the stability of the temperature sensor readings after the steady-state was achieved. Residual fluctuations were on the order of 10 mK. A slow drift of the two temperature sensors with respect to each other is also visible, which is likely caused by a small temperature gradient across the PCB that may vary with ambient conditions. However, this drift corresponds to a breakdown voltage variation of less than 1 mV and was therefore not investigated further. To shield the setup from ambient light and to provide sufficient isolation from the ambient laboratory temperature, the setup was placed in a light-tight black suitcase. To ensure a solder free and yet reliable electrical connection of the mounted SiPMs, spring probes[14] were employed. The probe tip consists of an array of small spikes that ensures a proper electrical contact when the

[14] Harwin P13-4023

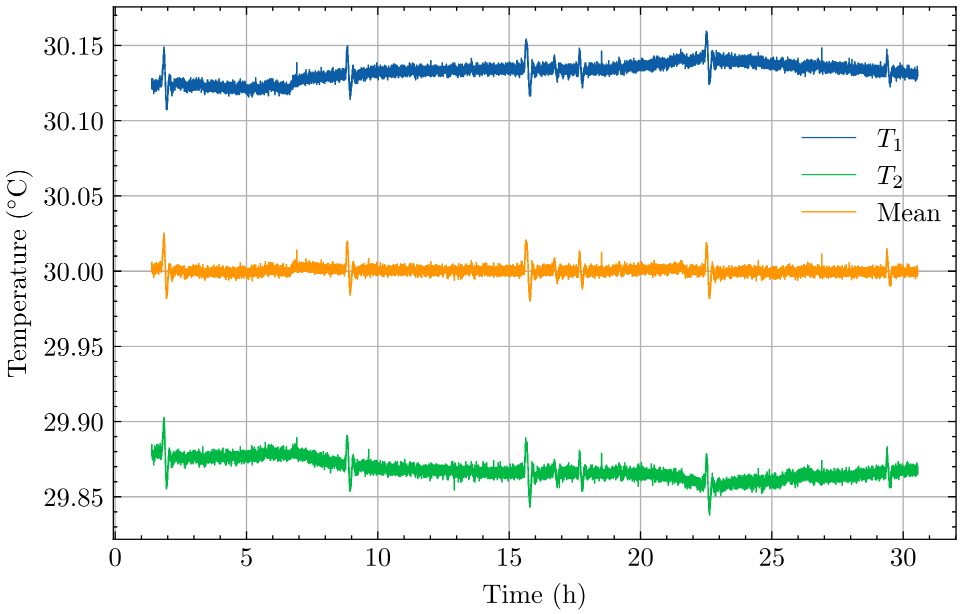

FIGURE 4.18: Temperature measured by temperature sensors T_1 and T_2 of the SiPM test setup while the temperature was actively stabilized by the PI controller. The mean of both sensor readings is ued as the controller reference.

SiPM contact pads are pressed onto them. To enable precise and convenient mounting of the SiPMs onto the spring probes, a three part 3D-printed mounting structure was designed and dimensionally optimized as part of a student project. The resulting setup allows quick exchange of the SiPM samples while maintaining a solder free connection. Within the SiPM test setup, the Raspberry Pi serves as the central control unit. In addition to temperature stabilization, it communicates with the sourcemeter via the local network. This enables automated triggering of IV-sweeps and direct retrieval of measurement data, resulting in a largely automated test procedure.

Breakdown voltage extraction

A large number of (semi-)heuristic methods have been used to extract the SiPM breakdown voltage from IV-curves. The differences between these approaches are e.g. discussed in [Kla19; Ott+17; Nag+17] and typically yield systematically different results. However, since no absolute voltage measurement is required in the present study (see above), these systematic differences are not critical. In the following, the method recommended by the manufacturer of the SiPM model used in this detector and one of the methods from the literature are discussed. The manufacturer recommends determining the breakdown voltage from the x-axis intercept of a linear fit to a plot of the square root of the current versus the bias voltage [Sem23]. An exemplary IV-curve for SiPM sample 0 is shown in Fig. 4.19. SiPM sample 0 served as a reference device for setup validation and was periodically remeasured to verify consistent operating conditions. As evident from the figure, the IV-curve does not exhibit perfect linear behavior above the breakdown voltage. Consequently, the slope of the linear fit and thereby the extracted breakdown voltage depend on the chosen fit range. Since the position of the breakdown voltage relative to the fit window differs between SiPMs, a fixed fit window would introduce a systematic bias. Furthermore, this method only yields results comparable to those provided by the manufacturer if the baseline conditions match those during the manufacturer's characterization. For these reasons, this

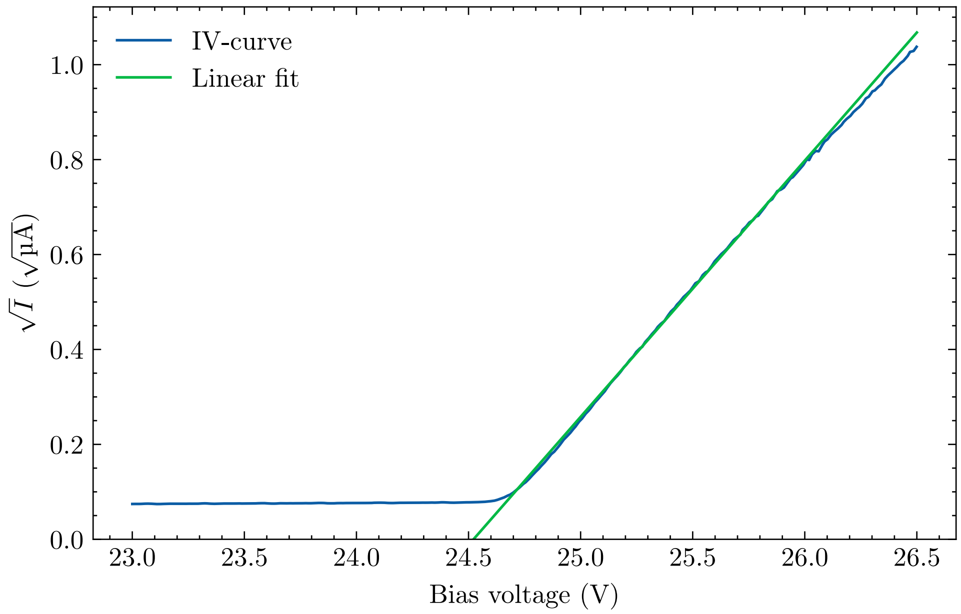

FIGURE 4.19: Square root of the current as a function of the bias voltage for an IV-curve of SiPM sample 0. Following the manufacturer's recommendation, a linear function was fitted to the data between 25 V and 26 V.

method was not adopted for the extraction of the breakdown voltage in this study. Alternative methods to extract the breakdown voltage commonly rely on a logarithmic representation of the IV-curve. An example for SiPM sample 0 is shown in Fig. 4.20. On the logarithmic scale, the sharp increase in current beyond the breakdown voltage is clearly visible. At the breakdown voltage, the slope of the logarithmic IV-curve changes abruptly. The breakdown voltage can therefore be determined by locating the maximum of the second derivative of $\ln(I(V_b))$, where I denotes the current and V_b the bias voltage. A detailed description of the numerical procedure used to determine this maximum is provided in Appendix D.2, including the required low-pass filtering of the first derivative. To assess the short term reproducibility, 1000 IV-curves were measured for SiPM sample 0. The distribution of the extracted breakdown voltages is shown in Fig. 4.21, yielding a standard deviation of 1.4 mV. In addition to short term statistical fluctuations, a gradual drift in the measured breakdown voltage of SiPM sample 0 was observed over approximately one year of operation.

To quantify this long term effect, the measurements of SiPM sample 0 were grouped into 11 subsets corresponding to distinct time intervals during which the device was installed in the setup. The subset sizes ranged from 5 to 1000 and the acquisition duration varied between 1 h and 12 d. The corresponding subset means are shown in Fig. 4.22. The drift magnitude was estimated by computing the standard deviation of the subset means, resulting in 3.3 mV. Two potential effects that may contribute to this drift are considered. First, the temperature sensors used for temperature stabilization are not directly attached to the SiPM. Variations in ambient conditions may therefore cause changes in the temperature offset between the SiPM and the sensors. Since the breakdown voltage changes by 21.5 mV/°C [Sem22], a variation of only 0.15 °C would account for a 3.3 mV shift. Verification of this hypothesis would require an additional temperature sensor mounted in close thermal contact with the SiPM, which was mechanically impractical and therefore not further pursued. Second, a long term drift of the sourcemeter voltage reading due to aging or environmental

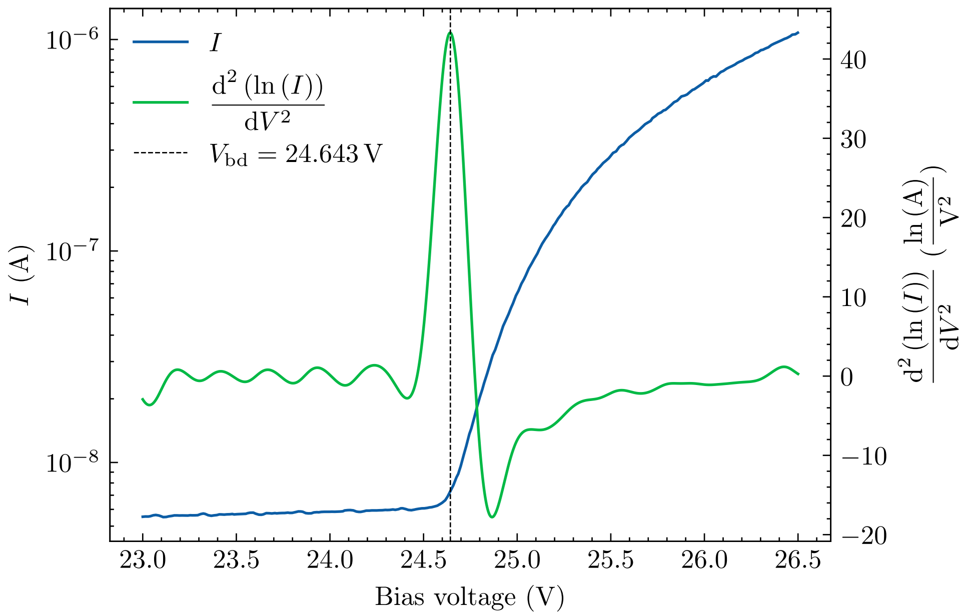

Figure 4.20: Logarithmic representation of the IV-curve of SiPM sample 0 recorded at 30 °C. The green curve shows the second derivative of the logarithmic IV-curve.

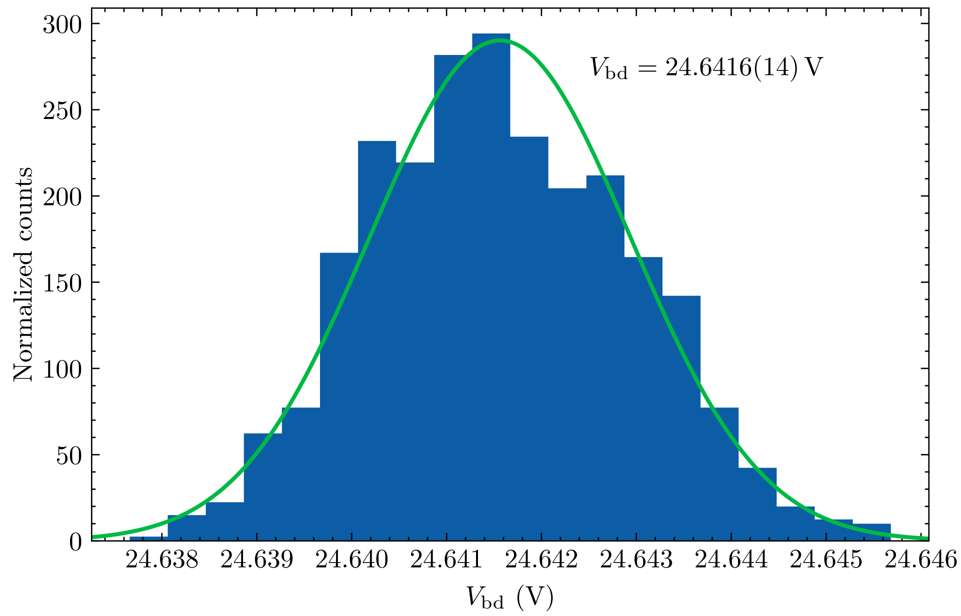

Figure 4.21: Histogram of 1000 consecutive breakdown voltage measurements of SiPM sample 0. A Gaussian distribution with the sample mean and standard deviation is overlaid to provide an estimate for how well the histogram is described in this way.

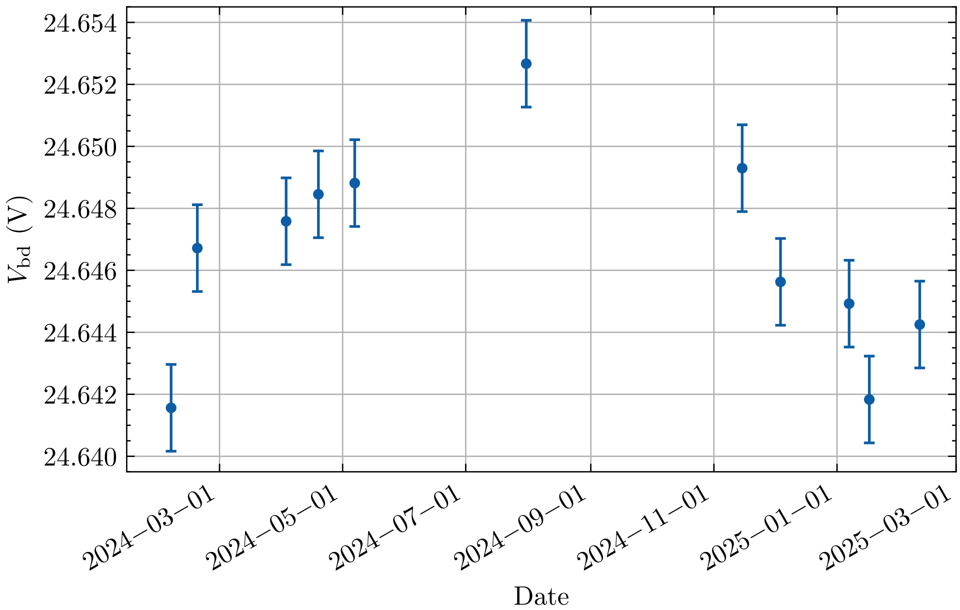

Figure 4.22: V_{bd} measurements of SiPM sample 0 grouped into 11 subsets corresponding to distinct measurement periods. Each data point represents the mean of one subset and the uncertainty is given by the short term fluctuations.

effects cannot be excluded. The specified accuracy in the relevant voltage range is 0.015 % + 10 mV one year after calibration [Tek22], referring to deviations from an absolute voltage reference. As demonstrated by the short term stability measurements, the relative precision on shorter timescales is substantially better. However, the datasheet does not specify long term drift behavior within the stated accuracy limits. Investigating this effect would require a highly stable voltage reference in the relevant voltage range, which was not available. Combining both short and long term contributions yields a total uncertainty of $\sigma_{bd} = 3.6$ mV, which remains well below the required precision of 10 mV. Consequently, the drift was not analyzed further.

Results and discussion

In total, 89 SiPM samples (including the reference sample) were characterized using the test setup. All measurements were performed at a stabilized temperature of 30 °C. For each IV-curve, 250 data points were recorded between 23.0 V and 26.5 V, with a delay of 0.3 s per point to allow the sourcemeter reading to settle. At least five IV-curves were recorded per sample and the corresponding breakdown voltages were averaged for Fig. 4.23. For SiPM sample 41, the IV-curves exhibited a significantly increased baseline noise level and a correspondingly a significantly larger spread in extracted breakdown voltages compared to all other samples. This sample was therefore classified as defective and excluded from further consideration. The measured breakdown voltages cannot be directly compared to the range quoted in the datasheet, as the present measurements were conducted at 30 °C rather than room temperature and a different breakdown voltage extraction procedure was used (see above). Nevertheless, the overall spread can be evaluated. The measurements span approximately 270 mV between minimum and maximum, compared to 500 mV specified in the datasheet [Sem22]. Thus, the observed variation is well within the manufacturer's specification. During the measurement campaign, SiPMs from three different production batches (bags), ordered over a period of approximately 2 years, were tested.

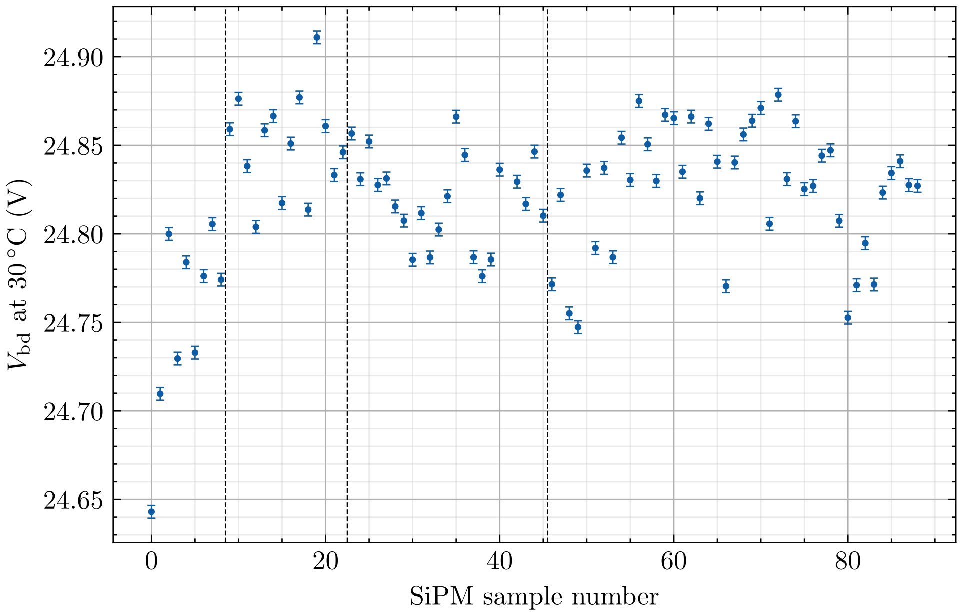

Figure 4.23: Breakdown voltage measurements of all tested SiPM samples at 30 °C. Sample 41 was excluded due to an unusually large spread in extracted breakdown voltages (see main text for details). The black dashed lines indicate when a new bag of SiPMs was opened.

The dashed black lines in Fig. 4.23 indicate the transition between batches. SiPMs from the first batch appear to exhibit slightly lower breakdown voltages than those from later batches. However, the statistical significance of this observation is limited by the comparatively small sample size. Apart from this minor trend, no systematic differences between batches were identified. For the detector prototype, the 64 SiPMs with the most similar breakdown voltages were selected from the available pool[15]. To achieve this, all samples were ordered by breakdown voltage and the contiguous subset with the smallest range between minimum and maximum values was chosen. This reduced the total spread to 76 mV, corresponding to a reduction by a factor of approximately 3.6 compared to a worst-case selection. The resulting maximum gain variation due to breakdown voltage differences is therefore expected to be approximately 3 %. This effect can be further mitigated by arranging SiPMs with similar breakdown voltages in close proximity within the detector. In this configuration, the impact depends on which specific SiPMs detect scintillation light in a given event.

In summary, systematic measurement and selection of a larger SiPM sample set enabled a minimization of breakdown voltage differences within the detector SiPM array and facilitated the identification and exclusion of a potentially defective sensor. Beyond optimizing the present detector prototype, the described test setup is useful infrastructure for future SiPM based detector developments within the collaboration.

4.7 Readout system

To read out signals from a large number of SiPM channels, a dedicated readout system is required. The primary challenges associated with many parallel channels are minimizing the physical footprint of the electronics and reducing the required data storage

[15]Samples 2, 4, 6, 8, 37 and 39 had been used for other projects.

disk space. Application-specific integrated circuits (ASICs), designed explicitly for multi channel photodetector readout, provide a compact and efficient solution to both constraints. The readout system[16] selected for the detector includes an ASIC[17] capable of reading out up to 64 SiPM channels. The ASIC is connected to a set of dedicated PCBs that enable direct communication with a computer via Ethernet. The manufacturer also provides the source code of the DAQ software of the readout system, allowing system level customization. Despite its compact footprint of only a few centimeters, the chip integrates everything from signal amplification to digitization electronics. To reduce the required disk space, the ASIC incorporates an internal event detection and trigger system. Only time intervals surrounding validated events are processed, eliminating the need to continuously record waveform samples at high frequency over extended periods, as is currently done in the τSPECT experiment. To further limit data volume, the ASIC does not store full signal waveforms. Instead, it evaluates either the signal integral or the time for which the signal exceeds a predefined threshold. These two modes are referred to as charge to digital conversion (QDC) and time over threshold (ToT), respectively and are discussed in detail in Section 5.2. Additionally, the time at which the threshold crossing occurs is recorded. For each channel event, the ASIC transmits the channel number, the event time and the corresponding QDC or ToT value to the computer, resulting in a total data size of 80 bits per channel event [PET24]. In total, the ASIC provides 64 SiPM signal inputs, four SiPM bias voltage lines and three additional lines for readout of a temperature sensor.

To interface the detector with the ASIC, two custom PCBs were required. The first, referred to as SiPM PCB, hosts the 64 SiPMs and is directly attached to the light guide. The second, a so-called fan-out PCB, routes the SiPM signals to the ASIC and connects bias voltage and temperature sensor lines to the respective SiPM PCB lines. Because the ASIC requires convective cooling, it cannot be mounted directly onto the SiPM PCB, as the detector operates in vacuum. The electrical connection between the two PCBs is therefore established using four 20-line micro coax-cables[18]. Each cable carries 16 channel signals and one bias voltage line. Additionally, one of the temperature sensor lines is routed through three of the cables. The SiPM PCB is partitioned into four quadrants, each containing 16 SiPMs supplied by one of the bias lines. A photograph of the assembled PCB is shown in Fig. 4.24. To prevent ground loops, no electrical ground connection is present between the individual quadrants. A temperature sensor[19] is mounted at the center of the PCB underside, opposite the center of the SiPM array. It monitors the PCB temperature to assess whether environmental conditions allow for stable SiPM operation. A complete pinout table for the custom PCBs is provided in Appendix G.

4.8 Mechanical assembly

4.8.1 Detector stack assembly

A photograph of the fully assembled detector stack is shown in Fig. 4.25. During assembly, the ^{10}B-coated scintillator sheet and the SiPM PCB were attached to the top and the bottom surfaces of the light guide, respectively. The scintillator sheet

[16]PETsys TOF ASIC E. KIT
[17]PETsys TOFPET2 ASIC
[18]I-PEX-Cabline-VS 20P
[19]TI LMT86

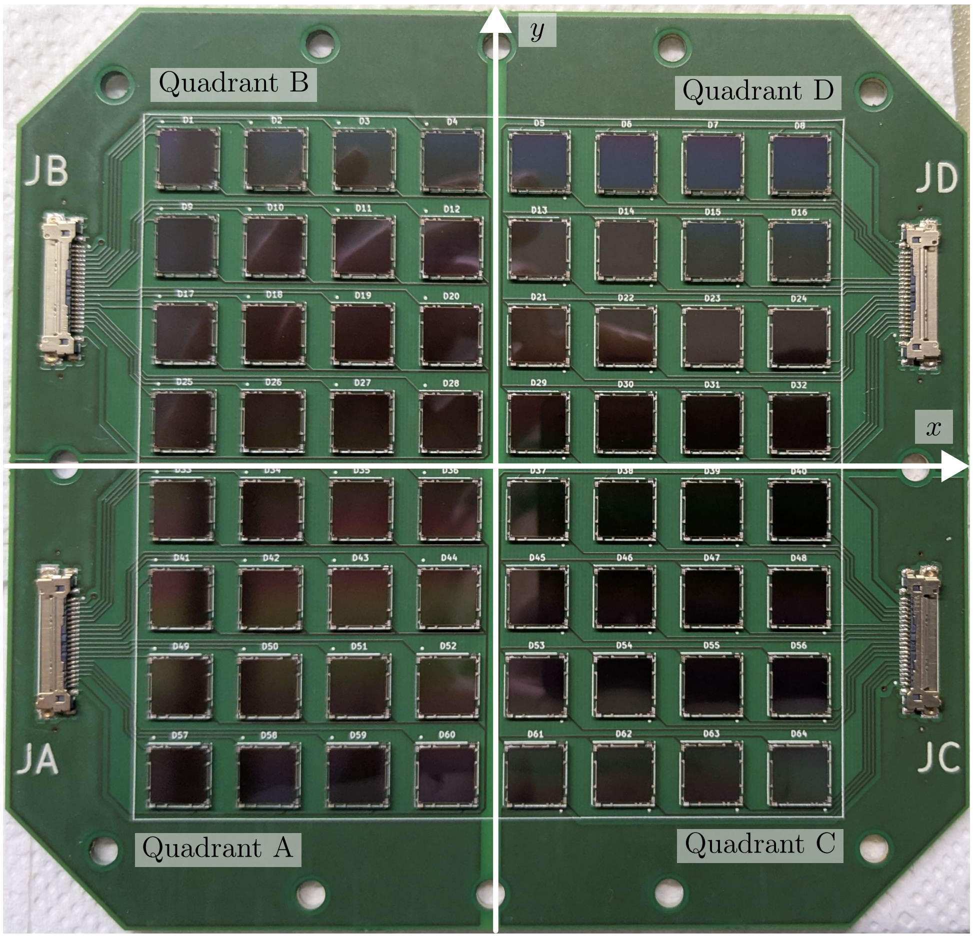

FIGURE 4.24: Fully assembled SiPM PCB of the detector. The PCB is divided into four quadrants, each labeled according to the corresponding bias voltage line from the ASIC PCB. To prevent ground loops, there is no ground connection between the quadrants. A temperature sensor is mounted in the center of the underside of the PCB. The SiPMs have a physical size of 7×7 mm^2 with a light sensitive area of 6×6 mm^2. The white lines indicate the detector coordinate system.

was bonded to the light guide using an optically transparent UV-curing glue[20] with a refractive index of 1.54 [Nor]. Compared to two component adhesives, UV-curing glue does not require mixing, significantly reducing the risk of introducing air bubbles into the bond layer. To prepare the gluing process, a boron coated scintillator sheet approximately 20 % larger than the light guide surface was used. The sheet was placed on a Plexiglass board with the boron-coated surface facing downward to protect the conversion layer from scratches and contamination. The Plexiglass board was selected due to its smooth surface and because it was covered by a protective film before use. The glue was applied in an X-shaped pattern to the uncoated side of the scintillator sheet, after which the light guide was pressed onto it. A few small air bubbles formed in the process. By applying a controlled pressure gradient across the light guide, most of the bubbles could be pushed out of the area between scintillator sheet and light guide. However, a few small air bubbles remained. Additionally, excess glue accumulated at the edges of the light guide and was removed using a piece of cleanroom tissue. Nevertheless, a thin, uneven layer of glue remained. After UV-curing, the excess scintillator sheet was trimmed with a sharp blade. In some locations near the

[20]NOA68

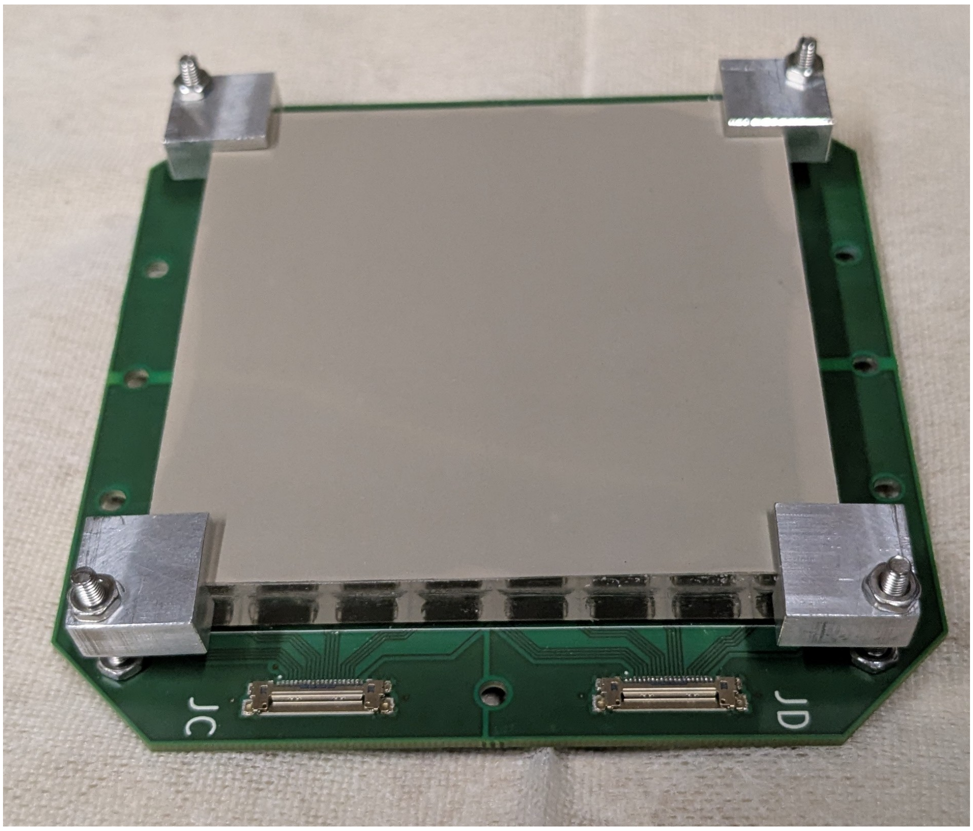

FIGURE 4.25: Photograph of the fully assembled detector. At the top the boron layer, which is coated onto the ZnS(Ag) scintillator sheet, is visible. The scintillator sheet is glued to the light guide, which has an outer dimension of approximately $80 \times 80 \times 10\ \mathrm{mm}^3$. Optical coupling between the light guide and the partially visible SiPMs is achieved using optical grease. The light guide is secured by four aluminum clamps attached to the SiPM PCB.

edges, the cutting process partially separated the bonded interfaces. The gluing procedure had been tested multiple times beforehand using uncoated polyethylene sheets and Plexiglass samples. Nevertheless, a small number of air bubbles, glue residue on the light guide edges and occasional bond separation during trimming could not be fully avoided. However, using an oversized scintillator sheet successfully prevented contamination of the boron surface by excess glue. The remaining layer of glue on the light guide edges and air gaps introduced by the cutting process may alter the optical properties near the light guide edges in a manner that is difficult to quantify, potentially degrading spatial resolution and detection efficiency in these regions. Similarly, residual air bubbles will locally reduce optical coupling efficiency. For the present prototype, the few remaining air bubbles with an estimated diameter of less than 1 mm are not expected to significantly affect overall performance. Moreover, due to the geometry of the prototype vacuum chamber (see Section 4.8.2), most detected events are anticipated to occur away from the detector edges. Consequently, the thin glue layer at the edges is not considered critical for the prototype either. For future assemblies, the gluing procedure should be further optimized. Another issue is that the adhesive used to bond the scintillator sheet to the light guide is not compatible with cryogenic environments. A different glue[21] compatible with these environments is available and was used for the τSPECT detector. However, it is a two component

[21] Acrifix 2R 0190

glue and even more prone to introducing imperfections in the bond layer. A significant improvement could therefore be achieved if both the scintillator and the ^{10}B layer could be coated onto the light guide directly, eliminating the need for adhesive bonding.

After gluing the scintillator sheet, the light guide was mechanically fixed to the SiPM PCB using aluminum clamps (see Fig. 4.25). Optical coupling between the light guide and the SiPMs was improved by applying optical grease[22] to each SiPM prior to mounting. The presence of small air bubbles within the grease cannot be excluded, as most SiPMs are no longer visible once the scintillator sheet is attached. The aluminum clamps, which are connected to the PCB, press the light guide onto the SiPM grid ensuring stable mechanical and optical contact.

4.8.2 Vacuum chamber

To characterize the detector prototype, a dedicated vacuum chamber was constructed. The setup must allow connection to a UCN guide both vertically and horizontally and must accommodate radioactive sources for test measurements. Additionally, feedthroughs for a total of 80 lines are required (see Section 4.7). A cross-sectional view of the main flange is shown in Fig. 4.26 and a photograph of the detector vacuum chamber is shown in Fig. 4.27. The main detector flange is a DN160CF aluminum blind flange. A 70.5 mm circular opening was milled into its center to accommodate a 70 mm UCN beam guide. Towards the interior of the flange, the diameter reduces by 4 mm, forming a mechanical stop that prevents accidental contact between the beam guide and the detector surface during assembly. The outer side of the flange contains threaded holes for mounting a Wilson flange, enabling a vacuum tight connection between the detector chamber and the beam guide. Aluminum was selected to minimize the overall chamber weight. The detector is mounted on the inside of the main flange using M3 spacers that attach the SiPM PCB to four dedicated threaded holes. The remainder of the chamber consists of a DN160 CF to ISO-K adapter, a DN160 ISO-K to DN63 adapter and a DN63 ISO-K T-piece. These commercially available standard components provide sufficient volume for pumping ports and electrical feedthroughs. The electrical feedthrough is implemented using two PCBs glued into slits machined into an ISO-K 63 blind flange with black epoxy[23]. This configuration provides vacuum

[22] Eljen technology EJ-550
[23] Loctite Stycast 2850FT CAT 24LV

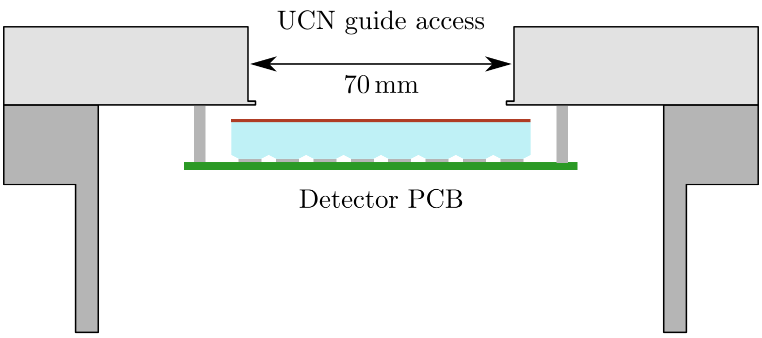

Figure 4.26: Cross sectional schematic of the main flange of the detector vacuum chamber. The flange includes an access port for 70 mm UCN guides, featuring a mechanical stop to prevent accidental contact with the detector surface. The detector assembly is mounted on the vacuum facing side of the flange.

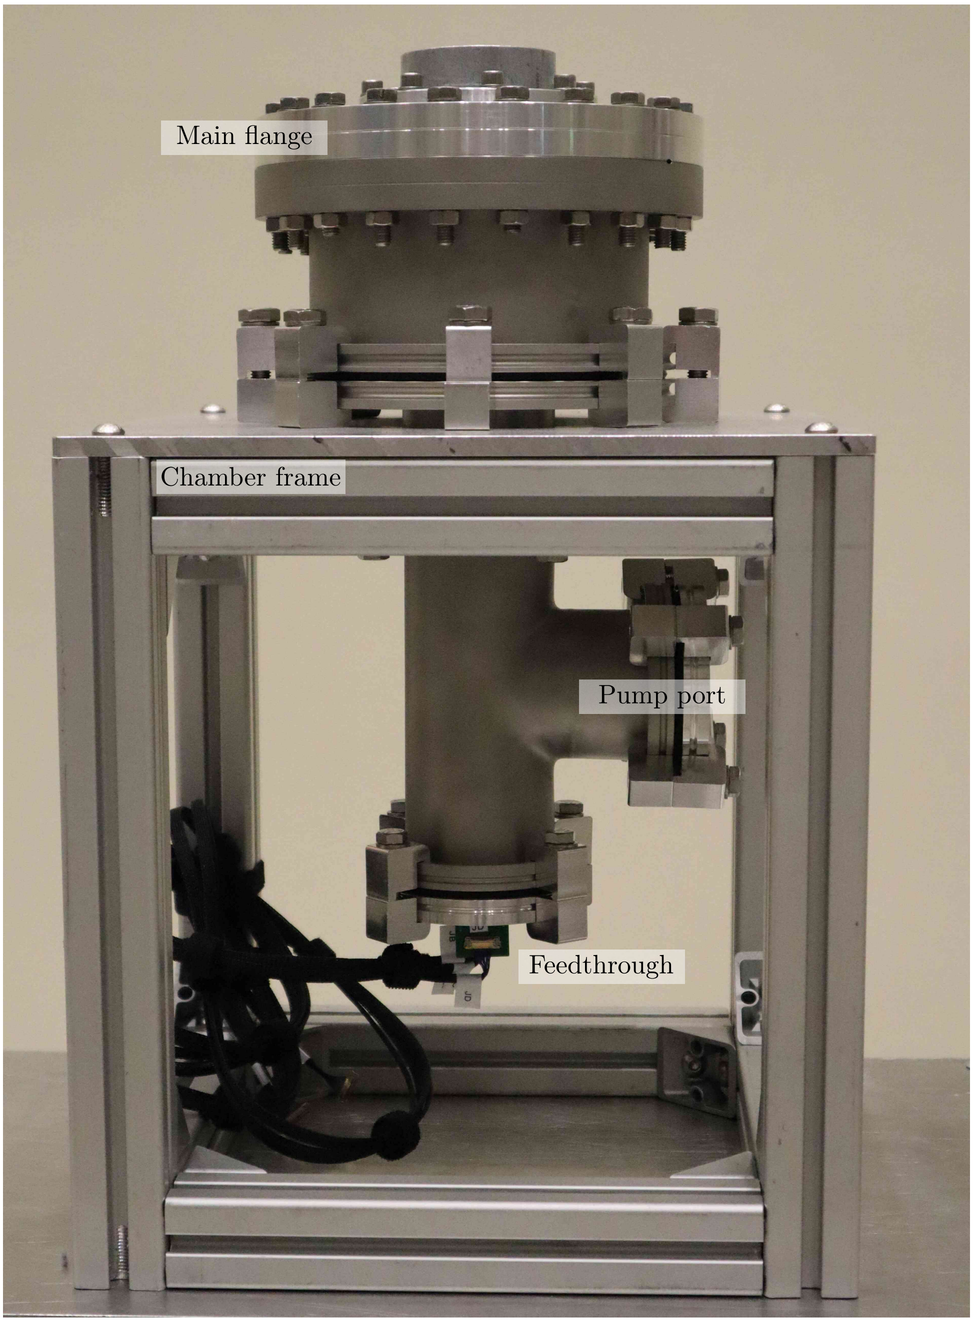

FIGURE 4.27: Photograph of the detector vacuum chamber mounted within an aluminum extrusion support frame. The main flange, to which the detector is attached on the vacuum side, is visible at the top of the image. The UCN beam guide port is covered by a blind flange in this image. The feedthrough flange, including the feedthrough PCBs, is shown at the bottom. A cross sectional schematic of the main flange is provided in Fig. 4.26.

tight and light tight sealing. The PCBs, originally developed for a different project by another collaboration member, provide 40 impedance-controlled lines at $50\,\Omega$ for two 20-line micro-coaxial cables each. Four 115 cm micro-coaxial cables connect the SiPM PCB to the feedthrough assembly. This length allows the main detector flange to be positioned adjacent to the chamber without having to disconnect cables first.

Chapter 5

Detector Characterization

After the detector was successfully assembled detailed characterization and test measurements were conducted to evaluate it's performance. This chapter describes the experimental setups used for the characterization measurements, the readout system settings and the results obtained from these measurements.

5.1 Characterization setups

To characterize the detector, both an ^{241}Am α-source and the UCN source at beam port C of the research reactor TRIGA Mainz were used. The ^{241}Am source is a custom-made source produced by a researcher from another group at the research reactor TRIGA Mainz at JGU. It consists of a 1 mm thick aluminum plate with dimensions of 2×2 cm^2, onto which a grid of droplets from an ^{241}Am solution is deposited. The technique used to produce this source is described in detail in [Haa+17]. In total, the source has an activity of approximately 33 kBq, distributed across the droplet grid on its surface. To avoid contamination of the detector surface through direct contact between source and detector, the source was mounted in a 3D-printed holder that creates a separation of approximately 2 mm between source and detector and smooths out spatial variations in the irradiation caused by the droplet structure of the source. The dominant decay channels of ^{241}Am produce α-particles with a kinetic energy of approximately 5.48 MeV [Bas06], corresponding to a range of about 4 cm in air [Ber+05]. Consequently, the 2 mm gap is not expected to stop a large fraction of the emitted α-particles at atmospheric pressures. The combined size of the source and the 3D-printed holder is sufficiently small to allow it to be placed on the detector surface when the vacuum chamber is oriented such that the surface faces upwards. A blind flange was used to provide a light-tight seal of the main detector flange. With this setup, tabletop characterization measurements can be performed without requiring access to a UCN source.

However, the majority of the characterization measurements were performed using UCNs at beam port C of the research reactor TRIGA Mainz. A detailed description of the reactor can be found in [EK00] and is summarized in the following paragraph. A simplified schematic cross section of the UCN source is shown in Fig. 5.1. Unlike the PSI source, the TRIGA source uses nuclear fission in the TRIGA reactor as its neutron source rather than a spallation source. The reactor employs so-called fuel-moderator elements composed of an alloy of uranium-zirconium-hydride (UZrH), which generate fast neutrons. Due to the hydrogen content of the fuel elements, they simultaneously act as a moderator and together with the water in the reactor pool, thermalize the fast neutrons. As in the UCN source at PSI (see Section 3.4), these thermal neutrons are subsequently converted to cold and ultra-cold neutrons in the deuterium crystal. In contrast to the PSI source, however, the TRIGA source does

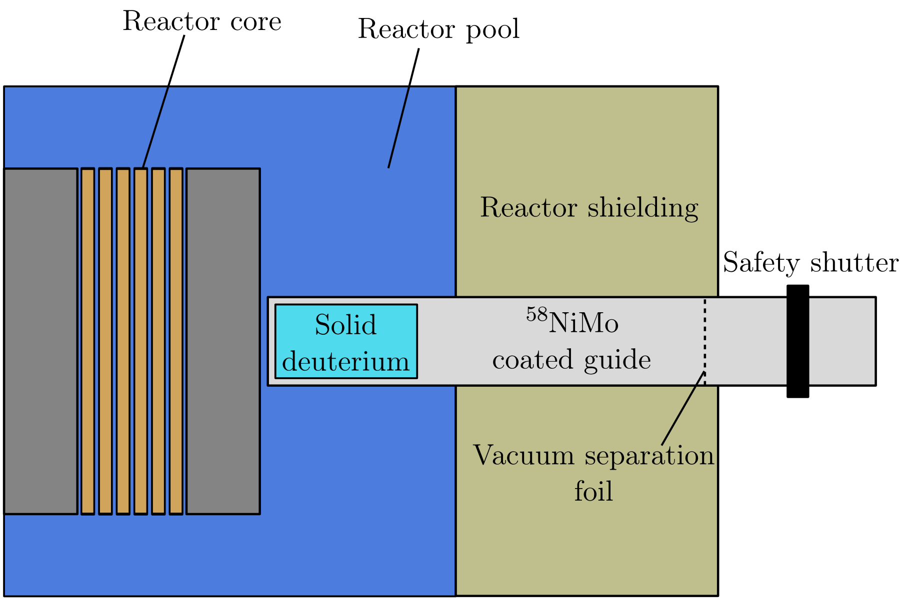

Figure 5.1: Simplified schematic cross section of the UCN source at beamport C of the research reactor TRIGA Mainz. Thermal neutrons originating from the reactor core are converted to UCNs in the solid deuterium crystal. The UCNs are subsequently delivered to the beam port via a ^{58}NiMo (85:15) coated guide. A vacuum separation foil and a safety shutter are installed to separate source and experiment vacuum. In addition, the safety shutter can also be used to interrupt the UCN flux if required.

not include a storage vessel and supplies only a single beam port. As a result, UCNs propagate directly towards the beam port after exiting the deuterium crystal. Because of the Fermi potential of solid deuterium (107 neV, see Appendix A), neutrons exiting the deuterium crystal experience a kinetic energy boost of 107 neV. Since the beam port is at the same height as the deuterium crystal, this value also corresponds to the minimum kinetic energy at the beam port. The TRIGA reactor can be operated in two modes: pulsed mode and steady state (DC) mode. In pulsed mode, the TRIGA reactor is typically operated with a pulse period of 12 min, producing pulses with a total thermal energy of approximately 10 MWs and a full width at half maximum of about 30 ms. This mode yields a high UCN count rate over a short time interval and is therefore particularly well suited for UCN experiments relying on UCN storage for extended periods of time. For the characterization measurements presented here, pulsed operation was used to assess the detector performance at high count rates. In DC mode, the reactor is operated at a constant thermal power of 100 kW. Based on previous unpublished measurements with other detector systems, a constant count rate of approximately 500 s^{-1} can be expected under these conditions. This mode was used for most characterization measurements, as it allows for the accumulation of a larger number of events under varying test conditions.

The experimental setup used for the characterization measurements is shown in Fig 5.2. To guide the UCNs from the beam port to the detector, a series of stainless steel UCN guides was employed. The standard guides available in Mainz have an outer diameter of 70 mm, whereas the guide at beam port C has an outer diameter of 80 mm for historical reasons. Therefore, an adapter flange narrowing the diameter from 80 mm to 70 mm is required. Downstream of the adapter flange, an S-shaped beam guide section consisting of two 45° bends and a straight guide was used to guide the UCNs upwards

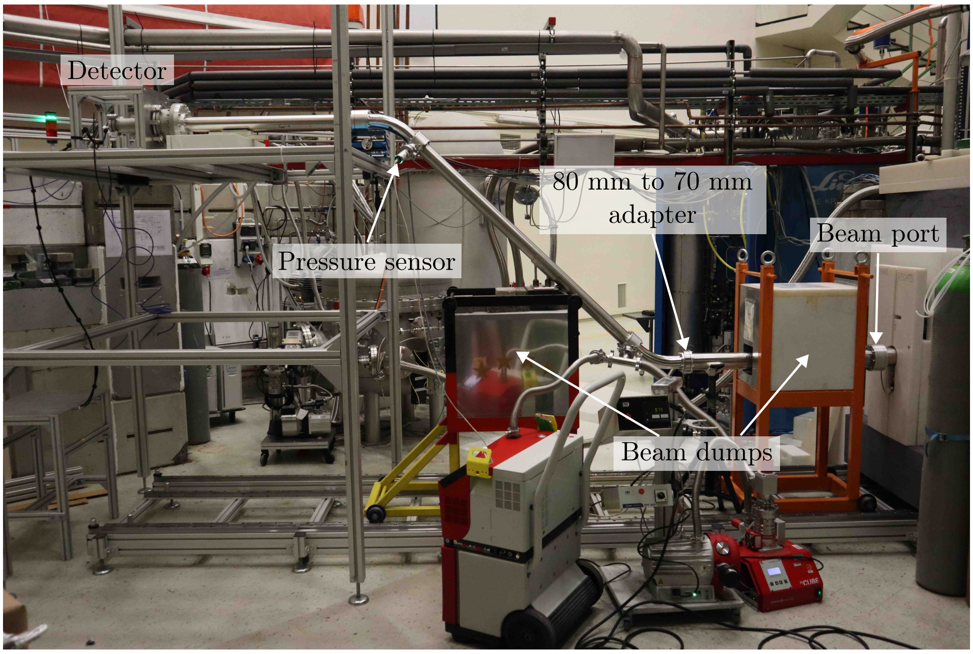

Figure 5.2: Photograph of the measurement setup used for the characterization measurements at beam port C of the research reactor TRIGA Mainz. UCNs exit the source at the beam port on the right-hand side and are guided towards the detector by stainless steel guides with an outer diameter of 70 mm. An S-shaped section containing two 45° bends guides them upwards by 125 cm. A pressure sensor is mounted near the top. Two beam dumps are used to shield the experimental area from radiation.

by about 125 cm. This vertical displacement decelerates the neutrons gravitationally by about 128 neV, compensating for the minimum kinetic energy of 107 neV at the beam port. The detector was connected to the top of the S-shaped section using a straight guide with a length of 90 cm. A shutter that can be used to separate the beam port from the experimental vacuum, is installed at the beam port for safety reasons and is therefore referred to as safety shutter. Closing this shutter allows the UCN flux to be interrupted while the reactor remains in operation. This capability enables dedicated background measurements to estimate the ratio of UCN to background events. The beam guides were evacuated using a turbomolecular pump in combination with a roughing pump, both connected to a pumping port at the lower 45° bend. To monitor the vacuum conditions in the beam guides, a pressure sensor was installed at a second pump port at the upper 45° bend. After approximately 12 h of pumping, a typical pressure of 1.5×10^{-4} mbar was achieved. In addition to UCNs from the UCN source, the reactor also produces a continuous flux of thermal neutrons and γ-radiation, which constitute background for the characterization measurements. To reduce the thermal neutron flux at the detector location, additional shielding composed of cadmium plates and polyethylene bricks containing a small fraction of boron were installed around the detector.

5.2 Signal processing

5.2.1 Readout system architecture and measurement modes

This section provides an overview of the configurable parameters and operating principles of the SiPM readout system (for technical details, see ASIC datasheet [PET24]). A schematic overview of the signal processing chain within a single ASIC channel is shown in Fig. 5.3. Each channel of the SiPM readout system ASIC comprises three parallel processing branches, referred to as T, E and Q branches. A replica of the SiPM signal S is fed into each branch and amplified by an adjustable gain factor (G_{T}, G_{E} and G_{Q} respectively). The T branch contains two discriminators referred to as T1 and T2 discriminator, which produce a logic high when the input signal exceeds the discriminator thresholds $V_{\mathrm{th,T1}}$ and $V_{\mathrm{th,T2}}$, respectively. The E branch contains a third discriminator with threshold $V_{\mathrm{th,E}}$. In contrast to the T branch the E branch operates at a gain that is about one order of magnitude lower, enabling it to process stronger signals that would saturate the T branch. The Q branch contains a charge-to-digital converter (QDC) that integrates the input signal. Depending on the ASIC configuration, either a fixed or a variable integration window can be selected. Additionally, two time-to-digital converters (TDCs), referred to as T TDC and E TDC, are available and provide a timing precision of approximately 30 ps [PET24]. The T TDC timestamps the rising edge of the input signal and defines the channel event

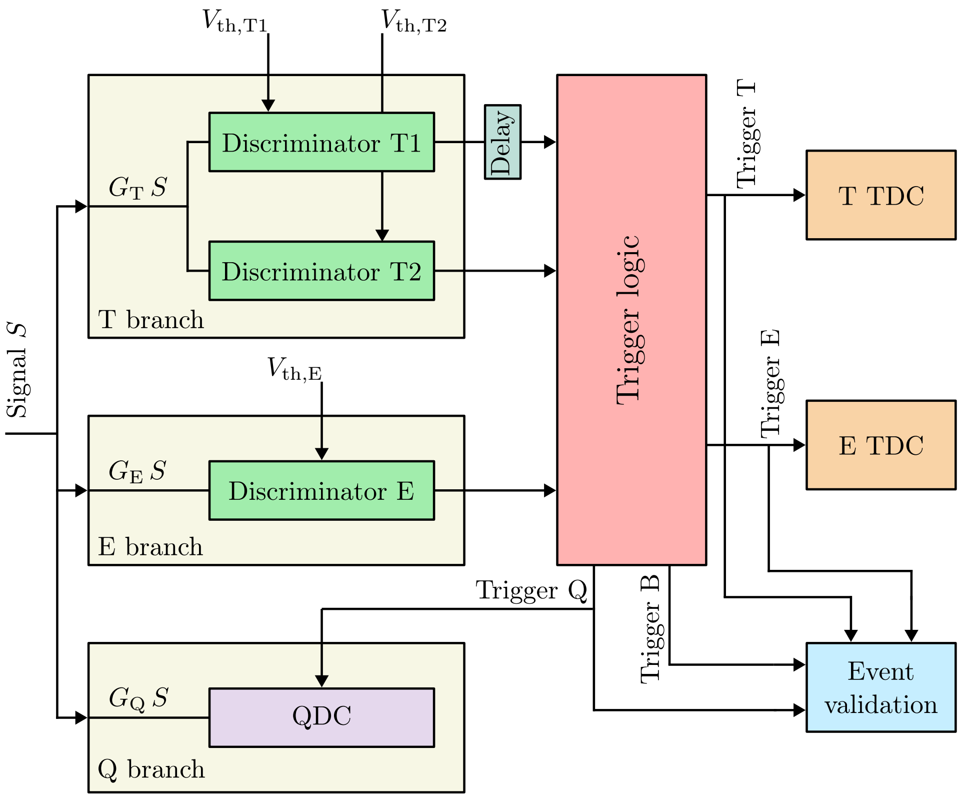

Figure 5.3: Schematic depiction of the internal logic of a single SiPM readout system channel. A replica of the SiPM signal is fed into branches T, E and Q, where it is amplified by adjustable gain factors. The outputs of the three discriminators with configurable thresholds are processed by the programmable trigger logic, which generates four trigger signals. These trigger signals control QDC and TDCs of the respective channel and determine whether a signal is accepted as a valid event. For further details refer to the main text.

time, while the E TDC records the time at which the signal falls below the selected threshold. In combination with the 200 MHz system clock, the TDC outputs allow the ASIC to determine the time over threshold (ToT). While the T TDC is always active, only one of the QDC and the E TDC can be used at a given time. This choice is made on a per-channel basis, defining the operating modes referred to as QDC mode and ToT mode, respectively. The discriminator outputs are processed by a programmable trigger logic, which generates four trigger signals referred to as triggers T, E, Q and B. Trigger B is asserted as soon as the input signal exceeds any discriminator threshold, indicating that a potential channel event has occurred and is deasserted once the signal falls below all thresholds. Triggers T, E and Q activate T TDC, E TDC and QDC, respectively and can be configured to respond to different logical combinations of the discriminator outputs. When trigger B transitions to logic low, the channel enters the validation stage. If triggers T, E and Q were active at any time while trigger B was asserted, the event is accepted and digitized. In QDC mode the E TDC is disabled and in ToT mode trigger Q is ignored.

The described scheme allows the readout system to be adapted to a wide range of use cases. Multiple discriminator thresholds can e.g. be employed to achieve precise timing near the baseline, define the start of signal integration and ensure that only sufficiently strong signals are digitized. Additionally, a feature, referred to as fast dark count rejection, is available. It relies on a configurable delay line inserted between the output of discriminator T1 and the trigger logic, which can be set to precisely known delays between 0 ns (bypassed) and 8 ns. By requiring that trigger T is activated only when discriminator threshold T1 and one of the other discriminator thresholds are exceeded, precise event timing is maintained while only sufficiently strong signals trigger the T TDC. This approach prevents logic dead time that would otherwise occur during the T TDC reset. However, precise timing is only ensured, if the signal crosses the second threshold within the configured delay interval. A further issue with single threshold operation is that noise on the signal tail can cause multiple re-triggers while the signal fluctuates around the threshold. By employing a dual threshold scheme, the signal must fall below both thresholds and subsequently cross both thresholds again before a new trigger is generated, effectively suppressing such behavior. In addition to channel specific parameters, there is also a number of global parameters that apply to all channels. These include settings that define the usable range of the digital-to-analog converters (DACs) setting the discriminator thresholds. For the initial test measurements these parameters were set to values providing the maximum dynamic range specified in the datasheet, allowing the full discriminator threshold range to be utilized. Furthermore, the trigger logic was configured to use a dual trigger threshold scheme based on discriminators T1 and T2 with fast dark count rejection, minimizing re-triggers caused by noisy signal tails. The T branch gain G_{T} was kept at its nominal value, which is maximum gain. Additionally, a large signal gain was ensured by setting the SiPM overvoltage to 4.5 V for all channels, close to the maximum recommended overvoltage of 5 V [Sem22]. A breakdown voltage of 24.5 V was assumed for all SiPMs, corresponding to the expected mean value at room temperature [Sem22]. Given that the breakdown voltage spread of the selected set of SiPMs is only 76 mV (see Section 4.6.2), this assumption is not expected to introduce significant gain variations between channels. In combination, these settings result in a gain close to the maximum achievable gain, maximizing the probability of detecting events during the initial test phase.

Finally, the measurement mode was selected. The available QDC integration windows range from (250 ns, 500 ns) to (1500 ns, 2100 ns), depending on the configured Q branch gain [PET24]. In the ToT mode, the coarse timing is obtained by polling a 10-bit counter that increments with every clock cycle, i.e. every 5 ns. The counter therefore wraps around every 5120 ns, which also defines the maximum measurable ToT. The lower limit of the dynamic range in the ToT mode is not explicitly specified in the datasheet and is likely constrained primarily by the precision of the ASIC timing setup. Consequently, the ToT mode offers a substantially larger dynamic range in terms of signal width than the QDC mode. Because measurements of signal width are of particular interest for assessing the optimal measurement mode, the ToT mode was selected for the initial test phase. For the first test measurements, the detector was installed in the TRIGA reactor setup (see Section 5.1). Prior to recording neutron signals, an ASIC calibration was performed using the software provided by the readout system manufacturer. During the calibration, the TDCs, QDCs and discriminators of all channels are calibrated. The resulting calibration parameters are stored in dedicated files and are used by the software to correct the data transmitted by the ASIC during normal operation. Notably, the calibration procedure also determines the baseline level of each channel in units of the discriminator threshold DACs. This enables the software to set the discriminator thresholds relative to the baseline value, ensuring consistent threshold settings across all channels.

5.2.2 Time over threshold correction

When recording neutron signals with the detector prototype, the raw ToT distribution shown in Fig. 5.4 was obtained. The distribution exhibits three distinct regimes (A, B and C), with transitions marked by the dashed vertical lines. This structure originates

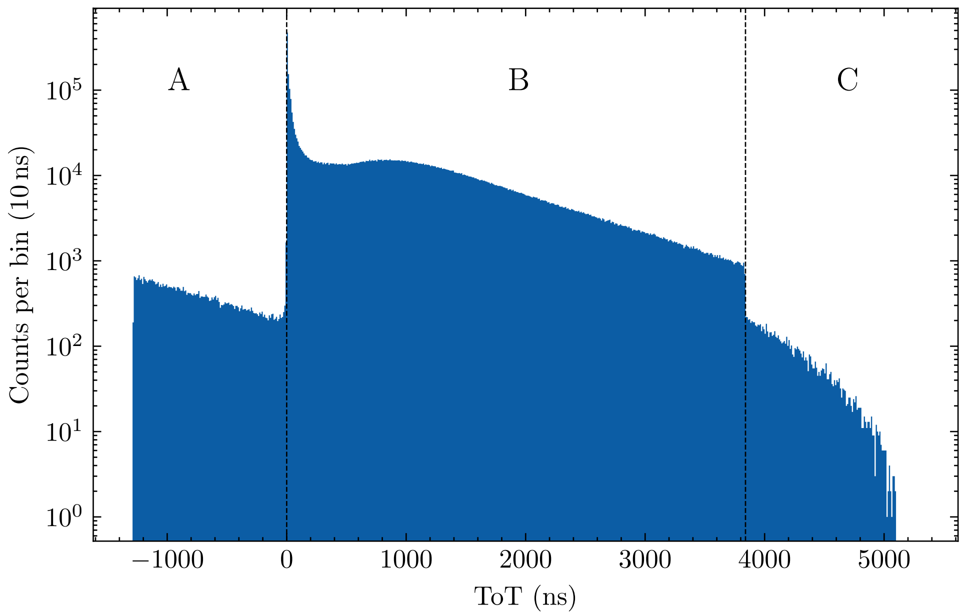

FIGURE 5.4: Raw time over threshold (ToT) distribution obtained from UCN signals in dual threshold mode with $V_{\mathrm{th,T1}} = 2\,\mathrm{DAC}$ units and $V_{\mathrm{th,T2}} = 13\,\mathrm{DAC}$ units. The three distinct regimes are indicated by the dashed vertical lines and labeled by A, B and C, respectively.

from the way the readout system software processes the ASIC output[1]. For each event, the ASIC polls the 10-bit timing counter twice: once when the signal crosses the threshold (t_{start}) and once when it falls below the threshold again (t_{end}). The software computes the raw ToT T_{raw} by subtracting these counter values, which can yield a negative result if the counter wrapped around during the event. In such cases, the ToT is corrected by adding the maximum counter value, corresponding to 1024 clock cycles or 5120 ns. The corrected ToT T_{corr} is therefore given by

$$T_{\text{corr}} = t_{\text{end}} - t_{\text{start}} + 5120\,\text{ns} \tag{5.1}$$

Subsequently, the ToT is further refined using the TDC outputs in combination with the ASIC calibration files. By default, however, the software applies the wrap around correction only if

$$t_{\text{end}} - t_{\text{start}} < -1280\,\text{ns} \tag{5.2}$$

applies. This condition is intended to alert the user to events approaching the upper limit of the ToT dynamic range and explains the three observed distinct regimes. In regime B, the readout system behaves as intended and negative ToTs are corrected. Regimes A and C contain events with ToTs exceeding 3840 ns. For events in regime A, a counter wrap around occurred during the event, whereas this was not the case for events in regime C. Accordingly, the corrected ToT distribution shown in Fig. 5.5 is obtained by applying the correction in Eq. (5.1) to all negative ToT values. Unlike the raw distribution, the corrected distribution shows no signs of abrupt regime transitions. However, a pronounced increase at ToTs above approximately 5000 ns remains, which deviates from the overall trend. A plausible explanation for this behavior is that these values result from the additional TDC based correction applied by the software. For events with very short ToTs, this correction can yield a negative

[1] This information was obtained via private communication with the manufacturer's customer support and verified in the source code that was shipped together with the readout system.

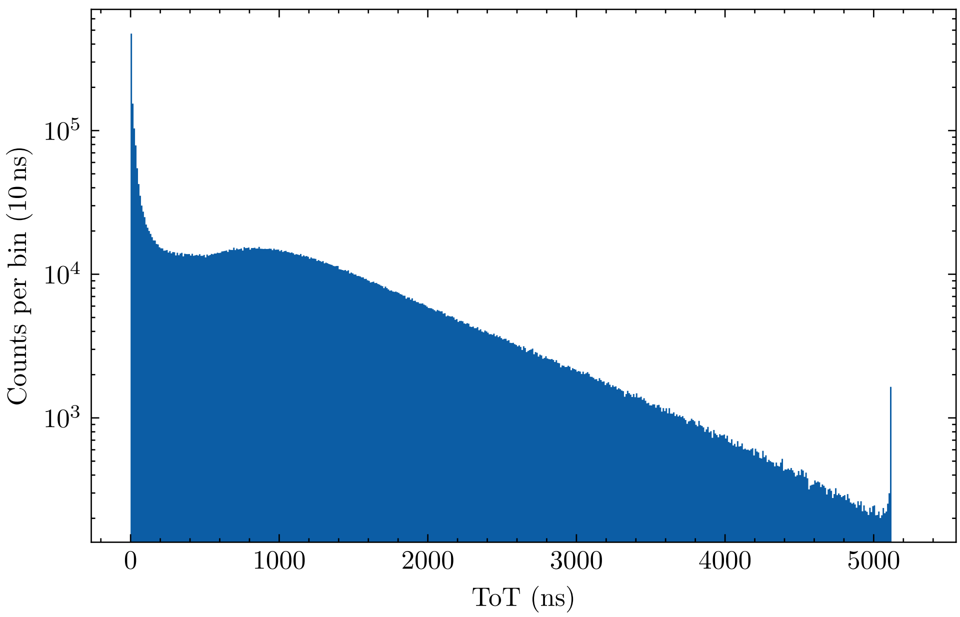

FIGURE 5.5: Corrected time over threshold distribution for neutron signals recorded in dual threshold mode with $V_{\text{th,T1}} = 2\,\text{DAC units}$ and $V_{\text{th,T2}} = 13\,\text{DAC units}$. The dataset was acquired over a measurement time of 990 s.

value, that is subsequently interpreted as a long ToT, because negative values are mapped into the long ToT region by the wrap around correction. To mitigate this effect, only ToTs smaller than $-50\,\text{ns}$ were corrected, while events with ToTs in the interval $(-50\,\text{ns}, 0\,\text{ns})$ were discarded as invalid.

5.2.3 Discriminator configuration and neutron event definition

In the following, neutron events are identified from individual SiPM channel events using a set of selection criteria. Channel events are first required to satisfy a minimum signal width in terms of ToT. Neutron events are then identified as temporal clusters of such channel events occurring within a short time window and involving multiple SiPM channels. The specific parameter choices for the ToT cut, clustering time window and minimum number of active channels are motivated and evaluated in the remainder of this section.

As evident from the ToT distribution, signals spanning the full dynamic range of the ToT mode are recorded. Extrapolating from the tail of the distribution in Fig. 5.5, a fraction of events is expected to exceed this range. For such events, the timing counter wraps around more than once, causing the ToT to be misinterpreted as significantly shorter than its true value. This effect could be reduced by raising the discriminator threshold, but at the expense of losing weaker signals e.g. from SiPMs located farther from the neutron event position. To evaluate different threshold configurations, the distribution of the number of active channels per neutron event was analyzed. For this assessment, neutron events were required to involve signals in at least two SiPM channels. This requirement effectively implements a coincidence criterion that strongly suppresses noise. Multi-channel events are identified by clustering channel events in time, based on their start times t_s. A channel event is assigned to the same cluster if the condition

$$t_{\text{s},i} - t_{\text{s},i-1} < t_\text{th} \tag{5.3}$$

is satisfied. Here, $t_\text{th} = 5\,\mu\text{s}$ was chosen, as most ToTs are of the order of $1000\,\text{ns}$. This choice also accounts for rare cases in which a channel is re-triggered by noise after the signal has fallen below the threshold. In such cases, the ToT of initial and re-triggered channel events were summed. While channel events with ToTs exceeding the chosen clustering time window are observed, such cases are rare and typically confined to individual channels. In these situations, a delayed re-trigger would affect only a single channel and therefore fail the minimum channel number requirement for neutron event identification. Consequently, potential re-triggering of long ToT channels is very unlikely to lead to the misidentification or splitting of neutron events. During the test measurements at the TRIGA reactor, a count rate of about $500\,\text{s}^{-1}$ was expected, corresponding to an average inter event spacing of $2\,\text{ms}$. It is therefore highly unlikely that channel events from different neutron events are merged into a single cluster. Cluster assignment was terminated once the condition defined by Eq. (5.3) was violated. The number of active SiPM channels per event was then extracted and compared for different discriminator threshold configurations. For three representative threshold configurations, both the active channel distribution and the ToT distribution are shown in Fig. 5.6. For event position reconstruction, a minimum of three active channels is required (see Section 5.4.1). Lowering the trigger threshold increases the number of active channels but eventually leads to a substantial rise in noise events, which predominantly populate the two channel bin. This behavior is evident for the $V_{\text{th,T1}} = 2\,\text{DAC units}$, $V_{\text{th,T2}} = 10\,\text{DAC units}$ configuration. Overall the $V_{\text{th,T1}} = 2\,\text{DAC units}$, $V_{\text{th,T2}} = 13\,\text{DAC units}$ configuration represents a reasonable

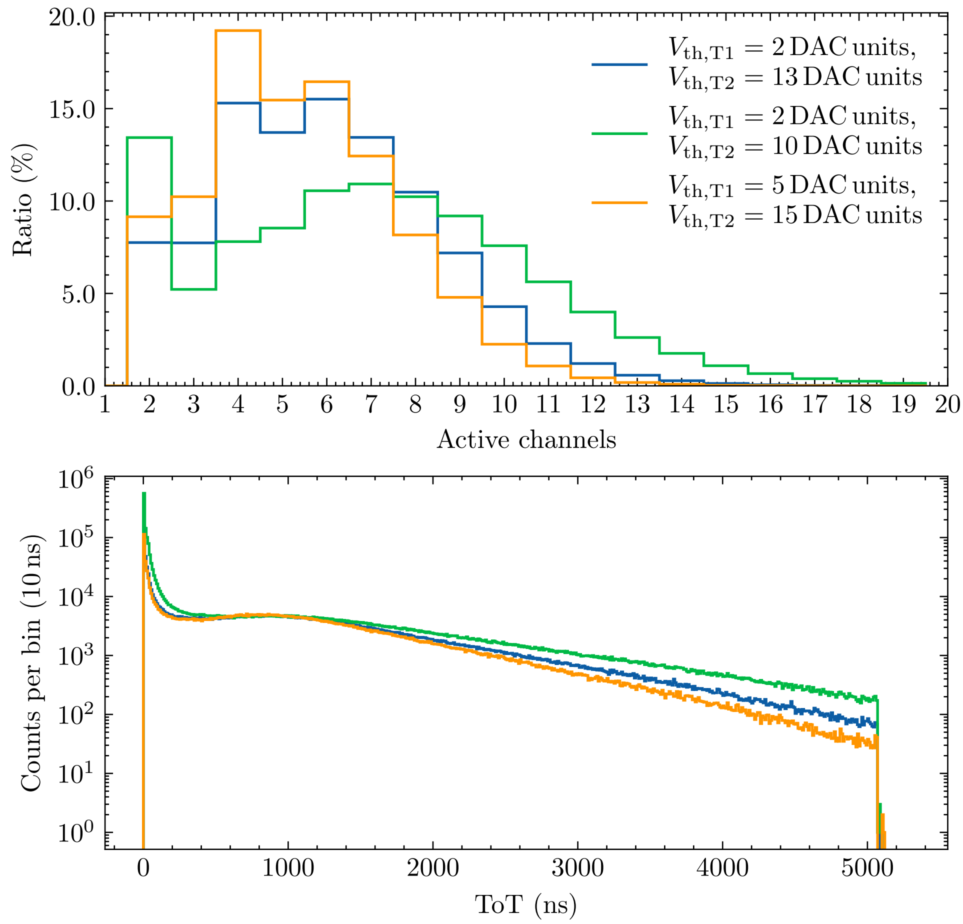

FIGURE 5.6: Active channel and time over threshold distribution for different trigger threshold configurations. A minimum of two active channels is required for an event to be classified a neutron event. For each configuration, neutron data were recorded for 300 s.

compromise, exhibiting the smallest fraction of two channel events and was therefore used for all subsequent characterization measurements. Due to limited beam time, no dedicated optimization of the T1 threshold was performed. Given the role of the dual-threshold logic discussed above, moderate changes of the T1 threshold are not expected to significantly affect the presented results. As indicated by the ToT distribution, a fraction of channel events exceeding the dynamic range of the ToT mode is still expected with the chosen configuration. For the initial characterization, however, maximizing the number of reconstructible neutron events was prioritized over completely avoiding such cases. Moreover, these events are expected to be associated with neutron events producing signals in many channels. In principle, some of these cases could therefore be identified by analyzing the spatial distribution of ToTs across neighboring channels.

To further refine the event identification and assess noise contributions to neutron events, the ToT distribution associated with neutron events was examined. This distribution includes only channel events that satisfy the neutron event definition introduced above and is shown in Fig. 5.7. To preserve visibility of the broad peak around 1000 ns, only channel events with ToTs greater than 50 ns were included. The distribution exhibits a pronounced dip between a sharp peak at small ToTs and the

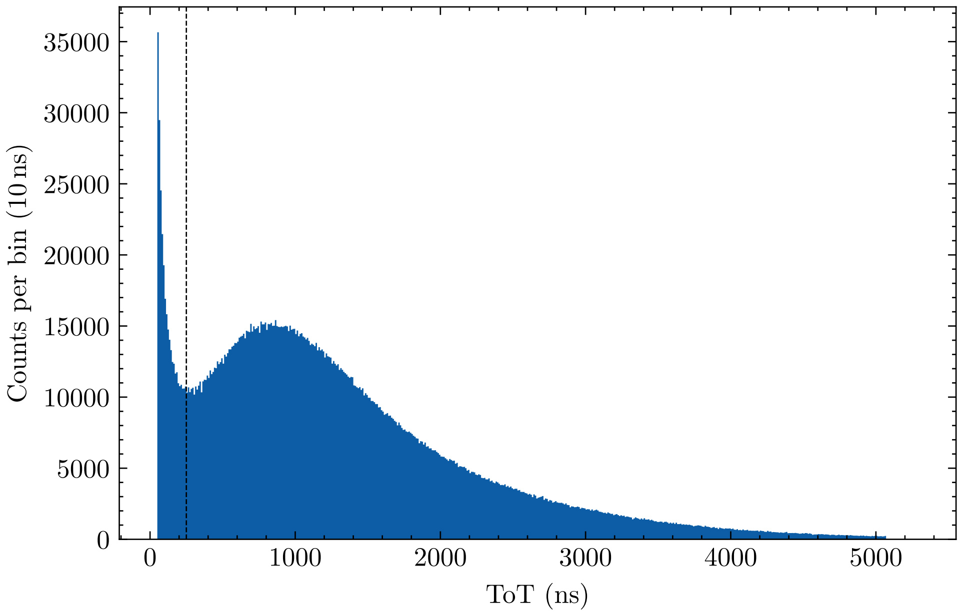

Figure 5.7: Time over threshold distribution for neutron events. The dashed vertical line marks a ToT of 250 ns.

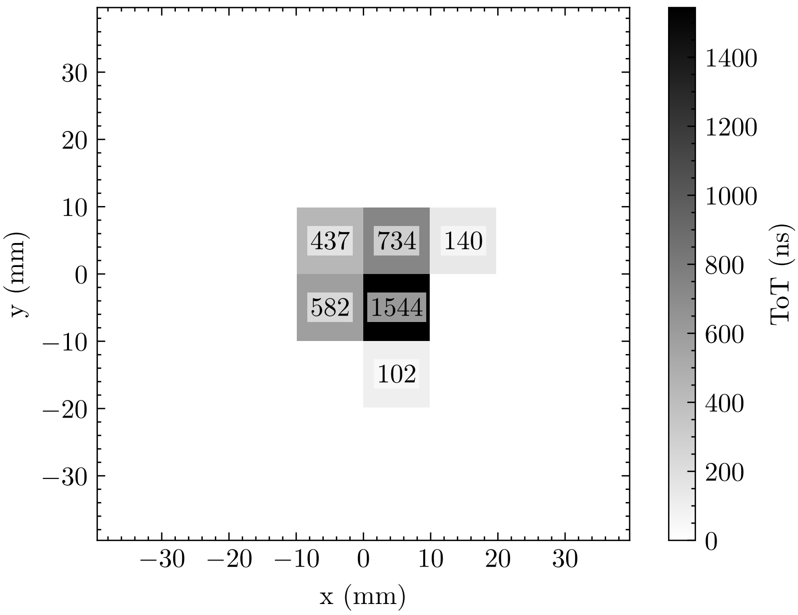

Figure 5.8: Spatial distribution of active SiPM channels for a representative neutron event. The squares represent the effective area covered by an SiPM channel. The numerical values denote the measured ToT in nanoseconds and are encoded by the grayscale intensity.

broad maximum at larger ToTs. This structure suggests the presence of two classes of signal shapes with markedly different widths. One possibility is that the short ToT component originates from noise processes that produce signals unrelated to neutron interactions. If such noise were uncorrelated with neutron events (e.g. a γ-event), a spatially disjoint channel pattern consisting of a neutron event cluster and isolated noise hits would be expected. Although such patterns are observed, their occurrence rate is not sufficient to explain the strong peak at small ToTs. Instead, channel events with small ToTs are typically part of spatial clusters attributed to neutron events, as shown in Fig. 5.8. This observation indicates that most short ToT signals are correlated with neutron events. One possible explanation is that channels near the neutron associated cluster receive insufficient light to exceed the threshold, but nevertheless get close enough to the threshold to be triggered by signal fluctuations. In this case, the measured ToT reflects the properties of a noise spike rather than the neutron induced signal and therefore carries no meaningful event information. An alternative explanation is that the scintillator response consists of multiple components with significantly different decay constants. In such a scenario, only a fast component would exceed the threshold for weak signals, whereas both fast and slow components contribute for stronger signals. However, analysis of typical neutron signal waveforms (see Section 5.3) indicates that short ToTs are most likely caused by noise. Consequently, a ToT cut at 250 ns was introduced to reject these channel events. This value was chosen because it lies at the center of the dip separating the two maxima in the ToT distribution.

5.2.4 Summary and conclusion

In summary, the initial test measurements demonstrate that the observed scintillator signal widths extend well beyond the dynamic range of the QDC mode and in some cases even beyond that of the TOT mode. As a result, the ToT mode represents the only viable measurement mode and was used for all subsequent measurements. Moreover, a neutron event identification algorithm based on clustering channel events in time was established and reasonable discriminator threshold settings were identified. Based on the considerations above, neutron event identification proceeds as follows:

1. Only channel events with a ToT exceeding 250 ns are considered.
2. Channel events are grouped into clusters based on their start times, with successive events separated by less than 5 µs assigned to the same cluster. If a channel is re-triggered within this interval, the corresponding ToTs are summed.
3. Only clusters containing at least two active channels are accepted as neutron events.

While the ToT mode accommodates a wider range of signal widths than the QDC mode, its primary disadvantage is that the ToT is not directly proportional to the amount of scintillation light detected by an SiPM channel. Such information is, however, required for accurate event position reconstruction. One possible approach to address this limitation is to record full signal waveforms and extract an average signal shape. Using this information, the signal integral could then be estimated from the measured ToT. Waveform measurements would also provide insight into the underlying signal time constants, which the ToT distribution already suggest to differ significantly from the values stated in the datasheet. Furthermore, detailed waveform information is useful for verifying the interpretation that the sharp peak at small ToTs originates from noise rather than from distinct scintillator decay components. This signal shape investigation is described in Section 5.3.

5.3 Signal shape

5.3.1 Shape model

Both SiPM microcell signal shape and scintillator light output as a function of time exhibit a characteristic exponentially decaying tail (see Sections 4.4 and 4.6). Therefore, both can be modeled by an exponentially decaying function

$$f_{\text{exp}}(t) = A \exp\left(-\frac{t}{\tau}\right) \tag{5.4}$$

under the assumption that the rise time of the signal is negligible compared to τ. Here, $A > 0$ denotes the amplitude, t the time and τ the exponential time constant. This assumption is justified, since both the excitation of a thin scintillator by high-energy ions and the SiPM microcell avalanche develop on nanosecond timescales, whereas the observed decay constants are of the order of tens of nanoseconds to microseconds. The scintillation light signal measured by the SiPM can be interpreted as a superposition of individual SiPM microcell signals, distributed according to the temporal light output profile of the scintillator. In the continuous limit, this is described by

$$f_{\text{sig}}(t) = (f_{\text{exp},1} * f_{\text{exp},2})(t) = \int_0^t A_1 A_2 \exp\left(-\frac{t'}{\tau_1}\right) \exp\left(-\frac{t-t'}{\tau_2}\right) dt', \tag{5.5}$$

which corresponds to the convolution of two exponentially decaying functions. Solving the integral yields

$$f_{\text{sig}}(t) = A_1 A_2 \frac{\tau_1 \tau_2}{\tau_1 - \tau_2} \left(\exp\left(-\frac{t}{\tau_1}\right) - \exp\left(-\frac{t}{\tau_2}\right)\right). \tag{5.6}$$

Assuming $\tau_1 > \tau_2$, this expression can be reparametrized as

$$f_{\text{sig}}(t) = A' \left(\exp\left(-\frac{t}{\tau_1}\right) - \exp\left(-\frac{t}{\tau_2}\right)\right) \tag{5.7}$$

with $A' > 0$. As discussed in Sections 4.4 and 4.6, both SiPM and scintillator signals commonly exhibit more than a single exponential decay component. For the recorded waveforms, a single exponential description was found to be insufficient, leaving a systematic structure in the fit residuals. Introducing a second exponential component for both SiPM and scintillator was the minimum extension required to remove these deviations and to obtain a waveform parametrization suitable for calibration. Accordingly, the SiPM and scintillator signals are modeled as

$$P(t) = P_1 \exp\left(-\frac{t}{\tau_{\text{P1}}}\right) + P_2 \exp\left(-\frac{t}{\tau_{\text{P2}}}\right) \tag{5.8}$$

and

$$S(t) = S_1 \exp\left(-\frac{t}{\tau_{\text{S1}}}\right) + S_2 \exp\left(-\frac{t}{\tau_{\text{S2}}}\right) \tag{5.9}$$

respectively. This leads to the following expression for the detector response function

$$R(t) = A_1 \exp\left(-\frac{t}{\tau_{\text{S1}}}\right) + A_2 \exp\left(-\frac{t}{\tau_{\text{S2}}}\right) - A_3 \exp\left(-\frac{t}{\tau_{\text{P1}}}\right) - A_4 \exp\left(-\frac{t}{\tau_{\text{P2}}}\right) \tag{5.10}$$

where

$$R(t) = P(t) * S(t). \tag{5.11}$$

This assumes $\tau_{P1}, \tau_{P2} < \tau_{S1}, \tau_{S2}$ and $A_i > 0$, which was empirically found for the recorded waveforms. The total amount of scintillation light detected by an SiPM is then given by

$$\int_0^\infty R(t)dt = A_1\tau_{S1} + A_2\tau_{S2} - A_3\tau_{P1} - A_4\tau_{P2}. \tag{5.12}$$

Here, the relative contributions of the exponential components are assumed to be constant for varying scintillation signal amplitude. This follows from the assumption that the amplitude ratios S_1/S_2 and P_1/P_2 remain constant at varying amplitude for SiPM and scintillator signals, respectively. The validity of this assumption is verified empirically in Section 5.3.3. Under this assumption, any scintillation signal $S'(t)$ can be written as $S'(t) = BS(t)$, where B is a time independent scaling factor. The corresponding response function $R'(t)$ is then given by

$$R'(t) = P(t) * S'(t) = P(t) * BS(t) \tag{5.13}$$

and since B is time independent,

$$R'(t) = BR(t) \tag{5.14}$$

holds. This demonstrates that the relative contributions of the exponential components remain unchanged for different scintillation signal amplitudes. Consequently, the signal amplitude is proportional to the amount of detected scintillation light.

To infer the signal amplitude from the measured ToT, the amplitude normalized response function

$$\hat{R}(t) = \frac{R(t)}{R_{\max}}, \tag{5.15}$$

is used, where $R_{\max}$ denotes the maximum of the response function. The amplitude a_{T} at which the width of $\hat{R}(t)$ is equal to the measured ToT can subsequently be determined. The signal amplitude A_{s} in units of the detector threshold is then given by

$$A_{\mathrm{s}} = \frac{1}{a_{\mathrm{T}}} \tag{5.16}$$

and is thus proportional to the amount of scintillation light detected by the corresponding SiPM. In summary, the presented model provides an analytical description of the detector response and enables the reconstruction of a quantity proportional to the amount of detected scintillation light from the measured ToT.

5.3.2 Measurement setup

By default, the SiPM readout system does not provide access to the full signal waveform for any channel. However, for debugging purposes, ASIC channel 63 can be configured such that the amplified T-branch signal and $V_{\mathrm{th,T1}}$ are routed to test pads on the ASIC PCB. To record the signal waveforms, the ASIC was configured as described in [PET24] and the signal was read using an oscilloscope spring probe[2] connected to an oscilloscope[3]. Unfortunately, this readout option was not considered during the design of the detector SiPM PCB and ASIC channel 63 was connected to SiPM channel 8. This SiPM is located in the corner of quadrant D of the PCB (see Fig. 4.24) and is largely covered by one of the clamps pressing the light guide onto the SiPM

[2]Sensepeek SP200
[3]Siglent SDS5104X

PCB. Owing to the design of the fanout PCB and the SiPM PCB, it was possible to swap the cables of quadrants C and D (see Appendix G for details on pinout). As a result, ASIC channel 63 was connected to SiPM channel 40, which is still located at the edge of the light guide but not covered by a clamp. As a consequence, readout of the temperature sensor on the SiPM PCB was no longer possible. However, previous measurements showed no significant temperature fluctuations, allowing the detector to be operated reliably without temperature monitoring for the duration of this measurement. Waveform measurements were performed both at the TRIGA reactor and using the ^{241}Am α-source setup (see Section 5.1). This enabled a direct comparison between neutron induced signals and ^{241}Am α-induced signal shapes. Demonstrating similarity between these signal shapes would allow future characterization studies to be performed in the laboratory setup, eliminating the need for reactor operation and significantly reducing preparation time. For each waveform, 5000 samples were recorded over a time window of 50 µs, while the oscilloscope trigger threshold was varied. The horizontal position was chosen such that 10 µs before the rising edge were recorded. For neutron signals, oscilloscope trigger thresholds of 580 mV, 640 mV and 700 mV were used, whereas a threshold of 600 mV was applied for the ^{241}Am α-signals. Additionally, background signals were recorded at the TRIGA reactor by closing the safety shutter, which stops the UCN flux. For this measurement, an oscilloscope threshold of 640 mV was used.

5.3.3 Time-walk correction and waveform filtering

Fig. 5.9 shows a representative set of detector SiPM waveforms recorded at an oscilloscope threshold of 580 mV. To compute an average waveform, the recorded waveforms were filtered to remove noise events and signals saturating the amplifier of the readout system. First, the signal baseline was determined by averaging the first 5 µs of all recorded waveforms. Since neutron and ^{241}Am α-signals were recorded on different

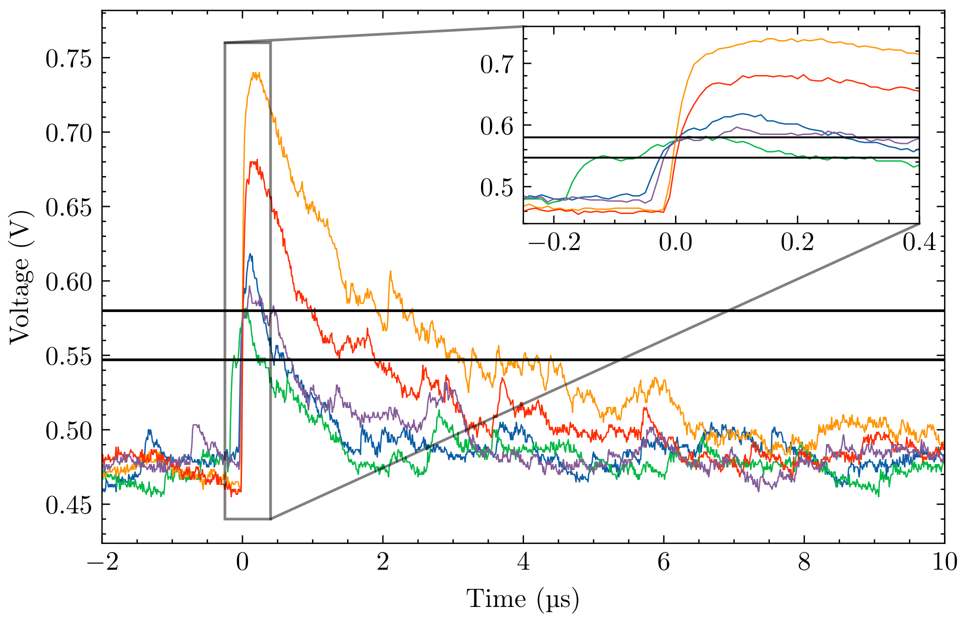

Figure 5.9: Exemplary set of detector SiPM scintillator signals recorded at the TRIGA reactor. The inset shows a magnified view of the rising edges. The bottom and top horizontal black lines indicate the detector threshold and the oscilloscope trigger threshold, respectively.

days with different ASIC calibrations, baselines were determined separately for the two datasets. The baseline was subsequently subtracted from all waveforms prior to further analysis. A second notable feature, highlighted in the inset of Fig. 5.9, is the time walk of the rising edge. This effect arises because the oscilloscope triggers at a fixed threshold and all timing information is referenced to the threshold crossing time. As a result, the apparent position of the rising edge on the time axis depends on the signal amplitude and varies between waveforms. This effect is particularly pronounced for signals with peak amplitudes only slightly above the trigger threshold. Averaging waveforms without correcting for time walk would distort sharp features such as the rising edge and peak. To mitigate this effect, the time at which each signal crossed a threshold of 500 mV, close to the baseline, was determined and used as a timing reference. Since the peak amplitude of all signals significantly exceeds this threshold, this method provides a reliable estimate of timing offset. To avoid misidentifying noise spikes as rising edges, the ToT was computed for each threshold crossing segment and the rising edge associated with the largest ToT was selected. To identify background signals, the dedicated background dataset was analyzed. Although closing the safety shutter stopped the UCN flux, the detector still recorded a small number of neutron events originating from thermal neutrons. Potential differences in signal shape between neutron and background events were therefore investigated. All waveforms were normalized to their respective amplitudes and the relative time over threshold T_R at a relative amplitude threshold of 0.35 was determined. Fig. 5.10 shows the T_R distribution for the background dataset and a dataset of neutron signals recorded using the same oscilloscope trigger threshold. A pronounced peak around 500 ns is observed in the background distribution, allowing these events to be clearly identified as noise. The remainder of the distribution, which is in good agreement with the neutron signal distribution, can be attributed to thermal neutrons. This

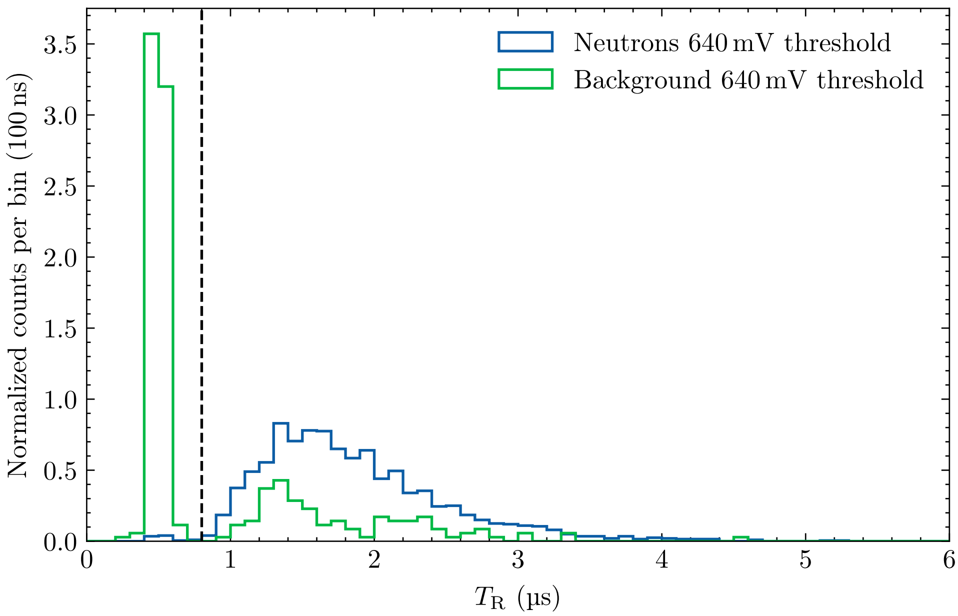

FIGURE 5.10: Relative time over threshold distributions for background and neutron events recorded at an oscilloscope trigger threshold of 640 mV. Both datasets were acquired at beam port C of the research reactor TRIGA Mainz operated in DC mode. Background data were recorded with the safety shutter closed. The vertical black line indicates $T_R = 800$ ns, which was chosen as the separation threshold between neutron and background events.

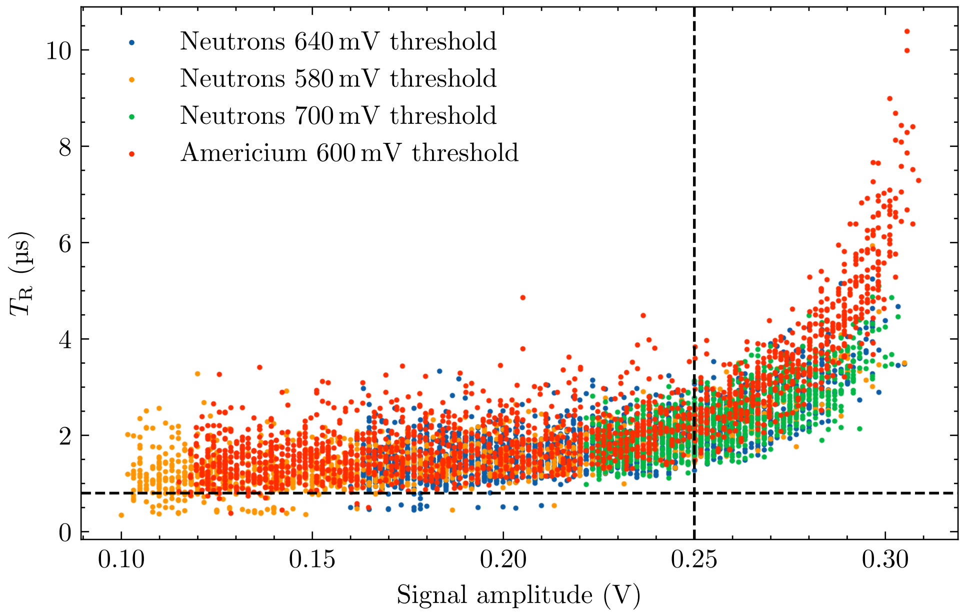

FIGURE 5.11: Relative time over threshold T_R as a function of signal amplitude for neutron and ^{241}Am α-signals recorded at different oscilloscope trigger thresholds. The datasets are shown in different colors. The horizontal and vertical black dashed lines indicate $T_R = 800$ ns and a signal amplitude of $A_w = 0.25$ V, respectively. Only waveforms located in the upper left quadrant are accepted after applying these cuts.

indicates that noise events produce significantly shorter signals, which would distort the average waveform if not filtered out. Consequently, a cut at 800 ns was applied to separate neutron events from background signals.

To identify the onset of amplifier saturation, T_R was plotted as a function of the waveform maximum A_w (see Fig. 5.11). At larger signal amplitudes, T_R increases sharply, indicating signal broadening without a corresponding increase in peak amplitude, which is characteristic of amplifier saturation. Accordingly, a cut at $A_w = 0.25$ V was applied to exclude saturating signals. Together with the T_R cut, this restricts the accepted waveforms to the upper left quadrant shown in Fig. 5.11. Below the saturation threshold, the data points exhibit only a weak dependence of T_R on the signal amplitude and show no indication of a significant systematic change in waveform shape, as would be expected if the assumption of constant exponential component amplitude ratios were violated. Within this range, the observed slope is small compared to the intrinsic scatter of the data points and therefore does not affect the validity of the calibration approach.

As a final filtering step, the time over threshold at the detector threshold T_d was computed. The detector threshold was measured to be approximately 0.547 V using the second ASIC PCB test pad and an additional oscilloscope spring probe. For each waveform segment crossing this threshold, the first data points above and below the threshold were used to determine T_d. If multiple threshold crossings occurred, the segment with the largest ToT was selected. In addition, the ToT integral, which is defined as the signal integral across the signal section above threshold, was calculated. A scatter plot of ToT integral as a function of T_d is shown in Fig. 5.12. For better visibility, the ^{241}Am α dataset was omitted and the background cut was already applied to remove outliers. Because the detector threshold used to compute T_d is identical for

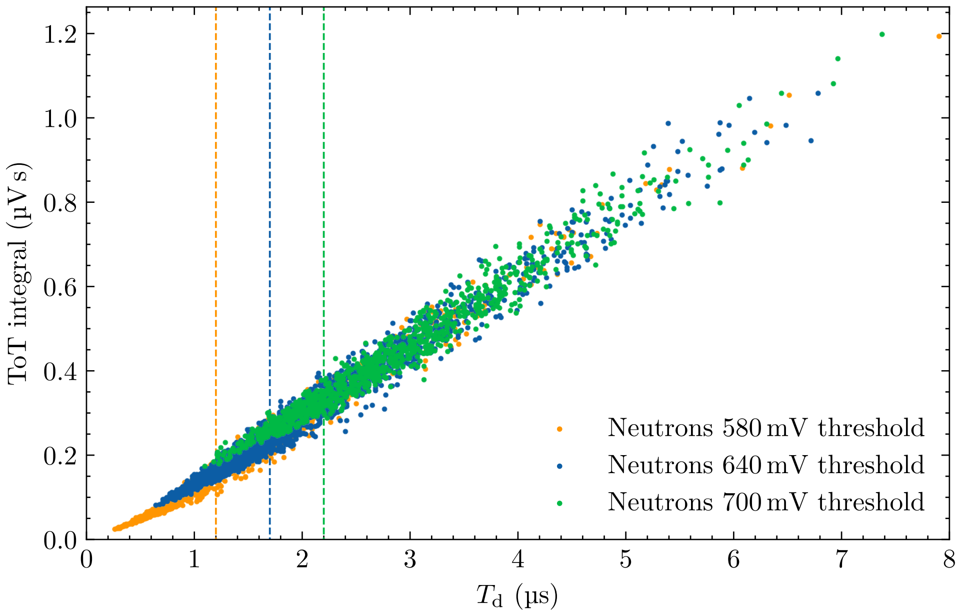

FIGURE 5.12: Scatter plot of ToT integral as a function of the time over threshold T_d for neutron signal datasets recorded at different oscilloscope thresholds. The detector threshold of 0.547 V was used to compute T_d. The dashed vertical lines indicate the threshold cuts selected for the respective dataset.

TABLE 5.1: Filter cuts applied to the SiPM scintillator waveforms recorded for neutron and ^{241}Am α-signals at different oscilloscope thresholds $V_{th,osc}$.

Signal source	$V_{th,osc}$	Background cut	Saturation cut	ToT cut
Neutron	580 mV	$T_R > 800$ ns	$A_w < 0.25$ V	$T_d > 1.2$ µs
Neutron	640 mV	$T_R > 800$ ns	$A_w < 0.25$ V	$T_d > 1.7$ µs
Neutron	700 mV	$T_R > 800$ ns	$A_w < 0.25$ V	$T_d > 2.2$ µs
^{241}Am α	600 mV	$T_R > 800$ ns	$A_w < 0.25$ V	$T_d > 1.4$ µs

all datasets, the datasets overlap for large T_d values. At smaller T_d, datasets recorded with higher oscilloscope trigger thresholds exhibit earlier cutoffs. However, the overlap between the datasets reduces closer to the cutoff of the respective datasets. This indicates a trigger induced bias for signals with amplitudes close to the oscilloscope threshold. In this regime, only sharply peaked signals with large ToT integrals are recorded. To mitigate this bias, an oscilloscope trigger threshold dependent T_d cut was introduced, as indicated by the vertical dashed lines in Fig. 5.12. Only waveforms with a T_d value exceeding these thresholds were retained. An analogous T_d cut was applied to the ^{241}Am α dataset based on the same considerations. All applied filter cuts are summarized in Tab. 5.1.

5.3.4 Waveform averaging and fit

After applying all filter cuts, the three neutron signal datasets were combined, yielding a total of 1067 waveforms used to compute the average neutron signal. For the ^{241}Am dataset, 662 waveforms remained after filtering. For waveform averaging, the time walk correction for each waveform was computed first as described above and the waveforms were shifted accordingly along the time axis. Since each waveform was shifted by a different amount, interpolation was required to evaluate all signals on a

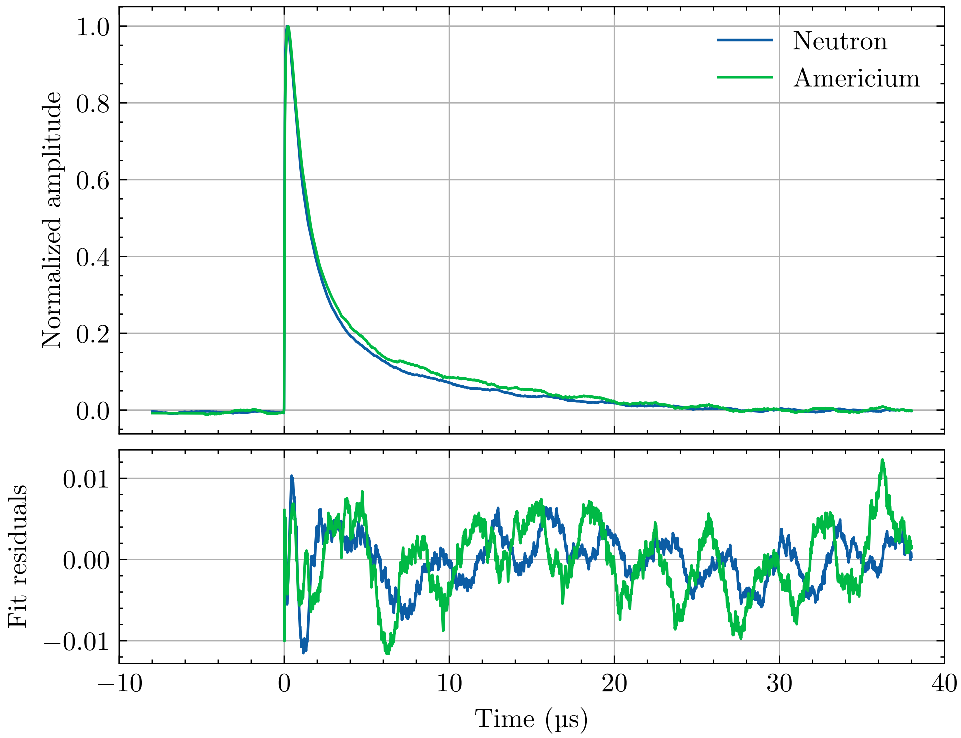

Figure 5.13: Top: normalized average waveforms for neutron and ^{241}Am α-signals. Bottom: residuals of least-squares fits to the average waveforms using the fit function defined in Eq. (5.17).

common time grid. The time walk corrected waveforms were interpolated using cubic spline interpolation[4] and evaluated on a grid of 5000 equidistant points on the time axis between −8 μs and 38 μs. Each waveform can be interpreted as the superposition of many SiPM microcell signals distributed according to the scintillator light output profile. Therefore, the average waveform was obtained by summing all interpolated waveforms. This approach implicitly assigns a larger weight to higher amplitude signals, which is appropriate because such signals correspond to a larger number of contributing SiPM microcell signals. Finally, the average waveform amplitude was normalized to its maximum amplitude. The normalized average waveforms for neutron and ^{241}Am α-signals are shown in the upper panel of Fig. 5.13. The primary goal of fitting the averaged waveforms is to extract a smooth analytical function that relates the ToT to the signal amplitude. As established in Section 5.3.1, the averaged waveform is described by Eq. (5.10) and the amount of detected scintillation light is proportional to both signal integral and amplitude. The signal amplitude A_s can be inferred from the ToT by determining the relative amplitude a_T at which the waveform width is equal to the ToT and subsequently computing A_s using Eq. (5.16). For this purpose, a least squares fit[5] to the averaged waveforms was performed starting at time $t = 0$ μs. The fit function is given by

$$F(t) = R(t - t_\mathrm{off}) + a_\mathrm{off}, \tag{5.17}$$

where R is defined by Eq. (5.10), t_off represents a temporal offset and a_off is a constant amplitude offset. In addition, the fit yields estimates for the exponential time

[4]Python package: scipy.interpolate.CubicSpline (v. 1.15.2)
[5]Python package: scipy.optimize.curve_fit (v. 1.15.2)

Table 5.2: Fit parameters obtained from least squares fits to the averaged neutron and ^{241}Am α waveforms using the fit function defined in Eq. (5.17). The quoted parameter spreads were estimated using a bootstrap procedure and should be interpreted as robustness estimates rather than rigorously derived statistical uncertainties.

Fit parameter	Neutron signal	^{241}Am α-signal
A_1	0.97(1)	0.99(3)
A_2	0.290(4)	0.298(5)
A_3	0.8(1)	0.9(4)
A_4	20(9)	15(4)
τ_{S1}	1.13(2) µs	1.17(2) µs
τ_{S2}	7.5(1) µs	8.9(2) µs
τ_{P1}	94(3) ns	106(5) ns
τ_{P2}	18.6(3) ns	20.4(5) ns
t_{off}	−40(5) ns	−40(10) ns
a_{off}	−0.0043(4)	−0.0078(7)

constants characterizing the signals. However, since the primary objective was to extract a smooth function rather than parameter inference, no detailed statistical uncertainty model for the waveform data points was developed. Instead, equal weights were assigned to all data points for the fit. The residuals, shown in the lower panel of Fig. 5.13, exhibit an oscillatory structure. These oscillations likely originate from small oscillatory features present in the individual waveforms that are not fully suppressed by averaging. Traces of such oscillations are also visible in the tail region of the averaged waveforms. Because the oscillations appear to involve multiple frequency components, no attempt was made to incorporate them into the fit model. To assess the stability of the fit parameters, a bootstrap approach was employed. A total of 1000 artificial datasets of size $N_\text{d} = 5000$, equal to the original dataset size, were generated by sampling with replacement from the averaged waveform. The fit procedure was repeated for each artificial dataset and the sample standard deviation of the resulting parameter distribution was used as a stability estimate. Since the rising edge region contains only about 10 data points, artificial datasets containing fewer than five data points within the rising edge region, defined as amplitudes between 0.05 and 0.90, were rejected. This rejection was necessary, because insufficient sampling of the rising edge led to non converging fits or large outliers in parameters associated with the fast SiPM time constant, which is particularly sensitive to the rising edge shape. The described procedure does not yield statistical confidence intervals in a strict sense. Rather, the resulting parameter spreads should be interpreted as data driven estimates of the robustness and stability of the fit parameters. The final fit results and stability estimates are summarized in Tab. 5.2.

5.3.5 Extraction of calibration functions

The fitted average waveforms were subsequently used to extract a calibration function. To achieve this, a functional relationship between signal width (ToT) T_w and relative amplitude a_T was established. The signal width at relative amplitude a_T is defined as

$$T_\text{w} = t_2(a_\text{T}) - t_1(a_\text{T}), \tag{5.18}$$

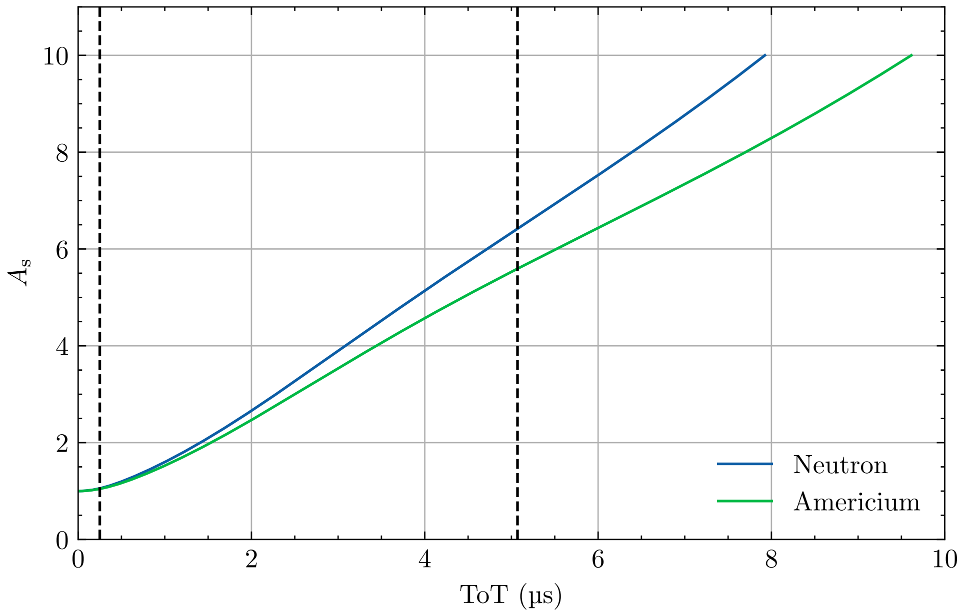

FIGURE 5.14: Relation between signal amplitude, expressed in units of the detector threshold, and ToT for neutron and ^{241}Am α-signals. The dashed vertical lines mark ToTs of 250 ns and 5070 ns, corresponding to the accepted ToT range in post processing of the detector readout system output.

where t_1 and t_2 denote the times at which the normalized response function $\hat{R}(t)$ crosses a_T on the rising and falling edges, respectively. Here, $t_1(a_T)$ corresponds to the inverse of $\hat{R}(t)$ for $t < t_{max}$ and $t_2(a_T)$ to the inverse of $\hat{R}(t)$ for $t > t_{max}$, where t_{max} denotes the time at which $\hat{R}$ reaches its maximum. In a first step, t_{max} and R_{max} were determined numerically using a minimizing algorithm[6]. Next, the normalized response function $\hat{R}(t) = R_{fit}(t)/R_{max}$ was computed. The fit function was subsequently sampled on an equidistant grid of 5000 time points t_i between 0 µs and 20 µs, yielding the corresponding relative amplitudes $a_{T,i}$. Based on t_{max}, the samples were split into two subsets with $t_i < t_{max}$ and $t_i > t_{max}$. The functions $t_1(a_T)$ and $t_2(a_T)$ were then obtained by cubic spline interpolation[7] of the respective subsets. Using these interpolations, 5000 pairs $(a_{T,j}, T_{w,j})$ were computed for equidistant values of $a_{T,j} \in [0.1, 1]$. This dataset was interpolated again using cubic splines to obtain the inverse relation $a_T(T_w)$. The signal amplitude in units of the detector threshold is then given by

$$A_s(T_w) = \frac{1}{a_T(T_w)}. \tag{5.19}$$

The resulting calibration functions are shown in Fig. 5.14. Here, a ToT of 0 ns corresponds to an amplitude $A_s = 1$, i.e., a signal exactly at the detector threshold.

5.3.6 Discussion

Overall, the waveform fits show that the relative weights of the exponential components are in good agreement between neutron and ^{241}Am α-signals. In contrast, differences are observed in the extracted time constants, most notably in the slow scintillator component. According to the SiPM datasheet, the microcell recharge constant is 95 ns [Sem22], which is in good agreement with the fitted value of τ_{P1} for

[6]Python package: scipy.optimize.minimize (v. 1.15.2)
[7]Python package: scipy.interpolate.CubicSpline (v. 1.15.2)

neutron signals. Given that the bootstrap based stability estimation provides only an order of magnitude spread estimate rather than a rigorously derived statistical uncertainty, the corresponding value obtained for the ^{241}Am α-signals is also considered consistent. Differences in τ_{P1}, τ_{P2} and τ_{S1} were found to have a negligible impact on the resulting calibration functions and were therefore not investigated further. The pronounced difference in the slow scintillator time constant τ_{S2} was anticipated, as the tail of the averaged ^{241}Am α-signal is visibly longer than that of the averaged neutron signal. However, the underlying physical mechanism responsible for this difference is not yet fully understood. One key difference in signal generation is the substantially higher kinetic energy of the ^{241}Am α-particles (approximately 5.48 MeV [Bas06]) compared to the energies of the ions produced in the ^{10}B neutron capture reaction (see Section 4.4). As a result, the range of α-particles emitted by ^{241}Am in the ZnS(Ag) scintillator of $d_{\mathrm{Am}} = 21\,\mu\mathrm{m}$, which was obtained from SRIM simulations [ZZB10], significantly exceeds the ranges of the neutron capture reaction products (see Tab. 4.2). Consequently, ^{241}Am α-particles are not fully stopped in the scintillation layer, which has an average thickness of 8 µm (see Section 4.4), at all incidence angles. Therefore, scintillation light generated by neutron capture products must, on average, traverse a longer path within the scintillator, increasing the probability of absorption and scattering. This may preferentially suppress wavelength components associated with the slow scintillation tail. If the emission spectrum components responsible for the slow decay overlap more strongly with the scintillator absorption spectrum, enhanced absorption for neutron signals would be expected. Fig. 5.11 is consistent with this interpretation, as the relative ToT T_{R} of ^{241}Am α-signals increases earlier near the amplifier saturation threshold than for neutron signals. Part of this effect, however, may also originate from small voltage offsets between the two measurements due to different ASIC calibrations. In summary, the precise origin of the observed differences in average signal shape remains unclear. Further investigations, such as repeating the measurements with varying scintillator thickness or performing time resolved measurements of the scintillator emission spectrum, could enable a test of the proposed hypothesis. Both approaches would require significant modifications of the measurement setup and were therefore not pursued within the scope of this work. To assess the practical relevance of the observed differences, the average discrepancy between reconstructed event positions obtained using the two calibration functions must be compared to the spatial resolution of the detector. This comparison is presented in Section 5.4.7. Ultimately, the differences between reconstructed positions obtained using the neutron based and ^{241}Am α based calibration functions were found to be negligible compared to the spatial resolution of the detector prototype. Accordingly, the calibration derived from neutron signals was used for the position reconstruction discussed in the subsequent sections.

5.4 Spatial resolution

5.4.1 Position reconstruction procedure

In order to assess which procedures are suitable for reconstructing the position of neutron events, a neutron signal dataset was analyzed. The dataset contains a total of 561 814 events, defined according to the event definition summarized in Section 5.2.4. Fig. 5.15 shows the distribution of active channels per event for this dataset. To

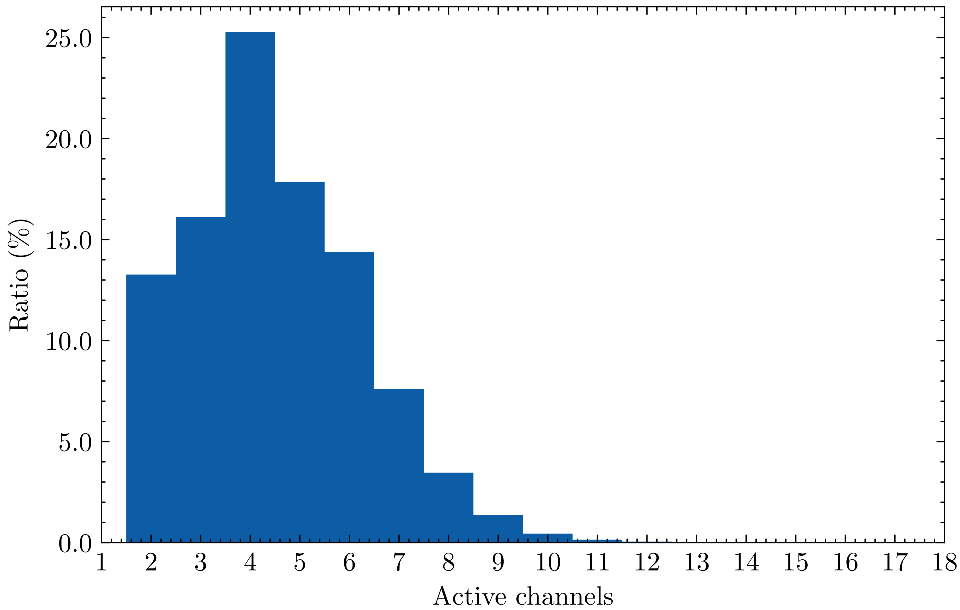

FIGURE 5.15: Distribution of the number of active channels per neutron event. Events were defined according to the criteria established in Section 5.2.4.

reconstruct an event position, the event position vector

$$\vec{e} = \begin{pmatrix} x_\mathrm{e} \\ y_\mathrm{e} \end{pmatrix}, \tag{5.20}$$

must be determined. Here, x_e and y_e denote the event position on the detector surface plane. The detector event coordinate system is defined such that its origin coincides with the center of the light guide. The orientation of x- and y-axes is defined as shown in Fig. 4.24, with both axes aligned to the SiPM grid. Consequently, SiPM channels 1 and 64 are in the upper left and lower right corners of the detector, respectively. All following figures show the detector surface from a perspective viewing the detector surface from the front. Every position reconstruction procedure must determine at least two free parameters. As a result, only events with three or more active channels allow for a meaningful position reconstruction. Many reconstruction methods, such as fitting analytical light intensity distributions to individual events, introduce additional free parameters and therefore require an even larger number of active channels. In addition to reconstruction accuracy, the required computation time must be considered. UCN experiments typically run continuously for several days and produce a large number of events in the process. Therefore, the reconstruction method must be capable of live or quasi live processing. At beam port C of the research reactor TRIGA Mainz, the event rate is approximately $500\,\mathrm{s}^{-1}$, corresponding to an upper limit of about $2\,\mathrm{ms}$ per event for the reconstruction time. Consequently, computationally expensive approaches, such as numerical modeling of the detection process for each event, are not feasible. Based on these considerations, three different reconstruction approaches were implemented and evaluated:

1. a weighted average approach,
2. a least-squares fit assuming an isotropic, point-like light source and
3. a least-squares fit assuming a two dimensional Gaussian light intensity distribution.

In the weighted average approach, the amount of light detected by each channel is assigned to the center of the corresponding SiPM. The reconstructed event position is then obtained by computing the weighted average of the SiPM positions, which is given by

$$\vec{e} = \frac{\sum_{i=1}^{64} w_i \vec{p}_i}{\sum_{i=1}^{64} w_i}, \tag{5.21}$$

where w_i is a weight proportional to the amount of detected light and $\vec{p}_i$ denotes the center position of the respective SiPM. The calibration functions $A_\mathrm{s}(T_\mathrm{w})$ for neutron and ^{241}Am α-signals established in Section 5.3.3 were used to estimate the weights w_i for each event. These calibration functions provide a value proportional to the absolute amount of detected light. However, due to the trigger threshold, the measured light distribution is effectively truncated. That is, channels located farther from the event position do not contribute to the average if their signal amplitude is below threshold. This leads to a systematic bias towards regions of higher light intensity. To compensate for this effect, the signal amplitude corresponding to the trigger threshold must be subtracted. Since the calibration function returns the amplitudes in units of the threshold level, the amplitude at trigger threshold is given by $A_\mathrm{s}(0\,\mathrm{ns}) = 1$. Additionally, the event identification algorithm only accepts channels with a ToT greater than 250 ns, which must be corrected for in the same way. The effective calibration function therefore becomes

$$A'_\mathrm{s}(T_\mathrm{w}) = A_\mathrm{s}(T_\mathrm{w}) - A_\mathrm{s}(250\,\mathrm{ns}). \tag{5.22}$$

The second reconstruction approach models each event as an isotropic, point like light source located at the event position. This is the same model assumed in the light guide simulations discussed in Section 4.5.1. The corresponding light intensity distribution is given by

$$I_\mathrm{ps}(x,y) = \frac{I_0}{\left((x-x_0)^2 + (y-y_0)^2 + z_0^2\right)^{3/2}}, \tag{5.23}$$

where I_0 denotes the total intensity, and x_0, y_0 and z_0 represent the event position in space (see Appendix E.1 for a derivation). In this model, the intensity observed by a given SiPM is obtained by integrating the intensity distribution over the channel area, which is given by

$$I_\mathrm{obs,ps}(x_1,x_2,y_1,y_2) = \int_{x_1}^{x_2}\int_{y_1}^{y_2} I_\mathrm{ps}(x,y)\,dy\,dx, \tag{5.24}$$

where x_1, x_2, y_1 and y_2 denote the SiPM active area edges. Evaluating the integral yields

$$I_\mathrm{obs,ps}(x_1,x_2,y_1,y_2) = \sum_{i,j}(-1)^{i+j} H(x_i,y_j), \tag{5.25}$$

with $i,j \in \{1,2\}$, where

$$H(x,y) = \frac{I_0}{2\pi}\arctan\left(\frac{(x-x_0)(y-y_0)}{z_0\sqrt{(x-x_0)^2+(y-y_0)^2+z_0^2}}\right) \tag{5.26}$$

is the corresponding indefinite integral of I_ps (see Appendix E.1 for a derivation). The indefinite integral was derived with the help of an online integral solver[8] based

[8] `https://www.integral-calculator.com/`

on an open source computer algebra system[9]. By performing a least squares fit of this model to the measured light distribution, the event position can be inferred. Since z_0 is known from the geometry of the light guide, it can be treated as a known parameter, reducing the number of free fit parameters to three.

Although an isotropic point source is a reasonable model for a transparent scintillator with a refractive index similar to that of the light guide, this assumption is not strictly valid for the ZnS(Ag) scintillation layer, which exhibits an opaque, polycrystalline structure as explained in Section 4.4.2. This structure introduces additional losses due to scattering and absorption, which depend on the photon path length within the scintillator. As a consequence, photons emitted at larger angles relative to the detector surface normal experience increased attenuation, leading to suppression of the intensity distribution tails. Accurately modeling these effects would require detailed knowledge of the optical properties of the scintillation layer, which was not available and is, in practice, difficult to obtain. Moreover, the resulting distribution is unlikely to be well described by a simple analytic function, which would substantially complicate any fitting procedure. In addition, the photon path length in the scintillator depends on the incidence angle of the ions produced in the ^{10}B neutron capture reaction, implying that both the amplitude and width of the intensity distribution may vary from event to event. Owing to the rough polycrystalline structure of the scintillator, even a position dependent asymmetric intensity distribution cannot be excluded. Since the true intensity distribution is expected to exhibit less pronounced tails than the point source model, a Gaussian intensity model was investigated as an approximation. In this approach, the light intensity distribution is given by

$$I_{\mathrm{G}}(x, y) = A \exp\left(-\frac{(x-x_0)^2 + (y-y_0)^2}{2\sigma^2}\right), \tag{5.27}$$

where A denotes the amplitude, x_0 and y_0 the event position and σ the width of the distribution. Due to the symmetry of the detector geometry, the distribution is assumed to be symmetric in x and y. The observed signal in each SiPM is obtained by integrating the distribution over the active area of the SiPM (see Appendix E.2 for a derivation), yielding

$$\begin{aligned} I_{\mathrm{obs,G}}(x_1, x_2, y_1, y_2) = \frac{A}{4} &\left(\operatorname{erf}\left(\frac{x_2 - x_0}{\sqrt{2}\sigma}\right) - \operatorname{erf}\left(\frac{x_1 - x_0}{\sqrt{2}\sigma}\right)\right) \\ &\left(\operatorname{erf}\left(\frac{y_2 - y_0}{\sqrt{2}\sigma}\right) - \operatorname{erf}\left(\frac{y_1 - y_0}{\sqrt{2}\sigma}\right)\right), \end{aligned} \tag{5.28}$$

where x_1, x_2, y_1 and y_2 denote the dimensions of the respective SiPM. By performing a least squares fit, the event position can be reconstructed. The fit function contains four free parameters. To reduce this number, the width σ was fixed to the most probable value obtained from fits to events with a large number of active channels. For both intensity distribution based approaches, the measured light amplitudes are obtained from the calibration functions established in Section 5.3.3. In contrast to the weighted average method, no offset correction is required, as truncation effects are intrinsically accounted for in the fit models.

[9] `https://maxima.sourceforge.io/`

5.4.2 Observed spatial distributions of events

To compare the different reconstruction approaches, the positions of a dataset of 487 261 neutron events recorded at the test measurement setup at the research reactor TRIGA Mainz were reconstructed. This includes only events exhibiting at least three active channels. The reconstructed event positions obtained for the three tested reconstruction procedures are shown in Fig. 5.16. In order to estimate the expected spatial distribution of the events, the UCN energy spectrum and the beam guide shape and material must be considered. The beam guides have an inner diameter of 66 mm, which corresponds to a gravitational potential difference of about 7 neV for UCNs. This is small compared to the UCN spectral energy range, which is limited by the Fermi potential ($V_\mathrm{F} = 184\,\mathrm{neV}$, see Appendix A) of the stainless steel guides. Consequently, an approximately uniform distribution of events across the neutron guide cross section is expected.

While the reconstructed event positions largely fall within the neutron guide cross section, the distributions are clearly non-uniform for all tested reconstruction procedures. Instead, the reconstructed positions exhibit clusters whose locations and shapes vary between the different methods. Additionally, a small number of events are reconstructed to positions outside the neutron guide cross section in all cases. For the weighted average approach, a sharp edge is visible in the reconstructed position distribution. The presence of a small number of events outside the neutron guide cross section can be explained by the approximately 5 mm wide gap between the end of the neutron guide and the detector surface. Due to this gap UCNs can be reflected from the guide walls at large angles and reach positions on the detector surface outside the geometric projection of the guide cross section. Additionally, as described in Section 5.1, a steady flux of thermal neutrons is present in the experimental area around the TRIGA reactor. Although the detection efficiency for thermal neutrons is significantly smaller than for UCNs (see Section 4.4), a small fraction of the thermal neutrons traversing the detector is still detected. Since thermal neutrons can penetrate the detector vacuum chamber at essentially any position, they are expected to produce a spatially uniform background across the detector surface. The sharp edge observed for the weighted average approach arises from an intrinsic bias introduced by the reconstruction method itself for events close to the edge of the light guide. Because no SiPM channels exist beyond the light guide boundary, the observed light intensity distribution is effectively truncated. As a result, there are no channels available to balance the weighting near the edge, which systematically shifts the reconstructed position towards the interior of the light guide. Since the observed light intensity is assigned to the center of the respective SiPM in the weighted average approach, no event can be reconstructed to a position outside the square defined by the centers of the outermost SiPMs. This geometric constraint produces the observed sharp edge in the reconstructed position distribution. In contrast, this limitation is not present in reconstruction approaches based on fitting an intensity distribution. In these methods, the fit naturally accounts for the absence of information outside the light guide. However, for events very close to the light guide edge, reflections from the boundaries are expected to significantly increase the observed intensity of the outermost SiPM channels. Since such effects are not included in the fit models, the fit based reconstruction approaches are also expected to exhibit biases in these regions. However, the beam guide cross section predominantly covers the central region of the light guide. Therefore, the number of events reconstructed close to the edge is expected to be small. Moreover, a future detector to be installed in the τSPECT experiment will

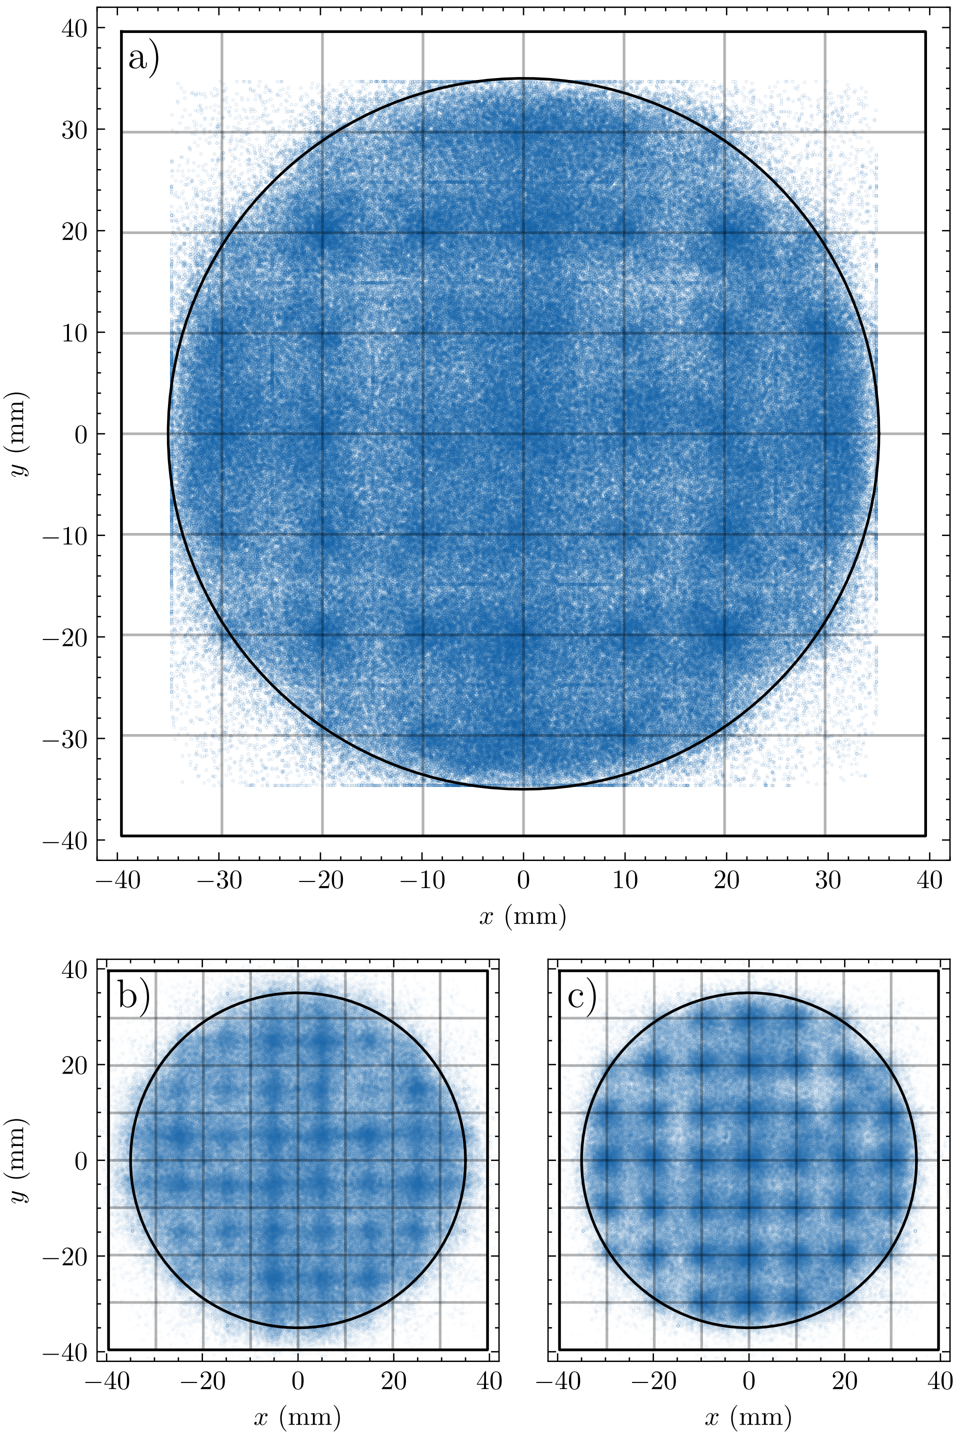

FIGURE 5.16: Reconstructed event positions of 487 261 neutron events exhibiting at least three active channels. The reconstruction is based on signal amplitudes estimated via the average neutron SiPM signal waveform fit. Each event is represented by a dot. The black square and circle mark the light guide and the beam guide edge, respectively, while the grid represents the light guide grid. The centers of the SiPMs coincide with the center of the grid squares. a) Weighted average approach, b) Gaussian least squares fit, c) isotropic point like light source intensity distribution least squares fit.

have a different geometry, which is likely to significantly modify edge related effects. For these reasons, the following sections focus primarily on the position reconstruction performance in the central region of the detector. A potential strategy to mitigate edge effects is discussed in Section 5.4.8.

5.4.3 Reconstruction bias

The most striking feature of the reconstructed event position distributions is the presence of clusters. In principle, such structures can originate from the actual UCN distribution, defects on the detector surface, variations in channel gain, the reconstruction procedure itself or a combination of these effects. However, the observation that the clusters exhibit a periodic structure aligned with the symmetries of the light guide and that their location changes between reconstruction procedures, strongly suggests that they are caused by the reconstruction process. To test this hypothesis, dedicated model calculations were performed. For this purpose, a light intensity distribution was assumed and analytically integrated across the area occupied by each SiPM to obtain the observed intensity for each detector channel. This corresponds to idealized conditions in which no refraction occurs and all SiPM channels exhibit a detection efficiency of 100 %. The observed intensity pattern was computed on an equidistant grid of 201 × 201 positions (x_i, y_i) covering the 16 innermost detector channels ($x_i, y_i \in [-19.8\,\text{mm}, 19.8\,\text{mm}]$). Subsequently, a threshold was introduced and all observed intensities below this threshold were set to zero. Event positions were then reconstructed using one of the reconstruction procedures introduced above. The model calculations were repeated for different light intensity distributions, distribution widths and thresholds. Only the central region of the detector was considered, as edge effects are expected to be negligible in this area. Furthermore, the model does not include any statistical noise, since all observed light intensities are computed purely analytically. As a consequence, any deviations from a uniform distribution in the reconstructed position patterns must originate from biases inherent to the reconstruction process.

For most model calculations, a Gaussian light intensity distribution was assumed and the weighted average approach used for reconstruction. In addition to setting observed intensities below threshold to zero, the threshold was subtracted from the remaining active channels for this approach, as described above. Under these conditions, a wide variety of clustering patterns in the reconstructed position distributions was found. As expected, the degree of clustering increased as the number of active channels decreased. Typically, two characteristic types of clustering patterns, referred to as types (a) and (b), emerged. Representative examples are shown in Fig. 5.17 (a) and (b), while subplot (d) shows the true positions of the probed events. The main difference between the pattern types lies in the location of the dominant clusters relative to the SiPM grid. For pattern type (a), the clusters are centered around the SiPM centers, whereas they are centered at the midpoints between diagonal neighboring SiPMs for pattern type (b). Pattern type (a) typically occurs for high thresholds, where the number of active channels drops below three for events at certain positions. As a result, the local density of reconstructed events is also influenced by the fact that events originating from these positions cannot be reconstructed and are therefore absent from the reconstructed position distribution. The true positions of such events for the exemplary type (a) pattern are shown in Fig. 5.17 (c). In contrast, pattern type (b) is typically observed for intermediate thresholds, where the number of active channels decreases significantly but does not fall below three at any position.

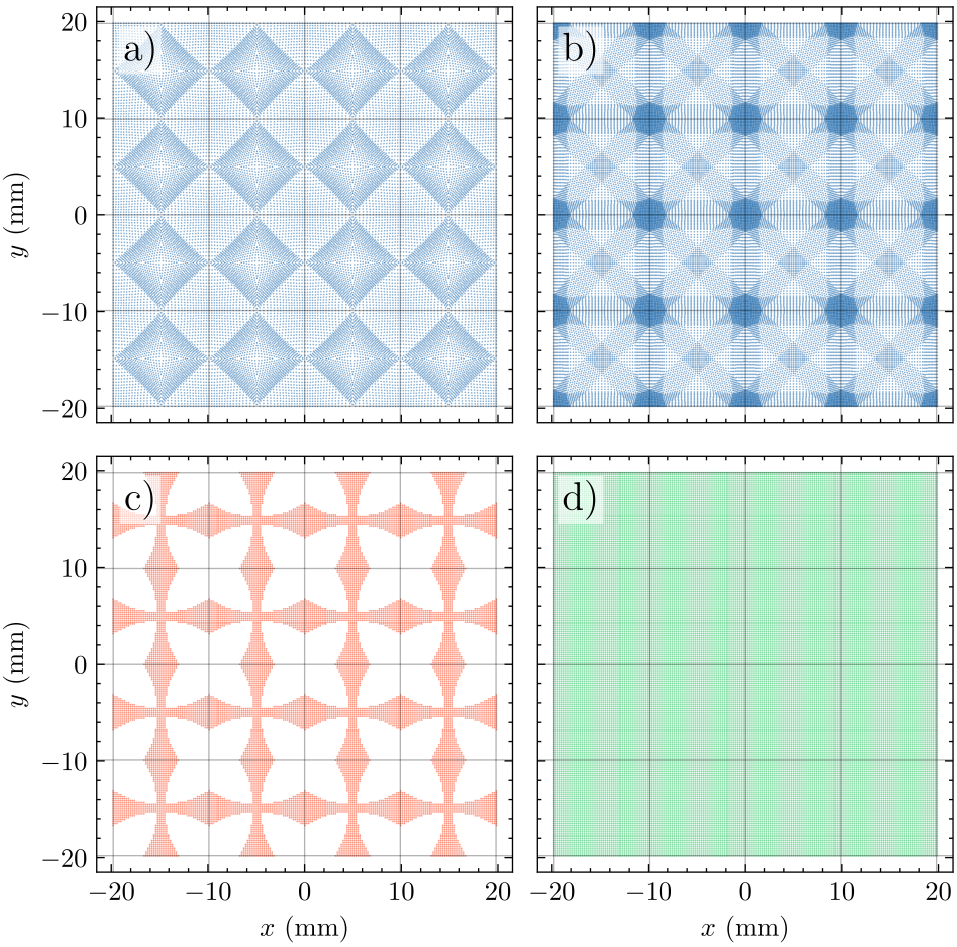

FIGURE 5.17: Representative reconstructed event position distributions generated using the weighted average approach and assuming a two dimensional Gaussian signal light intensity distribution with a width of $\sigma = 10\,\text{mm}$. The (reconstructed) position of each event is represented by a dot. The centers of the SiPMs are located at the center of the depicted grid squares. a) Reconstructed event position distribution for a threshold of 3.6 %. b) Reconstructed event position distribution for a threshold of 2.5 %. c) True positions of events for which the number of active channels is smaller than three for a threshold of 3.6 %. d) True position of all events.

While the exact shape of the patterns changes slightly with increasing width of the Gaussian light intensity model, the overall structure remains similar. The same holds when replacing the Gaussian distribution with a point like light source model.

Significant clustering and shifts of cluster centers are also observed when applying the point source fit approach to Gaussian model data and vice versa. However, the resulting patterns generally differ from those obtained with the weighted average approach. These distributions, as well as additional modeled reconstructed position distributions obtained for the weighted average approach with different thresholds are presented in Appendix E.4. Since the obtained patterns are relatively stable with respect to variations in the assumed light distribution type and width, it is reasonable to expect that similar bias patterns also occur in real data. In contrast to the model, however, the signal amplitude in real data follows an unknown distribution. This situation is equivalent to varying the detector threshold at a fixed amplitude. Consequently, the clustering pattern observed in the measured data can be interpreted as a superposition of bias patterns corresponding to different thresholds, weighted by the signal amplitude distribution. Based on the model calculations, it is therefore very likely that the periodic structures observed in the reconstructed event position distributions of the measured data originate from biases in the reconstruction procedures. Differences between the distributions obtained with different reconstruction approaches can plausibly be attributed to differences in the corresponding bias patterns. Moreover, the model calculations show that the reconstruction threshold at which fewer than three channels are active is position dependent. Events in these regions cannot be reconstructed in a meaningful way and are therefore omitted from the reconstructed event position distributions, leading to locally reduced densities. In general, the reconstruction bias is most severe for events with a small number of active channels. In such cases, significant bias can only be avoided if a fit model closely matching the actual light intensity distribution is employed. At present, however, the actual light intensity distribution is not known and as explained above may even vary between events. Additionally, fit approaches would benefit from a larger number of active channels, which would allow for a more robust determination of the fit parameters. A substantial fraction of detector events exhibits fewer than three active channels and therefore cannot be reconstructed in a meaningful way. That is, increasing the number of active channels would address multiple issues simultaneously. Possible detector design strategies to achieve this are discussed in Section 5.4.8.

5.4.4 Non-periodic spatial non-uniformities

In addition to the periodic clusters, the reconstructed event position distributions also exhibit non-periodic non-uniformities. One example is the region around $x = 15\,\mathrm{mm}$, $y = 15\,\mathrm{mm}$, where a reduced local event density is observed. Such features can plausibly be explained by variations in channel gain or optical coupling between neighboring SiPM channels. Variations in the optical coupling could e.g. arise from air bubbles in the optical grease. However, after mounting the light guide on the SiPM grid, only the optical coupling quality of the outermost SiPMs cannot be inspected visually. Therefore, this hypothesis has not been tested so far.

5.4.5 Performance of reconstruction approaches

In order to assess which reconstruction approach performs best, the uniformity of the reconstructed event position distributions obtained with the different methods was evaluated. While the true spatial distribution of events on the detector surface is not known a priori, the UCN phase space considerations discussed above suggest an approximately uniform distribution in the central detector region. A uniform distribution was therefore adopted as a benchmark for comparing the reconstruction approaches. For each approach, a symmetric two dimensional histogram with a bin width of 1 mm was constructed and compared to the bin counts expected for a perfectly uniform event density. The bin width was chosen to resolve spatial structures on the scale of the observed clustering while maintaining sufficient statistics per bin. For this comparison, all reconstructed events satisfying $x_\mathrm{e}, y_\mathrm{e} \in [-20\,\mathrm{mm}, 20\,\mathrm{mm}]$ were considered for each reconstruction approach. Within this region, the edge effects discussed above are expected to be negligible. Subsequently, the bin counts of each histogram were compared to the average bin count

$$\bar{N}_\mathrm{b} = \frac{N_\mathrm{tot}}{n_\mathrm{bin}}, \tag{5.29}$$

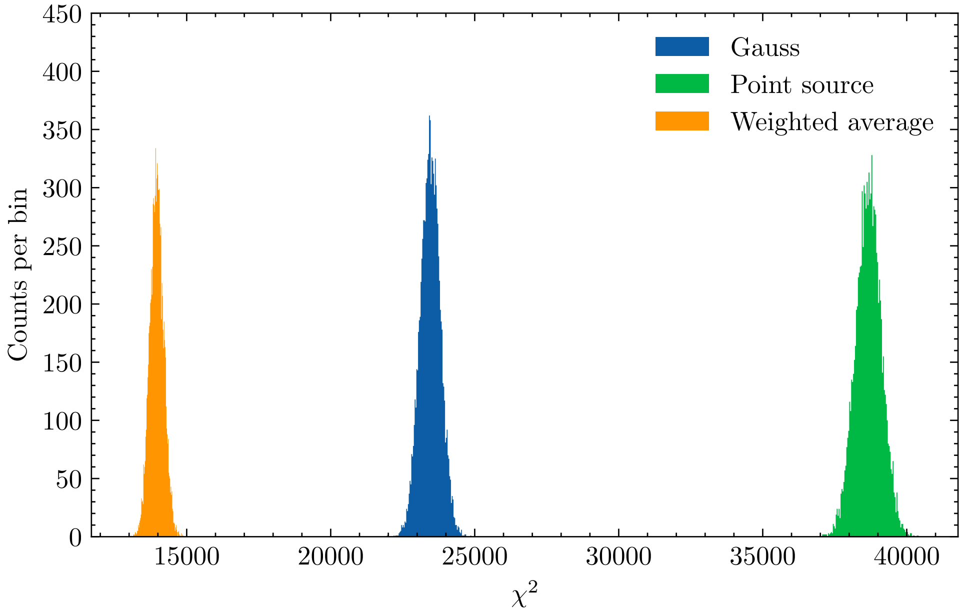

FIGURE 5.18: Distribution of the χ^2-metric used to estimate how well the different position reconstruction approaches recover the expected uniform distribution. The χ^2-distributions were generated by resampling the histogram bins derived from the reconstructed event positions, assuming Poisson distributed bin counts.

where N_{tot} denotes the total number of reconstructed events in the selected region and n_{bin} the number of bins. To quantify agreement between reconstructed distributions and a uniform distribution, a χ^2 goodness of fit test was performed for each reconstruction approach. During the measurements, each bin is populated at a constant average rate. Therefore, the number of counts per bin is assumed to be Poisson distributed. The corresponding test statistic is then given by

$$\chi^2 = \sum_i \frac{(N_i - \bar{N}_{\text{b}})^2}{\bar{N}_{\text{b}}}, \tag{5.30}$$

where N_i denotes the number of counts in the respective bin. To account for the statistical uncertainty of the individual bin counts, each histogram was resampled 10 000 times. For this procedure, the counts in each bin were assumed to be Poisson distributed with mean N_i. For each resampled histogram a χ^2 value was computed according to the definition above. The resulting distributions are shown in Fig. 5.18 and clearly indicate that the weighted average approach yields the event position distribution that is closest to uniform. Therefore, neglecting edge effects and within the central detector region, the weighted average reconstruction approach yielded the best performance and was used for all subsequent analyses.

5.4.6 Spatial resolution estimation

In order to assess the spatial resolution of the detector prototype, neutron events at well defined position must be generated. However, a narrow, focused UCN beam that can be directed onto arbitrary positions on the detector surface is technically challenging and not readily available at UCN facilities. Therefore, thin stainless steel masks with precisely defined cutouts were used instead. These masks block a large fraction of the UCN flux across the detector surface due to their comparatively high Fermi

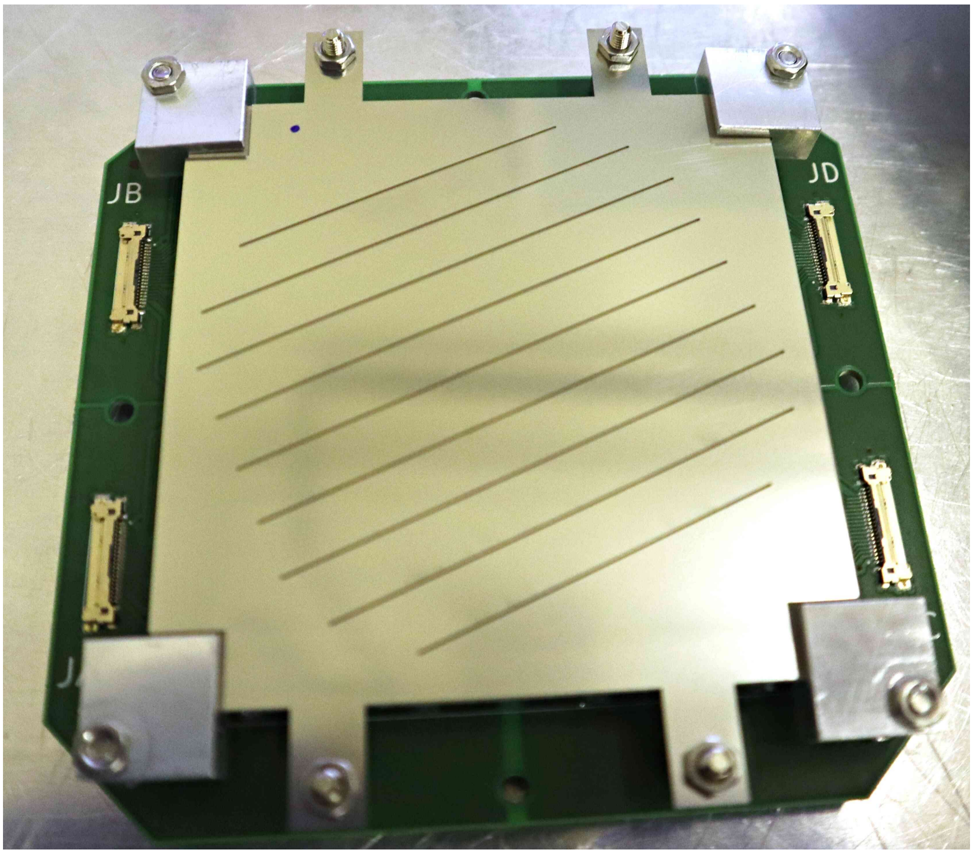

FIGURE 5.19: Photograph of the detector with a 0.1 mm thick stainless steel mask mounted on top of the detector surface. The cutouts have a diameter of 0.3 mm, a center to center spacing of 7.0 mm and are oriented at an angle of 22.5° with respect to the detector x-axis, which aligns with the bottom edge of the light guide in this image.

potential ($V_{\mathrm{F}} = 184\,\mathrm{neV}$, see Appendix A), while the flux remains at its nominal value at the cutout locations. In this way, the cutouts can be imaged with the detector and the reconstructed cutout shapes can subsequently be used to estimate the resolution of the detector.

Fig. 5.19 shows a photograph of the mask that was eventually used for the spatial resolution estimation. The mask is 0.1 mm thick and contains nine parallel line-shaped cutouts with a width of 0.3 mm and a center to center spacing of 7.0 mm. The cutout lines exhibit an angle of 22.5° with respect to the x-axis of the detector coordinate system and their length was chosen such that the lines cover the projection of the UCN guide cross section. Using a total of nine parallel lines also allows to identify potential local variations in spatial resolution across the detector. Line-shaped cutouts were chosen because deviations from linearity in the reconstructed shapes can be identified relatively easily. The spacing and orientation of the cutouts were selected such that they do not coincide with the symmetry of the light guide grid and that the reconstructed line shapes are not expected to overlap significantly. To accommodate the additional space required by the spacers that were used to mount the mask, the distance between the detector surface and the detector flange had to be increased by approximately 2 mm to 3 mm. For the measurement, the detector was installed at beam port C of the research reactor TRIGA Mainz with the mounted mask. A total of 417 316 events, defined according to the event definition summarized in Section 5.2.4, were recorded. The distribution of the reconstructed event positions is shown in Fig. 5.20. The reconstruction weights were obtained from the neutron signal calibration and only events with three or more active channels were considered for the analysis. In general, all nine cutouts are visible in the reconstructed event position distribution. However, a substantial background count rate is also present, which partially obscures the cutout shapes. As in the measurement without a mask, clusters associated with the light guide and SiPM grid symmetry are visible. In a few regions like e.g. around $x = 15\,\mathrm{mm}$ and $y = 15\,\mathrm{mm}$ deviations from an ideal straight line shape can be observed. These deviations typically occur in the same areas where deviations from the periodic pattern were already seen in measurements without masks. They are therefore likely caused by the same underlying effects, such as imperfections in the optical coupling between the SiPMs and the light guide, as discussed previously in Section 5.4.4. In practice, clearly identifying deviations from the ideal cutout shapes is challenging, since both event clustering at the light guide grid intersections and reconstruction induced edge biases are present simultaneously. These effects can either distort the reconstructed line shapes or wash out finer features. To estimate the

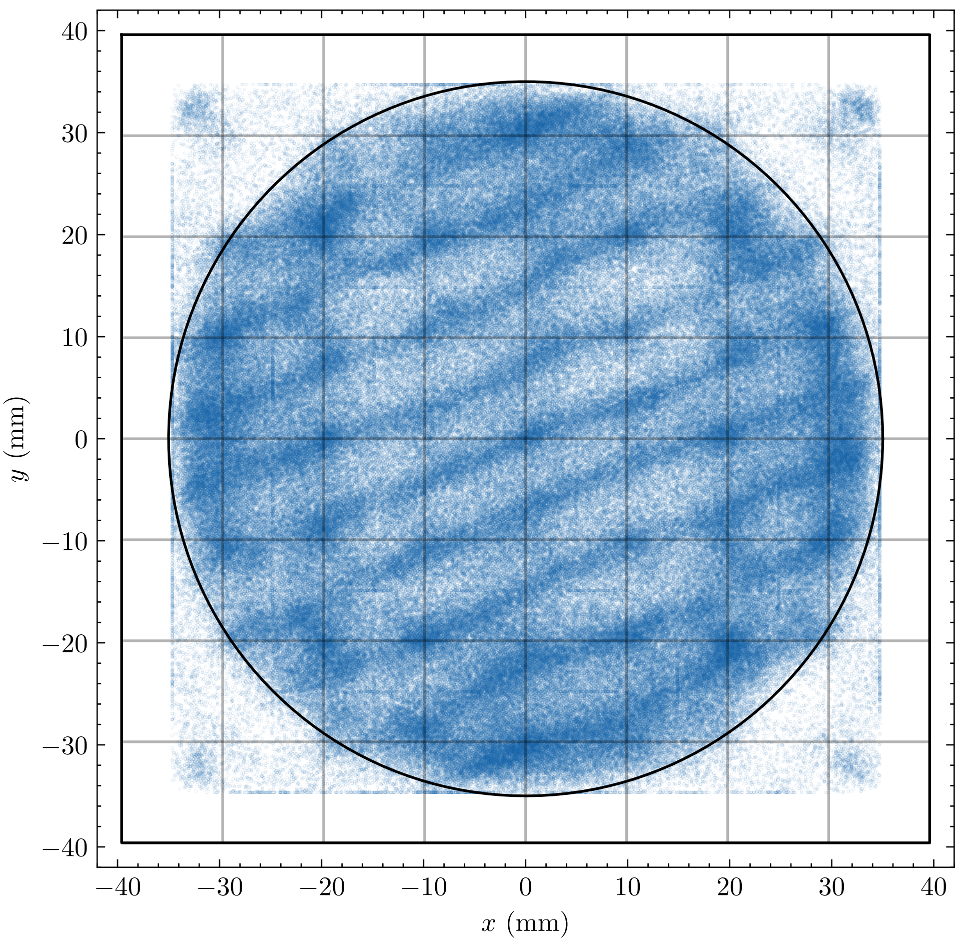

FIGURE 5.20: Reconstructed event positions for a measurement conducted with the stainless steel mask depicted in Fig. 5.19. Only events with at least three active SiPM channels were considered and the reconstruction weights were obtained from the neutron signal calibration. The black square and circle mark the light guide and the beam guide edge, respectively, while the grid represents the light guide grid. The centers of the SiPMs coincide with the center of the grid squares.

background event rate, a second stainless steel mask of identical thickness but without cutouts was used. The datasets from both mask measurements were then combined to estimate the width of the reconstructed cutout distributions. For this purpose, both datasets were projected onto an axis x' that is orthogonal to the mask cutouts and rotated by $-67.5°$ with respect to the x-axis. Only events within the diameter of the beam guide were considered, as defined by the condition

$$\sqrt{x_\mathrm{e}^2 + y_\mathrm{e}^2} < 35\,\mathrm{mm}, \tag{5.31}$$

to exclude regions with low local event densities. The reconstructed events were subsequently binned into intervals of 70 µm width to enable background subtraction. The bin size was chosen to provide fine spatial resolution while maintaining a sufficiently large number of events per bin to allow meaningful uncertainty estimation. In both measurements, the bins can be assumed to be filled at a constant rate. Consequently, the number of events per bin can be treated as Poisson distributed. The bin width was selected such that a minimum of 100 counts per bin was achieved for both datasets within the region of interest. Under this condition, the relative uncertainty is below 10 % and the Poisson distribution can be well approximated by a Gaussian distribution (see Appendix F). The background subtracted bin counts and their uncertainties are then given by

$$\hat{N}_\mathrm{b} = N_\mathrm{b,m} - N_\mathrm{b,bg}\frac{T_\mathrm{m}}{T_\mathrm{bg}} \tag{5.32}$$

and

$$\Delta\hat{N}_\mathrm{b} = \sqrt{N_\mathrm{b,m} + N_\mathrm{b,bg}\frac{T_\mathrm{m}^2}{T_\mathrm{bg}^2}}, \tag{5.33}$$

respectively. Here, $N_\mathrm{b,m}$ denotes the bin counts obtained with the cutout mask, $N_\mathrm{b,bg}$ the background bin counts, T_m the duration of the cutout mask measurement and T_bg the duration of the background measurement. The background subtracted counts projected onto x' are shown in Fig. 5.21. To estimate the width of the event position distribution for each cutout, a least squares fit[10] of the histogram was performed. The event position distribution is assumed to follow the convolution of a Gaussian distribution with a uniform distribution corresponding to the finite width of the line shaped cutouts. The resulting variance of the convolution is given by

$$\sigma_\mathrm{obs}^2 = \sigma_\mathrm{G}^2 + \sigma_\mathrm{c}^2, \tag{5.34}$$

where σ_G^2 and σ_c^2 denote the variance of the Gaussian and the uniform distributions, respectively. The contribution of the uniform distribution can be neglected if its variance, given by

$$\sigma_\mathrm{c}^2 = \frac{d_\mathrm{c}^2}{12}, \tag{5.35}$$

is small compared to the observed variance. Here, $d_\mathrm{c} = 0.3\,\mathrm{mm}$ is the width of the cutout. This yields $\sigma_\mathrm{c}^2 = 7.5\times10^{-3}\,\mathrm{mm}^2$, which is negligible compared to the estimated observed variance of approximately $\sigma_\mathrm{obs}^2 = 1\,\mathrm{mm}^2$. Therefore, the fit function can be approximated by a sum of Gaussian functions given by

$$\hat{G}(x') = \sum_{i=1}^{9} G(x', A_i, \mu_i, \sigma_i), \tag{5.36}$$

[10]Python package: scipy.optimize.curve_fit (v. 1.15.2)

with

$$G(x') = A \exp\left(-\frac{(x'-\mu)^2}{2\sigma^2}\right), \tag{5.37}$$

where A, μ and σ denote the amplitude, centroid and standard deviation of each Gaussian component, respectively. To ensure reliable uncertainty estimation, bins near the edges of the projection range were excluded. In these regions, the number of bin counts falls below the threshold of 100 established above for the background measurement. This reduction is a purely geometric effect arising from the projection of a circular detector area onto a single coordinate. Consequently, the histogram was restricted to the range from $x' = -31.5\,\mathrm{mm}$ to $x' = 31.5\,\mathrm{mm}$, corresponding to half the cutout spacing beyond the expected centers of the outermost cutouts.

The fit results and the corresponding residuals are shown in Fig. 5.21. The extracted fit parameters and the resulting χ^2 value are summarized in Tab. 5.3. In total, 900 bins were included in the fit and 27 parameters were determined. The resulting χ^2 value is therefore expected to follow a χ^2-distribution with 873 degrees of freedom. The symmetric 95 % confidence interval of this distribution is given by $[793, 957]$, which does not include the observed χ^2 value. This indicates that the fit model is not fully consistent with the data. While the residuals are not entirely random, they do not exhibit a clear structure that would suggest an improved fit model. Additionally, the

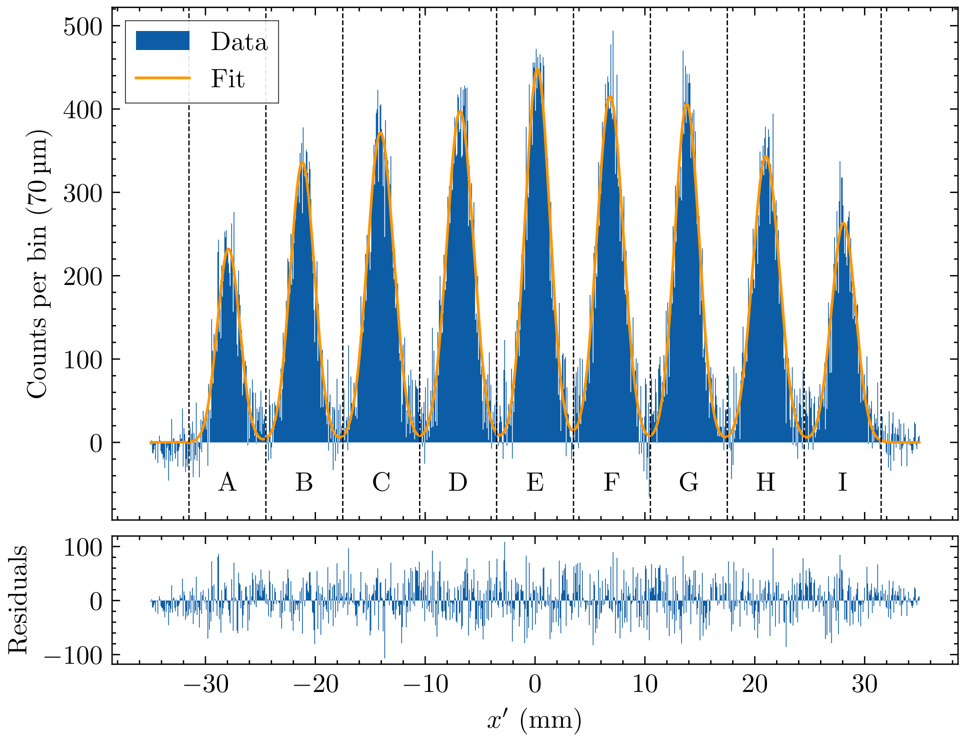

FIGURE 5.21: Background subtracted bin counts for a projection of the reconstructed event position distribution obtained with the stainless steel cutout mask (see Fig. 5.19) onto the x' axis. The x' axis is rotated by $-67.5°$ with respect to the x-axis and is orthogonal to the line-shaped cutouts. The dashed black lines mark the midpoints between the expected cutout centers and each cutout is labeled by a capital letter. The orange line shows a fit of a sum of nine Gaussian functions to the data and the bottom panel shows the corresponding residuals.

Table 5.3: Fit parameters obtained from a fit of a sum of nine Gaussian functions, as defined in Eq. (5.36), to the histogram depicted in Fig. 5.21. The cutouts corresponding to the individual Gaussian components are labeled according to Fig. 5.21.

Cutout	A	μ (mm)	σ (mm)
A	232(8)	−27.90(4)	1.02(4)
B	336(8)	−21.18(3)	1.10(3)
C	372(8)	−14.06(3)	1.20(3)
D	398(8)	−6.83(3)	1.21(3)
E	448(9)	0.22(2)	1.12(2)
F	415(9)	6.83(3)	1.19(3)
G	406(9)	13.81(3)	1.11(3)
H	343(8)	21.02(3)	1.22(3)
I	263(8)	28.07(4)	1.10(3)
Global parameters:			
$\chi^2_{\text{red}} = 1013/873$			

extracted parameters are not fully consistent with expectations. While absolute positioning of the masks is difficult to control, the distance between neighboring cutouts should be constant within uncertainties. Similarly, the extracted widths of the cutout distributions are expected to be consistent. One possible explanation for the observed variations, as well as the elevated χ^2, is that the stainless steel masks were not perfectly flush with the detector surface and may have slightly waved when mounted. In this case, variations in the distance between detector surface and mask could lead to the observed inconsistencies in both width and spacing. Furthermore, the reconstruction bias, which was discussed in Section 5.4.3, causes uneven distortions of the reconstructed cutout distributions. UCNs reflected from the cutout walls may also produce local shape distortions that are not accounted for in the Gaussian model. In some cases, reconstruction edge effects could additionally lead to systematic shifts of the cutout centroids. It therefore appears likely that the observed inconsistencies arise from a combination of these effects. Overall, the observed inconsistencies are on the order of at most 200 µm. This scale is small compared to the extracted widths of the cutout distributions, allowing the results to be used for an estimation of the detector resolution, with the dominant contribution arising from the intrinsic detector resolution rather than the observed systematic distortions. To obtain a single resolution estimate, the weighted mean of all extracted σ values was calculated and the sample standard deviation was used as a measure of the uncertainty. This yields

$$\bar{\sigma} = 1.14(7)\,\text{mm}. \tag{5.38}$$

As a final metric, the average 2σ half width was chosen as a measure of the spatial resolution of the detector prototype. Neglecting edge effects, the resolution towards the center of the light guide can thus be estimated as

$$x_{\text{res}} = 2.3(2)\,\text{mm}. \tag{5.39}$$

The influence of the reconstruction bias on the achievable resolution requires further investigation. Nevertheless, based on these results, it can be concluded that sub-pixel spatial resolution has been achieved.

5.4.7 Comparison of calibrations

The same dataset that was previously used to assess the performance of different reconstruction approaches (see Section 5.4.2) was employed to examine how strongly the reconstructed positions depend on the calibration functions extracted from the averaged neutron and ^{241}Am α-signals (see Section 5.3). To evaluate this, reconstruction weights for the weighted average position reconstruction were extracted from both calibration functions and the event positions were reconstructed using both sets of weights. For each event, the distance between the two reconstructed positions was then computed. The distance between the reconstructed positions is defined as

$$\delta = \sqrt{(x_\mathrm{n} - x_\mathrm{Am})^2 + (y_\mathrm{n} - y_\mathrm{Am})^2}, \tag{5.40}$$

where the indices n and Am refer to neutron and americium calibration, respectively. For this evaluation, only events with three or more active channels were considered. A histogram of the resulting distances is shown in Fig. 5.22. To characterize the distribution, the 95 % quantile

$$\delta_{95} \approx 50\,\mu\mathrm{m} \tag{5.41}$$

was determined. In order to assess whether the observed differences between the reconstructed positions obtained from the two calibration functions are significant, they must be compared to the spatial resolution of the detector. As estimated in Section 5.4.6, the spatial resolution of the detector prototype is approximately 2.3(2) mm. Since the observed deviations are more than an order of magnitude smaller than the spatial resolution, they do not contribute appreciably to the overall position uncertainty. It can therefore be concluded that the differences between reconstructed positions obtained from calibration functions based on averaged neutron and ^{241}Am α-signals are negligible for practical purposes.

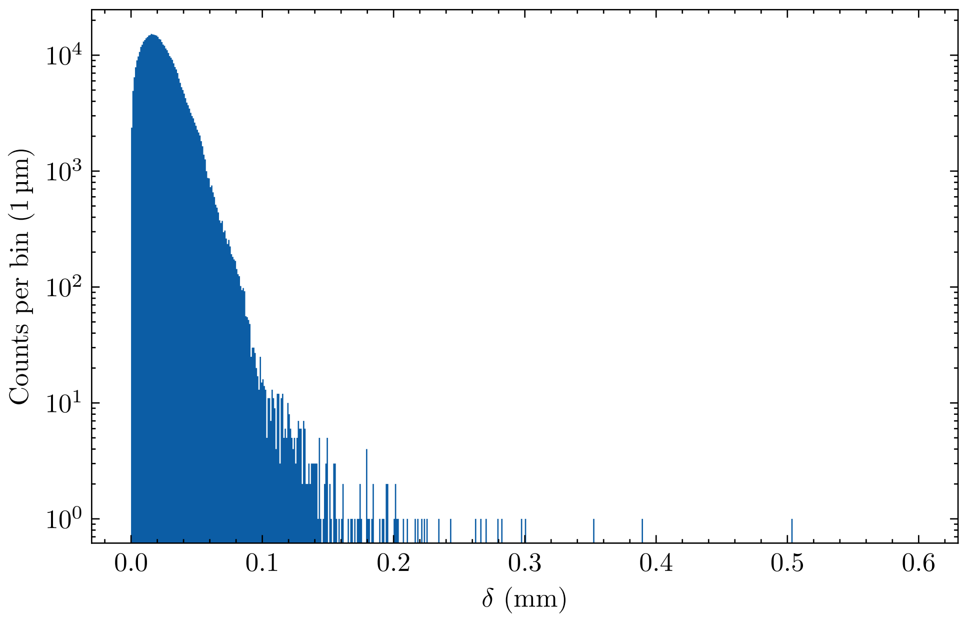

FIGURE 5.22: Distribution of the absolute distances between reconstructed event positions obtained using calibration functions derived from the averaged neutron and ^{241}Am α-waveforms. The distance δ is defined according to Eq. (5.40) and the weighted average approach was used for reconstruction.

5.4.8 Discussion of spatial resolution performance

As part of the study of the spatial resolution of the detector prototype, three different position reconstruction approaches were tested. Two of these approaches assume a scintillation light intensity distribution that is fitted to neutron events, whereas the third computes the average of the active SiPM channel positions, weighted by the amount of light observed by the respective channels. Ultimately, the weighted average approach was found to provide the best reconstruction of the uniform distribution expected for one of the characterization measurements and was therefore used to estimate the spatial resolution of the prototype. However, all reconstruction approaches exhibited reconstruction biases, which promoted the formation of clusters in the reconstructed event position distribution that aligned with the symmetry of the light guide and the SiPM grid. Additionally, about 13 % of all registered neutron events exhibited only two active channels and could not be reconstructed in a meaningful way. In general, the number of active channels was found to be significantly smaller than expected from the ray-tracing simulation performed in the design phase (see Section 4.5.1). This indicates that the scintillation light intensity distribution differs significantly from the assumed isotropic model and may even vary from event to event. Model calculations also showed that events with a small number of active channels were more strongly affected by the reconstruction bias. These results therefore provide a strong motivation to increase the mean number of active channels. This could be achieved by broadening the light intensity distribution, e.g. by increasing the distance between the scintillator and the SiPM grid through a thicker light guide. However, this would also reduce the light intensity observed by the SiPMs closest to the event, so a trade-off between distribution width and signal amplitude must be made. Alternatively, a denser SiPM grid could expose more channels to a sufficiently large fraction of the scintillation light, although this would require adding an additional ASIC to the SiPM readout system if the same detector cross section is to be covered.

To estimate the spatial resolution of the detector, a thin stainless steel mask with nine parallel line-shaped cutouts distributed evenly across the mask was mounted on top of the detector surface. Subsequently, the widths of the reconstructed cutouts were determined and used as a measure of the spatial resolution, which, for the purpose of this study, was defined as the average 2σ half width of the reconstructed cutouts. This yielded a spatial resolution of 2.3(2) mm. While this is sufficient for studies of UCN transport in UCN guides of diameters between 70 mm and 200 mm, which are typically employed at the research reactor TRIGA Mainz and at PSI, it does not fully meet the requirements of the τSPECT experiment. For this application, a resolution of approximately 1 mm or better would be ideal (see Section 3.8). To achieve this, both the statistical uncertainties and the systematic distortion introduced by the reconstruction procedure must be reduced. Statistically, the spatial resolution is in general limited by

1. the randomness of photon emission and detection in the scintillator and the SiPMs,
2. the detector geometry,
3. the SiPM signal noise affecting the signal ToT determination and
4. the uncertainties associated with the conversion of the ToT into a quantity proportional to the number of detected photons.

The uncertainties associated with photon emission and detection are determined by the light yield of the scintillator and the detection efficiency of the SiPMs, respectively. Although the photon detection efficiency varies slightly between different SiPM models and manufacturers, the underlying detection mechanism is the same. Therefore, no significant improvement in this respect can be expected by changing the SiPM model. Since ZnS(Ag), which is used as scintillation material, is an inorganic crystal, which generally have the largest light yield among known scintillation materials [Leo94, p. 165], no significant improvement is expected in this respect either. While the detector geometry itself does not introduce uncertainties, it determines how strongly variations in estimated photon counts affect the estimated event position. This mainly depends on the size of both the SiPMs themselves and the gap between them. Therefore, the spatial resolution can be improved by reducing both of these quantities. As discussed above, this would require adding an additional ASIC to the readout system and an optimization of the distance between scintillator and SiPM grid. The magnitude of the resulting improvement is difficult to estimate, because it also depends on the light intensity distribution, which, as mentioned above, is not sufficiently well known. To reduce the uncertainties originating from the signal ToT measurement and the subsequent conversion into a quantity proportional to the amount of detected light, a direct measurement of the signal integral would be beneficial, since the integral is proportional to the number of detected photons. This would avoid the conversion step altogether and additionally measure a quantity that is usually less sensitive to noise than the ToT. Although the SiPM readout system used in the detector prototype is capable of measuring the signal integral, the available integration time windows are not well suited for the comparatively long signals produced by the scintillator (see Section 5.2 and Section 5.3). Therefore, a faster scintillator would be required to leverage the signal integral mode. Additional benefits of a faster scintillator are further discussed in Section 5.6.5.

To address the reconstruction bias, multiple approaches can be pursued. As mentioned above, the magnitude of the effect could be reduced by increasing the average number of active channels. However, tackling this effect fundamentally within the framework of the tested reconstruction approaches requires better knowledge of the scintillation light intensity distribution. Such knowledge would also be valuable for modeling the detector more accurately in simulations. To measure this distribution, light guides with varying thickness could be employed and an α-source could be mounted at a fixed position in order to record the light distribution on the SiPM grid as a function of the light guide thickness. While such a measurement would provide the desired information, obtaining reliable results is expected to be fairly difficult, because it would require frequent changes of the light guide and reproducible optical contact between scintillator and light guide as well as between SiPM grid and light guide. Furthermore, such a measurement by itself cannot account for boundary reflections at the light guide edges or for channel to channel variations in optical coupling and gain. An alternative approach could be to use a well collimated α-source that can be moved to well controlled positions on the detector surface. By performing a 2D scan across the surface and recording events at each position, a direct relation between the ToT recorded by the individual channels and the source position could be established. Subsequently, a machine learning algorithm, such as a neural network, could be trained to reconstruct the source position from the recorded ToT pattern. Such an approach is expected to significantly reduce reconstruction biases and would also intrinsically account for gain variations, optical coupling variations and edge effects. Moreover, it would remove the need to use a calibration function to transform the measured ToT into a quantity proportional to the amount of detected light. A potential issue of this approach, however, is that predominantly α-particles with momenta parallel to the surface normal of the detector would be measured, whereas α-particles and ^{7}Li nuclei with momentum vectors pointing in all spatial directions are created in neutron events. To avoid this potential bias, a well collimated neutron beam would be required. However, the required neutron flux is available only at large research facilities, whereas the required α-source could be set up in a local laboratory. Therefore, it should first be tested whether using an α-source is sufficient. To partially account for this issue when using an α-source, the source and the detector could be rotated such that the α-particles enter the scintillation layer at a different angle. Since this approach would address several of the remaining limitations of the first prototype simultaneously, it should be investigated as a high-priority calibration strategy in the next stage of detector development.

5.5 Gamma sensitivity and intrinsic background rate

One of the key characteristics of a detector is the intrinsic background rate, i.e. the count rate in absence of radiation fields. In addition to that, a non-negligible level of ambient γ-radiation is typically present in most experimental areas where a UCN detector may be deployed. This radiation can vary over time and is often a significant contribution to the observed background rate. Therefore, both the sensitivity to γ-radiation and the intrinsic detector background are important detector performance parameters. A concrete example illustrating the relevance of the γ-sensitivity is provided by the UCN detector of the τSPECT experiment, where a periodic γ-radiation spike leads to a time-dependent background count rate (see Section 3.7.6).

To assess the sensitivity of the detector prototype to γ-radiation, dedicated test measurements using a γ-source were performed. For these measurements, the detector was moved to the same laboratory in which the test with the ^{241}Am source were performed. However, instead of the α-source, a ^{137}Cs γ-source, which was obtained from radiation protection at the research reactor TRIGA Mainz at JGU, was installed on top of the detector surface. Since the source is enclosed by a metal housing with relatively sharp edges, the detector surface was covered by a 0.1 mm thick stainless steel mask to prevent damage to the conversion layers. In total, measurement data were recorded for 14 h. Afterwards, the source was removed from the setup and an additional 14 h of background data were acquired to estimate the intrinsic background rate of the detector. ^{137}Cs decays predominantly into an excited state of ^{137}Ba, which subsequently deexcites to the ground state emitting γ-radiation with an energy of 661.7 keV [Bro+18a]. The ^{137}Cs source was selected because its form factor was compatible with the detector prototype vacuum chamber. A certificate specifying the initial activity of the source, however, was not available. Consequently, the activity was estimated using a hand-held radiation monitor[11]. During the measurement, the source was fully covered by the radiation monitor and a count rate of about 190 s^{-1} was observed. The sensitivity of the radiation monitor to γ-radiation from a ^{137}Cs source is approximately 35 % [Gmb22]. Since the radiation monitor covered roughly 50 % of the solid angle of the source, the total activity was estimated to be about 1090 s^{-1}. Throughout the measurement, the detector surface covered a solid angle comparable to that of the radiation monitor. Accordingly, it is expected to have been

[11] Graetz CoMo-170

exposed to a γ-rate of approximately $545\,\mathrm{s}^{-1}$. During the measurements, a γ-rate of

$$R_\gamma = 2.3(7) \times 10^{-3}\,\mathrm{s}^{-1} \tag{5.42}$$

at a background rate of

$$R_{\mathrm{bg}} = 11.3(5) \times 10^{-3}\,\mathrm{s}^{-1} \tag{5.43}$$

was measured. This corresponds to an estimated γ-efficiency of

$$\epsilon_\gamma = 4.3 \times 10^{-6} \tag{5.44}$$

for γ-radiation from a ^{137}Cs source, under the assumption that no significant fraction of the radiation was attenuated by the 0.1 mm thick stainless steel mask.

To put this result into context, an equivalent measurement with the τSPECT detector should be performed in the future. In addition, improving the precision of the estimate would require a γ-source with well known activity. While the γ-sensitivity cannot be directly compared to that of the τSPECT detector without an equivalent measurement, the intrinsic background rates of both detectors can still be compared. The measured R_{bg}, which represents a good estimate of the intrinsic background rate, is approximately two orders of magnitude lower than the background count rates typically observed with the current τSPECT detector under comparable conditions and therefore represents a significant improvement in this respect. This improvement is likely due to the neutron event identification algorithm of the prototype, which effectively applies a coincidence criterion by requiring at least two active channels (see Section 5.2.3). Such a criterion cannot be implemented in the current τSPECT detector, since it only provides a single readout channel.

5.6 Event pile-up

At high count rates so-called pile-up can occur in detector systems. Pile-up refers to the merging of multiple physical events into a single reconstructed event by the detector readout. In the case of the detector prototype and the current version of the event identification algorithm, pile-up occurs when a second event happens within the time interval during which channel events are still considered to be part of the same neutron events. To estimate the pile-up rate as a function of measured count rate, a small number of measurements were performed at the research reactor TRIGA Mainz while the reactor was operated in pulsed mode. Compared to DC operation, this generates a significantly higher count rate over periods of a few seconds.

5.6.1 Statistical model

At a constant count rate R, the number of events k occurring within a time interval t follows a Poisson distribution given by

$$P(k) = \frac{(Rt)^k}{k!} \exp(-Rt). \tag{5.45}$$

It can be shown (see Appendix E.3) that the time t required for N events to occur is distributed according to an Erlang distribution defined by

$$\rho(t) = \frac{R^N t^{N-1}}{(N-1)!} \exp(-Rt), \tag{5.46}$$

which reduces to an exponential distribution for $N = 1$. For a fixed event duration T, the probability for pile-up is given by the probability that a second event occurs within the time window T. The pile-up probability is therefore

$$p_{\text{pu}} = \int_0^T \rho(t, N = 1)dt = 1 - \exp(-RT), \tag{5.47}$$

corresponding to the cumulative distribution function (CDF) of the exponential distribution. The average time between observed events, δT_{obs}, can be derived from this distribution. It is given by

$$\delta T_{\text{obs}} = T + 1/R, \tag{5.48}$$

which is the sum of the event duration and the mean waiting time for the next event after the previous event window has ended. This yields

$$R_{\text{obs}} = \frac{1}{T + 1/R} \tag{5.49}$$

for the expected observed count rate. For a fixed total number of N events at constant count rate, the number of pile-up events k is assumed to follow a binomial distribution given by

$$P(k) = \binom{N}{k} p_{\text{pu}}^k (1 - p_{\text{pu}})^{N-k}. \tag{5.50}$$

5.6.2 Pile-up identification and rate dependence

Fig. 5.23 shows the temporal distribution of neutron events recorded by the detector following a reactor pulse. Data acquisition was manually initiated before the reactor pulse and was therefore not synchronized with the pulse, resulting in an arbitrary temporal offset of the time axis. The time distribution exhibits two distinct peaks

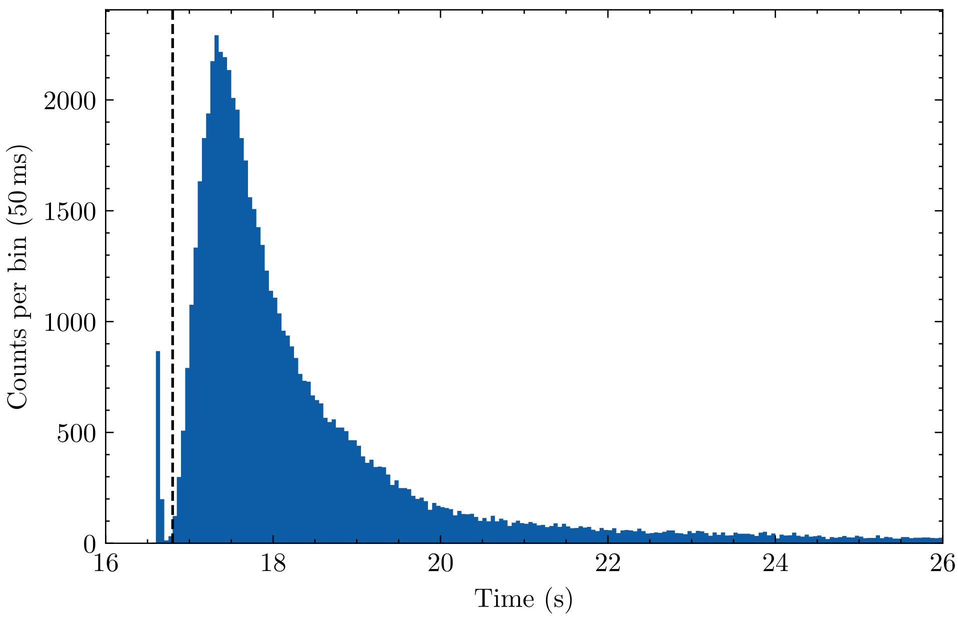

FIGURE 5.23: Temporal distribution of events recorded by the detector prototype following a TRIGA reactor pulse. The dashed black line indicates the visually determined separation between the initial peak attributed to thermal neutrons and γ-radiation and the second peak attributed to UCNs.

separated by a region in which the count rate drops close to background level. The first peak is very narrow and spans only a few time bins, whereas the second peak has a larger amplitude and exhibits a long tail. The first peak can be attributed to thermal neutrons and γ-radiation produced during the 30 ms reactor pulse, which naturally results in a sharp feature. UCNs produced during the pulse require a substantially longer time to travel from the source to the detector via the UCN guides, giving rise to a delayed and temporally spread out second peak. Although the number of thermal neutrons traversing the detector is likely much larger than the number of UCNs, the detection efficiency for thermal neutrons is expected to be significantly lower due to the thickness of the ^{10}B layer, which is optimized for UCN detection (see Section 4.4). Furthermore, as discussed in Section 5.5, the sensitivity to γ-radiation at energies around 660 keV is approximately 4.3×10^{-6}. Hence, it is reasonable to also assume a low overall sensitivity to γ-radiation. Consequently, the UCN peak exhibits a larger amplitude.

At high count rates, the event identification algorithm is expected to classify multiple neutron events as a single event. An exemplary case is shown in Fig. 5.24. The current version of the event identification algorithm groups channel events purely based on temporal coincidence and does not take the spatial arrangement of the SiPMs into account (see Section 5.2.4 for details). Since scintillation light from a neutron event is typically distributed over a relatively small number of neighboring channels, it is unlikely that two temporally close events trigger the same channels. As a result, pile-up events can often be identified by the presence of multiple local maxima in the spatial distribution of active channels. For the purpose of this analysis, a local maximum is defined as a channel whose ToT exceeds that of all its nearest neighbors, including diagonal neighbors. To suppress noise events that trigger only a single channel, it was additionally required that at least one of the neighboring channels exhibit a ToT above the event detection threshold of 250 ns. This approach is expected to

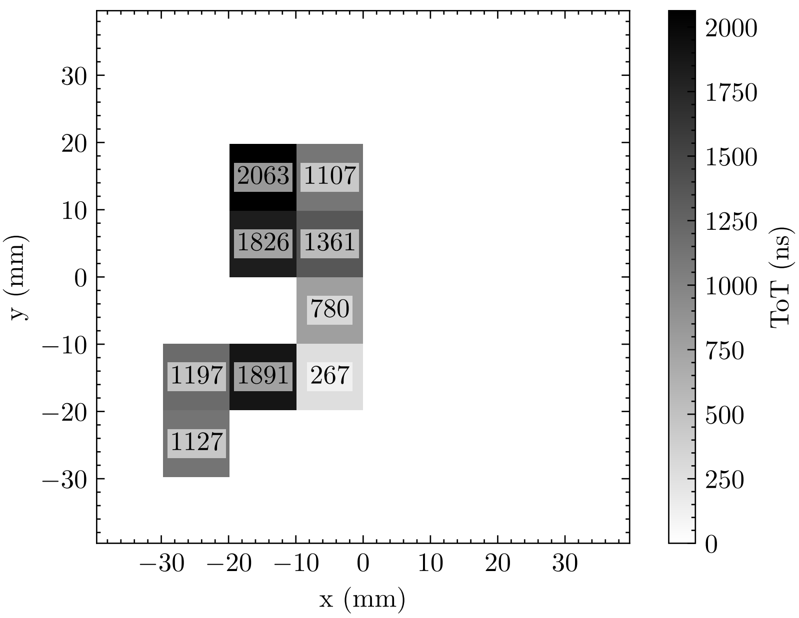

FIGURE 5.24: Exemplary spatial active channel distribution of an event containing more than one local maximum. The squares indicate the effective area covered by an SiPM channel. The numerical values shown on top of the squares denote the signal ToT in nanoseconds recorded by the respective channel.

perform reasonably well as long as the number of neutron events merged into a single reconstructed event remains small. At very high count rates approaching detector saturation, multiple events may overlap such that individual local maxima can no longer be resolved, causing this method to fail. To estimate the pile-up fraction as a function of count rate, the fraction of events containing multiple local maxima r_{m} was determined for varying count rates. For this purpose, events recorded during the reactor pulse were grouped into bins. To ensure sufficient statistical sensitivity, the bin edges were chosen such that each bin contained a constant number of $N = 2000$ events. This guarantees that even for low pile-up probabilities on the order of 1 %, a sufficient number of pile-up events are observed to allow a meaningful pile-up rate estimate. Because this binning strategy performs poorly when the count rate changes rapidly, the initial thermal neutron peak was excluded from the analysis. To achieve this, a separation time between the thermal neutron peak and the UCN peak was selected visually for each of the five recorded reactor pulses. For the pulse shown in Fig. 5.23, this separation is indicated by the dashed vertical line. After binning, the observed bin count rate

$$R_{\mathrm{obs}} = \frac{N}{t_{\mathrm{bw}}}, \tag{5.51}$$

and the pile-up fraction

$$r_{\mathrm{m}} = \frac{M}{N}, \tag{5.52}$$

were determined, where t_{bw} denotes the bin width and M the number of reconstructed events containing more than one local maximum. It should be noted, however, that r_{m} only represents a lower bound on the temporal pile-up probability p_{pu}, since it only counts pile-up events that are spatially resolvable through multiple local maxima. The uncertainties of R_{obs} and r_{m} were estimated from the statistical distributions of t_{bw} and M, respectively. As discussed in Section 5.6.1, the time required for N events to occur at a constant count rate follows an Erlang distribution (see Eq. (5.46)). Assuming that the count rate is approximately constant within each bin, t_{bw} is therefore also described by an Erlang distribution. For sufficiently large number of events N, such as $N = 2000$, the distribution can be approximated by a Gaussian (see Appendix F), yielding

$$\Delta t_{\mathrm{bw}} = \frac{\sqrt{N}}{R_{\mathrm{obs}}}. \tag{5.53}$$

Propagation of uncertainty then gives

$$\Delta R_{\mathrm{obs}} = \frac{R_{\mathrm{obs}}}{\sqrt{N}}. \tag{5.54}$$

Similarly, the number of events with multiple local maxima M is modeled as a binomial distribution with N trials and success probability r_{m}. The corresponding uncertainty is given by

$$\Delta r_{\mathrm{m}} = \sqrt{\frac{r_{\mathrm{m}}(1 - r_{\mathrm{m}})}{N}}. \tag{5.55}$$

This uncertainty can be interpreted as Gaussian only when the binomial distribution is reasonably well approximated by a Gaussian distribution. While this condition is satisfied for most data points, the approximation deteriorates for the very small values of r_{m} observed at low count rates (see Appendix F). Nevertheless, it remains sufficient to provide at least a rough estimate of the uncertainty.

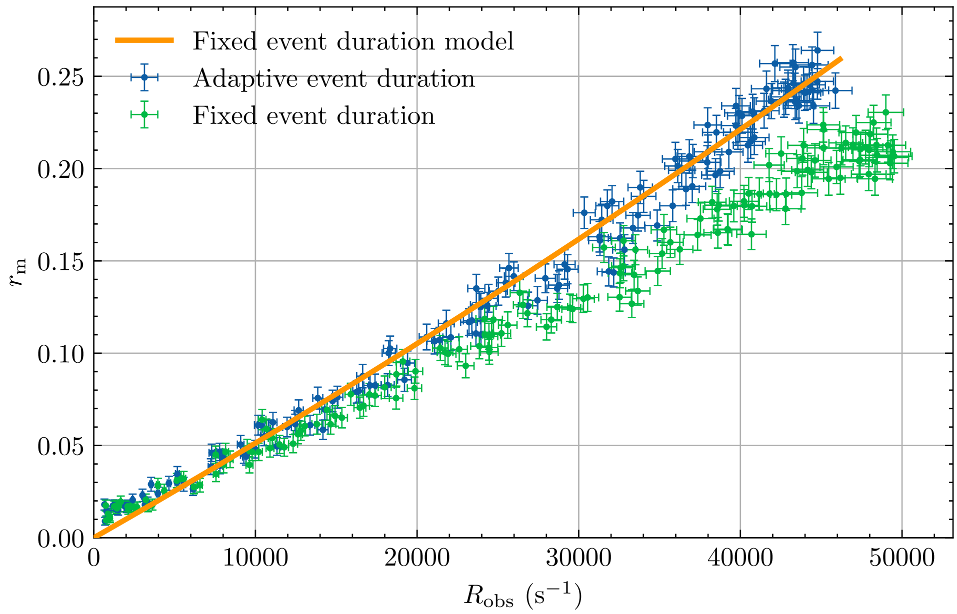

FIGURE 5.25: Pile-up fraction r_m as a function of the observed count rate R_{obs} for an event identification algorithm using an adaptive event duration (blue), a fixed event duration (green) and the fixed duration statistical model (orange).

To enable a direct comparison with the statistical model introduced in Section 5.6.1, the event identification algorithm was modified to use a fixed event duration instead of an adaptive one. For comparability, the fixed duration was set to 5 µs, corresponding to the minimum event duration of the adaptive algorithm (see Section 5.2). The results for both algorithm versions, along with the model predictions, are shown in Fig. 5.25. At first glance, the good agreement between the fixed duration model and the adaptive duration data appears counterintuitive. However, this is most likely coincidental. The discrepancy between the fixed duration model and the corresponding data is instead attributed to limitations in the pile-up identification procedure, which, as explained above, fails to identify pile-up events that are not spatially resolved. Such cases lead to a systematic underestimation of the pile-up fraction, consistent with the observed behavior. Because the fixed event duration was chosen to match the lower limit of the adaptive event duration, a higher pile-up fraction was expected for the adaptive algorithm at larger count rates. For the present parameter choice, the adaptive results happen to agree well with the fixed duration model. At low count rates below $R_{obs} = 10\,000\,s^{-1}$, both algorithms yield consistent results. However, the measured pile-up fraction exceeds the model prediction. This excess is likely caused by a constant background contribution that is not included in the model. Background events or electronic noise can generate additional local maxima within a spatial channel event cluster when coinciding with genuine neutron events, which mimics pile-up. The results indicate that a suitably chosen fixed event duration performs better at high count rates than the adaptive event duration currently employed.

5.6.3 Spatial separation of pile-up events

To estimate the potential reduction in pile-up achievable by incorporating spatial clustering into the event identification algorithm, the fraction r_{sep} of pile-up events that is cleanly separable was determined. In this context, events were classified as cleanly separable if no uninterrupted path of active channels existed between local

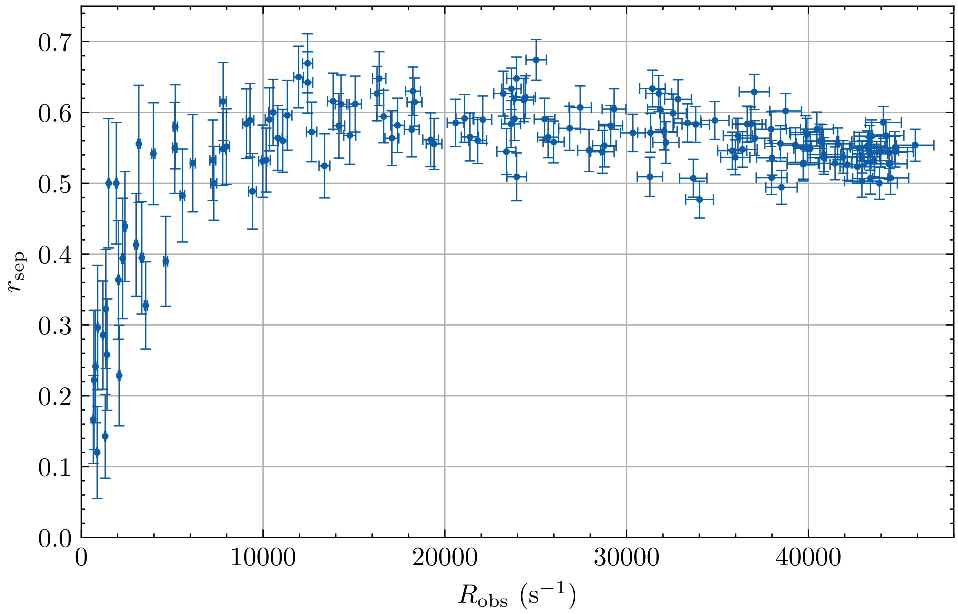

FIGURE 5.26: Fraction of spatially separable pile-up events r_{sep} as a function of the observed count rate R_{obs}. See main text for the definition of spatially separable events.

maxima. In such cases, the individual event positions could be reconstructed without additional constraints compared to single-event reconstruction. At larger count rates, pile-up events containing more than two neutron events become more frequent, reducing the likelihood of clean separation. Consequently, r_{sep} is expected to decrease with increasing count rate. An individual value of r_{sep} was therefore determined for each bin used in the pile-up analysis. From a statistical standpoint, each pile-up event constitutes a Bernoulli trial. The number of cleanly separable pile-up events M_{sep} thus follows a binomial distribution, yielding an uncertainty

$$\Delta r_{\mathrm{sep}} = \sqrt{\frac{r_{\mathrm{sep}}(1 - r_{\mathrm{sep}})}{M}}, \tag{5.56}$$

where M is the total number of pile-up events in the respective bin. Fig. 5.26 shows r_{sep} as a function of the observed count rate R_{obs}. As expected, a slight decrease of r_{sep} is observed at larger count rates. More strikingly though, a sharp increase of r_{sep} occurs at low count rates. This coincides with the range in which the measured pile-up fraction exceeds the model prediction in Fig. 5.25. These events are therefore likely dominated by single peak neutron events in which noise produces an additional local maximum that is part of the same cluster. Excluding bins with $R_{\mathrm{obs}} < 10\,000\,\mathrm{s}^{-1}$, it can be concluded that incorporating spatial clustering could reduce the pile-up fraction by at least 50 % across the observed count rate range.

5.6.4 Gamma pile-up

Finally, events recorded during the thermal neutron peak were examined. In addition to typical neutron events as e.g. shown in Fig. 5.8 or Fig. 5.24, a significant fraction of events exhibited multiple isolated active channels that were not part of any cluster. An example is shown in Fig. 5.27. During the reactor pulse a high flux of thermal neutrons and an intense γ-radiation field are present. However, the detection mechanism for thermal neutrons is identical to UCNs, implying similar event signatures. The absence

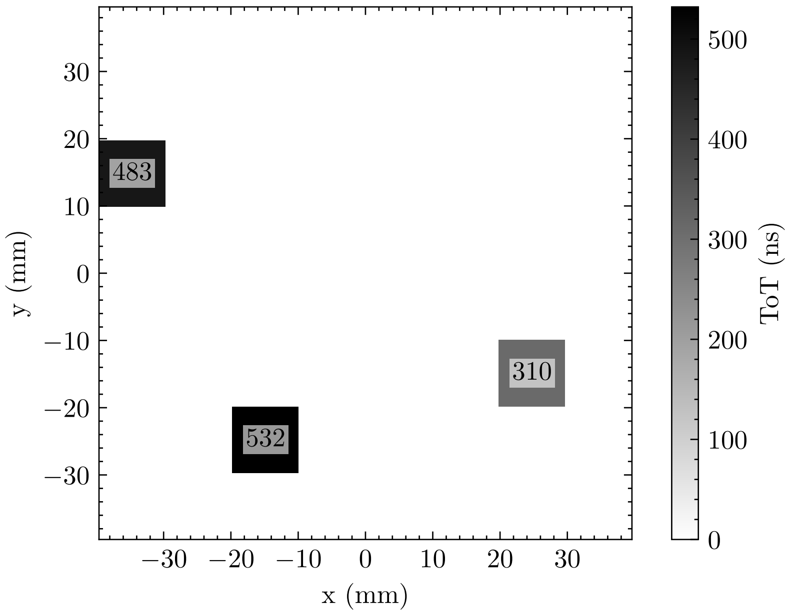

FIGURE 5.27: Exemplary event recorded during a TRIGA reactor pulse, characterized by high thermal neutron and γ-count rates. The squares indicate the effective area covered by the triggered SiPM channels, with the ToT encoded in the grayscale. The numerical values on top of the squares denote the recorded ToT in nanoseconds.

of spatial clustering in these events therefore strongly suggests that the observed single channel cluster events are produced by γ-radiation. Individually, such events would not pass the event identification algorithm, but at sufficiently high γ-rates they can overlap temporally, thereby satisfying the requirement of two or more active channels. Consequently, these events provide a useful test case for studying the characteristic signatures of γ-events. In addition to confirming the shorter signal widths already observed in waveform studies (see Section 5.3), these events demonstrate that γ-events typically produce clusters involving at most two active channels. Introducing spatial clustering in the event identification algorithm would therefore also significantly reduce the γ-background rate at high γ-rates, such as those present during TRIGA reactor pulses.

5.6.5 Discussion

While Fig. 5.25 provides a reasonable estimate of the pile-up behavior over a reasonable range of count rates, several limitations must be noted. The analysis assumes a constant count rate within each bin, which becomes questionable at the steep flanks of the UCN peak. Additionally, the pile-up identification relies on spatially resolvable local maxima and therefore fails when events overlap too strongly. This limitation is reflected in the comparison between the fixed duration analysis and the statistical model, indicating a systematic underestimation of the true pile-up fraction. An improved measurement of the pile-up rate would require a high yield DC UCN source with adjustable and well monitored output. By comparing the detector count rate to that of an independent monitoring detector, the true pile-up fraction, including unresolved events, could be determined. However, such a setup was not available for this study. When interpreting the results of this study, it should be noted that the event identification algorithm was developed using data recorded at low count rates of approximately 500 s^{-1}. As demonstrated here, this algorithm could be improved such that the pile-up fraction across the tested count rate range reduces by approximately 50 % by incorporating spatial clustering in addition to temporal grouping.

While reconstructing individual event positions becomes challenging when clusters overlap strongly, the correct number of detected neutrons may still be recovered by identifying local maxima in the spatial channel distribution. However, since noise can also produce additional local maxima, additional criteria such as a minimum prominence may be required. The impact of such criteria would need to be carefully evaluated, as events occurring between channels may produce a relatively uniform ToT across multiple channels in a cluster and could therefore be missed. Further improvements could be achieved by refining the temporal clustering logic, e.g. by optimizing the time interval during which repeated channel triggers are associated with the same event. Ultimately, however, the maximum sustainable count rate is limited by the intrinsic signal duration, which in the present setup is dominated by the scintillator decay time constants. A substantial reduction of pile-up at the investigated count rates would therefore require a faster scintillator. Suitable candidate materials have e.g. recently been investigated at Los Alamos National Laboratory [Kri+26]. However, the tested materials exhibit a scintillation light spectrum in the near ultraviolet range. Therefore, the transmittance in this section of the electromagnetic spectrum of all detector components involved in the light detection must be evaluated first to assess whether these materials are suitable for the present detector design.

5.7 Comparison to CASCADE detector

In order to assess the neutron detection efficiency of the prototype, it was compared to a commercially available UCN detector[12]. In contrast to the developed prototype, the CASCADE detector relies on a combination of ^{10}B, a counting gas and high voltage to detect UCNs [CDT18]. To separate the counting gas from the beam guide vacuum, a thin aluminum separation window ($V_F = 54$ neV, see Appendix A) is employed. This introduces an additional Fermi potential barrier that UCNs must overcome in order to be detected. To compensate for this, the CASCADE detector is typically operated with an additional 80 cm vertical guide, in which neutrons are gravitationally accelerated. Fig. 5.28 shows schematic cross sections of the setups used to perform the comparison measurements with both detectors. In addition to measurements of the raw count rate R_{raw}, dedicated measurements of the background count rate R_{bg} were performed with both detectors. For this purpose, the UCN flux was blocked at the beam port using the safety shutter. The UCN count rate is then given by

$$R_{\text{UCN}} = R_{\text{raw}} - R_{\text{bg}}. \tag{5.57}$$

All measured count rates are tabulated in Tab. 5.4. Overall, the measured count rate for the detector prototype is about 25 % larger than that of the CASCADE detector. Furthermore, the measured background rate is significantly smaller for the prototype. However, the position of the detectors relative to both shielding and reactor was different during the measurement. Therefore, the background rates should not be compared directly. In general, a comparison of detection efficiencies is not trivial, as it strongly depends on the probed UCN phase space and energy spectrum. In particular, the additional 90° bend used in the CASCADE measurement setup

[12] CDT CASCADE Detector Technologies GmbH, CASCADE-U

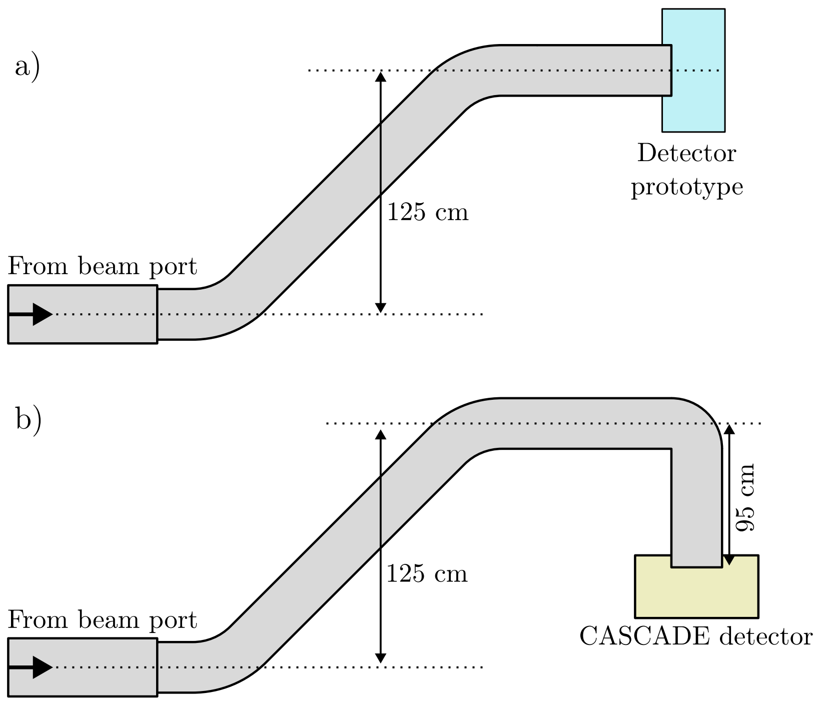

FIGURE 5.28: Schematic cross section of the beam guide setups used to compare the detector prototype (a) with a commercially available CASCADE detector (b). Both detectors were connected to beam port C of the research reactor TRIGA Mainz using a 80 mm to 70 mm beam guide diameter adapter. Downstream the adapter, an S-shaped guide section guides UCNs upwards by about 125 cm. For the CASCADE detector setup an additional 80 cm vertical guide and a 90° bend were added downstream of the upper horizontal neutron guide section.

complicates the comparison, since higher-energy neutrons that are still detectable in the prototype setup may be lost at this bend, effectively reducing the count rate measured with the CASCADE detector. To assess which fraction of the count rate measured with the detector prototype is produced by such neutrons, the count rate of a measurement conducted with a 0.1 mm stainless steel mask placed in front of the prototype surface was evaluated. During this measurement, a UCN count rate of $R_{\mathrm{UCN}} = 64.6(1)\,\mathrm{s}^{-1}$ was observed. These neutrons could not be stored in a stainless steel vessel for a significant period of time and can thus be regarded as non storable in the context of this measurement setup. Assuming that these neutrons would be lost in the 90° bend, this rate can be subtracted from the measured prototype count rate. The resulting rate remains about 11 % larger than that recorded with the CASCADE detector. Moreover, the 90° bend likely still transmits a certain fraction of these

TABLE 5.4: Count rates measured at beam port C of the research reactor TRIGA Mainz for a CASCADE detector and the developed detector prototype. For the masked configuration, the prototype surface was covered by a 0.1 mm thick stainless steel mask.

Detector configuration	R_{raw} (s^{-1})	R_{bg} (s^{-1})	R_{UCN} (s^{-1})
CASCADE	507.2(4)	57.0(1)	450.2(4)
Prototype	567.5(8)	3.81(5)	563.7(8)
Masked prototype	69.6(1)	4.96(7)	64.6(1)

neutrons. Therefore, the prototype's detection efficiency in this particular setup can be estimated to be between 11 % and 25 % larger than that of the CASCADE detector.

An interesting detail is that the observed background count rate for the prototype was about 25 % larger during the masked measurement compared to the unmasked one. This can likely be explained by the fact that the two measurements were conducted three days apart. In between the detector prototype was removed from the measurement setups multiple times. During this process, parts of the radiation shielding installed around the detector may have been slightly shifted, increasing the exposure of the detector. However, this hypothesis was not further tested, as this effect was only noticed after the beam time at the TRIGA reactor had already ended.

In summary, the detector prototype exhibits a higher detection efficiency than the CASCADE detector for the probed UCN phase space and energy spectrum. This can mainly be attributed to the fact that the prototype's ^{10}B surface is directly exposed to vacuum, thereby reducing transport losses. Furthermore, the prototype requires neither high voltage nor counting gas, which significantly reduces the required support infrastructure. However, the CASCADE detector is capable of handling count rates of up to $4 \times 10^7\,\mathrm{s}^{-1}$ [CDT18], whereas the count rate of the prototype is currently limited to below $1 \times 10^5\,\mathrm{s}^{-1}$, depending on the tolerable pile-up rate (see Section 5.6.2). In addition, the CASCADE detector also provides spatially resolved UCN detection. For this purpose, each event is assigned to one of 64 pixels with a size of $12.5 \times 12.5\ \mathrm{mm}^2$ [CDT18], which effectively limits its spatial resolution. Consequently, the spatial resolution of 2.3(2) mm estimated for the detector prototype is already superior to that of the CASCADE detector. The detection efficiency of the current τSPECT detector has also been compared with that of the same CASCADE detector by evaluating the number of neutron counts registered in both detectors after storing UCNs in an aluminum ($V_\mathrm{F} = 54\,\mathrm{neV}$, see Appendix A) storage bottle. In this measurement, the ratio of counts and therefore the relative detection efficiency was found to be 1.06(7) [Aul+24]. Since the UCNs were stored in a container with a Fermi potential significantly lower than that of the stainless steel UCN guides ($V_\mathrm{F} = 184\,\mathrm{neV}$, see Appendix A), the resulting UCN energy spectrum is expected to be shifted towards substantially lower energies than the energy spectrum probed in this study. Therefore, the results cannot be compared directly. Nevertheless, the measurement suggests that the detection efficiency of the prototype is at least comparable to that of the current τSPECT detector.

Chapter 6

Conclusion and outlook

6.1 τSPECT cleaning study

This work investigated marginally trapped neutrons in the τSPECT experiment as an important systematic effect for a neutron lifetime measurements, using different spectral cleaning configurations. Additionally, the dataset recorded for this study was used to examine differences in the temporal distribution of neutron events in the counting phase after different storage times. Such differences could for the first time be observed in the τSPECT experiment and part of the required corrections for different cleaning configurations were quantified. The observed differences can likely be attributed to an evolution of the phase space of the stored neutrons during the storage phase. Especially for deep cleaning, significant differences between the counting distributions were observed, dominated by an offset between the distributions along the time axis. This is consistent with deep cleaning depleting a region of phase space that is repopulated during longer storage times but remains underpopulated after shorter storage times. In addition to an offset on the time axis, the counting distributions also differ in shape. This can likely be explained by the fact that the stored neutron ensemble requires a certain time after filling to reach an equilibrium phase space distribution. Correcting for an offset between the distributions is relatively straightforward and introduces only small additional uncertainties. In contrast, correcting for shape differences between the distributions is more complex. Because the counting phase duration is a significant fraction of the neutron lifetime, shape corrections require an assumed neutron lifetime value and are only viable if their dependence on the assumed value remains negligible compared with the final uncertainty. For the present dataset, this requirement was found to be satisfied. However, for future, more precise determinations of the neutron lifetime the magnitude of this effect must be reduced. In a τSPECT like geometry, this could be achieved by shortening the counting phase, e.g. by moving the detector further into the trap or by modifying the magnetic field to accelerate the phase space evolution. In the current setup, measurements at intermediate storage times would help determine the phase space evolution time scale and assess whether an extended short storage time could reduce shape corrections sufficiently to compensate for the loss in statistical sensitivity. Additionally, further studies with extended counting phase duration are required to estimate the effect of a finite counting phase in the context of differing counting distributions shapes.

The derived shape corrections were subsequently applied to the data recorded in the cleaning study to determine a storage time constant for each configuration. However, the results showed neither a plateau nor a peak with increasingly deep cleaning, which were the two previously expected scenarios. Instead, the data suggest a plateau at moderately deep cleaning followed by a further increase at deeper cleaning. The statistical precision is at this time insufficient for a definitive conclusion, but the results motivate additional, higher-precision measurements in the region between cleaning positions of 1800 mm and 1760 mm. Furthermore, the results clearly showed that both shallow cleaning and short cleaning at a deep cleaning position are insufficient to effectively remove marginally trapped neutrons. As a next step, complementary measurements with the same cleaning configurations but different storage times could provide an important consistency check.

Overall, this work demonstrates that time dependent phase space evolution of the stored UCN ensemble can lead to measurable biases in neutron lifetime analyses if differences between the counting distributions after different storage times are not properly accounted for. The procedure developed in this work provides a correction strategy for the current τSPECT dataset at the target precision of 1 s. At the same time, the observed dependence of the corrections on the phase space evolution of the stored ensemble shows that a more detailed characterization of this evolution is required for future measurements aiming at substantially higher precision. One route towards such a characterization is a spatially resolving detector. Such a detector would enable more detailed studies of the stored phase space distribution by recording event positions in addition to event times. In the radially confining magnetic field, this would also provide sensitivity to the stored UCN energy spectrum and allow to test its long term stability throughout multiple months of beam time. For the current geometry, such a detector would require a spatial resolution of 1 mm to 2 mm, a rate capability up to $1 \times 10^4\,\mathrm{s}^{-1}$, and coverage of the full trap cross section of approximately 100 mm diameter.

6.2 Detector prototype

The second part of this thesis concerned the development and characterization of a detector satisfying the aforementioned requirements. The detector uses a ^{10}B conversion layer stacked with a ZnS(Ag) scintillation layer. The high energy ions created in the ^{10}B neutron capture reaction generate scintillation light in the scintillation layer, which is guided onto an array of SiPMs by a light guide. In a first step, geometric constraints of the light guide were derived from external constraints. Optimized light guide dimensions within these constraints were then determined using a ray-tracing simulation. However, because the optical properties of the scintillation layer are not well known and the required knowledge is difficult to obtain, the results must be treated with caution. The primary objective of the first prototype was therefore to perform characterization studies. For the first prototype, a geometrically simpler square light guide with an 8×8 SiPM grid was chosen. This avoided additional edge effects from a circular to square cross section transition while preserving a central detector region expected to be representative of the center of a light guide geometry suitable for τSPECT. A dedicated SiPM test setup was set up and used to measure the breakdown voltage of 89 SiPMs with a precision of approximately 3.6 mV. The 64 best matched devices were selected, limiting gain variations from breakdown voltage differences to approximately 3 %. As the final step of the design process, PCBs and a vacuum chamber were designed and the detector assembled.

Subsequently, the prototype was characterized at the UCN source at the research reactor TRIGA Mainz and via laboratory α-source measurements. A neutron event identification algorithm was established that primarily relies on clustering SiPM channel signals in time. To convert the time over threshold (ToT) values provided by the readout system into a quantity proportional to the amount of light detected by each SiPM, the SiPM signal shape in neutron events was analyzed in detail. In addition to providing the required ToT calibration, this analysis showed that the slowest scintillator time constant is approximately 7.5 µs, which is substantially longer than specified in the manufacturer's datasheet but consistent with values reported elsewhere in the literature. Direct measurements of the signal integral would avoid the ToT calibration and provide a quantity less sensitive to signal noise. However, the signal width is currently outside the dynamic range of the signal integrating branch of the readout system. The slow scintillation time constant also affects the rate capability, which is primarily limited by pile-up which occurs more frequently for longer scintillation time constants. The pile-up fraction was measured to remain below 5 % up to an observed count rate of $1 \times 10^4\,\mathrm{s}^{-1}$ and including spatial clustering in the event identification algorithm reduced the pile-up fraction by approximately 50 % over the probed range up to $5 \times 10^4\,\mathrm{s}^{-1}$. The detector is therefore capable of handling the count rates expected in τSPECT. However, the scintillation time constant remains the limiting factor for the rate capability and could limit the suitability of the detector in other applications and future experiments.

Spatial resolution is the central feature of the detector prototype and the main limitation of the current version was identified in this area. In total, three different reconstruction approaches were tested and a weighted average approach was found to perform better than approaches fitting a light intensity distribution to each event. However, all approaches were found to exhibit a significant reconstruction bias, which increased as the number of active channels decreased. Moreover, none of the approaches fully accounted for effects introduced by the outer light guide edges. Additionally, approximately 13 % of all events exhibit only two active channels, which is insufficient for a meaningful position reconstruction. Together, these findings show that the scintillation light is currently not shared among a sufficient number of channels. They also show that the scintillation light intensity distribution differs significantly from the isotropic point source model assumed in the ray-tracing simulation. By reconstructing the width of thin, line shaped cutouts in a stainless steel mask mounted on top of the detector surface, the spatial resolution of the prototype was estimated using the average 2σ half width of the reconstructed cutouts, yielding 2.3(2) mm. The prototype therefore achieved sub pixel resolution and a resolution sufficient to perform studies of UCN transport in guides, but it does not yet meet the approximately 1 mm resolution requirement of the τSPECT experiment.

In addition to spatial reconstruction, the suitability of the prototype for UCN measurements depends on its ability to suppress background. As part of this evaluation, characterization measurements of the γ-sensitivity and the intrinsic background rate were conducted. The sensitivity to the γ-radiation from ^{137}Cs was measured to be 4.3×10^{-6}. Since the source activity was not certified and had to be estimated independently, this value should be interpreted as an order of magnitude estimate. In the same measurements, the intrinsic background rate of the detector in absence of significant radiation fields was measured to be approximately $1 \times 10^{-2}\,\mathrm{s}^{-1}$, which is about two orders of magnitude less than observed in the current τSPECT detector. This can likely be attributed to the time clustering performed as part of the event identification, which acts as an effective coincidence criterion and is not available for a single channel detector such as the τSPECT detector. Finally, the prototype was compared with a CASCADE detector to estimate the detection efficiency. For the probed UCN spectrum, the prototype measured a count rate that was approximately

11 % to 25 % larger than that measured with the CASCADE detector, suggesting that the detection efficiency of the prototype is slightly better. It must also be considered that the detection efficiency is energy dependent and may therefore vary between different applications. The spatial resolution of the prototype is superior to that of the CASCADE detector, which assigns each event to one of 64 pixels with a size of 12.5×12.5 mm^2 [CDT18]. On the other hand, the rate capability of the CASCADE detector is still about three orders of magnitude better.

Based on the characterization measurements, several avenues for improvements were identified. One of the most important objectives for a future detector iteration is improved spatial resolution. The characterization measurements showed that this requires both a reduction of the reconstruction bias and a reduction of the statistical reconstruction uncertainty. A central step is therefore to increase the average number of active SiPM channels per neutron event. This could be achieved by increasing the distance between the scintillator and the SiPM grid plane, thereby broadening the scintillation light intensity distribution and by using a denser SiPM grid with smaller gaps and potentially also smaller SiPM size. This would require additional readout channels, which can be accommodated by adding additional ASICs to the readout system. Such an improvement would also reduce the statistical uncertainty of the position reconstruction, as more fine grained position information would be available. Currently, both the rate capability and the signal reconstruction are limited by the slow scintillation time constant. A faster scintillator would reduce the signal width, thereby decreasing the pile-up probability and potentially bringing the signal into the dynamic range of the integrating branch of the readout system. This would allow the signal integral to be measured directly instead of relying on a signal shape based ToT calibration. Since the signal integral is expected to be less sensitive to noise than the ToT, this would improve the precision of the position reconstruction. For these reasons, a faster scintillator would represent a significant overall improvement. Suitable candidate materials exhibiting shorter scintillation time constants have e.g. recently been tested at Los Alamos National Laboratory [Kri+26]. To reduce reconstruction bias, especially near the edges, additional calibration measurements with a collimated movable α-source could be conducted. By scanning the source across the detector surface, a direct relation between ToT or signal integral and the source position could be established. This calibration could be used to correct reconstruction approaches or to train a machine learning model, such as a neural network, to infer the event position. The latter approach would also intrinsically account for variations in channel gain or imperfections in the optical coupling between scintillator, light guide and SiPMs. Since this approach could address multiple remaining issues simultaneously, it should be pursued with high priority.

Overall, the prototype demonstrates that a ^{10}B-based scintillation detector read out by an SiPM array is a viable concept for spatially resolved UCN detection. The detector achieves sub pixel spatial resolution, a low intrinsic background rate, sufficient rate capability for the count rates expected in the τSPECT experiment and an efficiency that is slightly better than that of a CASCADE detector for the investigated UCN spectrum. The main remaining limitation is the spatial resolution, which does not yet meet the approximately 1 mm resolution ideally required by the τSPECT experiment. The current limitations are primarily insufficient light sharing between SiPM channels, reconstruction bias and edge effects. The characterization measurements provide a clear path for further development of the detector concept to address these limitations.

Appendix A

Fermi potential of selected materials

As derived in Section 2.1.1, the Fermi potential of a material is given by Eq. (2.12). Accordingly, only the number density n and the coherent neutron scattering length b are required to compute the Fermi potential of a given material. For composite materials, the number density of the individual components n_i can be computed and the Fermi potential of the material is then given by

$$V_\mathrm{F} = \frac{2\pi\hbar^2}{m_\mathrm{n}} \sum_i n_i b_i, \tag{A.1}$$

where b_i is the coherent neutron scattering length of the respective component and m_n is the neutron mass. For naturally occurring elements, averaged coherent scattering lengths and number densities are available such that the Fermi potential can be computed without explicitly accounting for the contributions of individual isotopes. For most of the materials listed below, the mass density ρ, rather than the number density, was obtained from the literature. The corresponding number density is then given by

$$n_i = \frac{a_i \rho N_\mathrm{A}}{\sum_i a_i M_i}, \tag{A.2}$$

where N_A is the Avogadro constant and M_i and a_i are molar mass and atomic fraction of the respective material. In cases where only the mass fraction w_i is available, the number density of the respective material is given by

$$n_i = \frac{w_i \rho N_\mathrm{A}}{M_i}. \tag{A.3}$$

For materials with an isotopic composition different from the natural abundance, the number density was assumed to be equal to that of the naturally occurring material. The only exception from this is solid deuterium, for which a density corresponding to an isotopically pure crystal could be identified in the literature (see Tab. A.2). The Fermi potentials of all materials, for which a Fermi potential is referenced in this thesis are summarized in Tab. A.1. For the computation of the listed Fermi potentials, a neutron mass of $m_\mathrm{n} = 1.674\,927\,500\,56 \times 10^{-27}$ kg [Moh+25] was assumed. All material properties required to compute the Fermi potentials listed in Tab. A.1 are listed in Tab. A.2. For most compound materials, manufacturer datasheets were consulted to obtain values for composition and density. In general components with a mass ratio w_i below 1 % were neglected. For the τSPECT stainless steel guides, the manufacturer[1] specified only the steel grade (1.4435). Material properties were therefore taken from

[1]NEUMO GmbH + Co. KG

Table A.1: Fermi potential of selected materials.

Material	V_F (neV)
Copper	171
Aluminum	54
Quartz glass (SiO_2)	91
Zinc sulfide (ZnS)	56
Stainless steel (1.4435, 316L)	184
^{58}NiMo (85:15)	313
Solid deuterium (D_2)	107
Boron (detector surface)	-0.4 - 39i

Table A.2: Material properties used to compute the Fermi potentials listed in Tab. A.1. Blank entries indicate that the respective quantity was not required for computing the Fermi potentials.

Material	ρ (g/cm^3)	b (fm)	M (g/mol)
Cu	8.96[A]	7.718[B]	63.546[C]
Al	2.6989[A]	3.449[B]	26.9815384[C]
Si		4.1491[B]	28.085[C]
O		5.803[B]	15.999[C]
Zn		5.680[B]	65.38[C]
S		2.847[B]	32.06[C]
Mn		-3.73[B]	54.938[C]
Cr		3.635[B]	51.996[C]
Mo		6.715[B]	95.95[C]
Ni		10.3[B]	58.693[C]
^{58}Ni		14.4[B]	
Fe		9.45[B]	55.845[C]
B	2.37[A,I]		10.81[C]
^{10}B		(-0.1 - 1.066i)[B]	
^{11}B		6.65[B]	
Quartz glass (SiO_2)	2.203[D]		
Stainless steel (1.4435, 316L)	8[E]		
Solid deuterium (D_2)	0.2059[F]	6.671[B]	2.0141[G]
ZnS(Ag)	4.090[H]		
NiMo(85:15)	9.1[J]		

A [Dav05]
B [Sea92]
C [Pro+22]
D [Rae24]
E [Thy20]
F [MS36]
G [Wan+21]
H [ELJ21b]
I Density of amorphous boron was assumed
J Density estimated by weighted average of component densities

a steel manufacturer's material datasheet (see Tab. A.2). Additionally, only ranges were provided for the mass fractions of the individual stainless steel components. For the computation of the Fermi potential, mass fractions of $w_{\mathrm{Si}} = 1\,\%$, $w_{\mathrm{Mn}} = 1\,\%$, $w_{\mathrm{Cr}} = 18\,\%$, $w_{\mathrm{Mo}} = 2.75\,\%$, $w_{\mathrm{Ni}} = 13.75\,\%$ and $w_{\mathrm{Fe}} = 63.5\,\%$ were assumed. For the boron used in the detector conversion layer, the company[2] performing the coating provided a certificate specifying a ^{10}B atomic fraction of 98.68 %. The remaining fraction was assumed to be ^{11}B.

[2] CDT CASCADE Detector Technologies GmbH

Appendix B

τSPECT

B.1 Statistical uncertainty model

As explained in sec. 3.7.1, the probability to count k neutrons after the storage phase is given by

$$P(N = k) = P(N = k \mid \hat{N}_0 = n)\, P(\hat{N}_0 = n). \tag{B.1}$$

Substituting Eq. (3.3) and Eq. (3.6) yields

$$P(N = k) = \sum_{n=k}^{\infty} \binom{n}{k} p_{\mathrm{s}}^k (1 - p_s)^{n-k} e^{-N_0} \frac{N_0^n}{n!}. \tag{B.2}$$

Using $\binom{n}{k} = \frac{n!}{k!(n-k)!}$ and rewriting $N_0^n = N_0^k N_0^{n-k}$ this expression can be rearranged to

$$P(N = k) = \frac{N_0^k}{k!} e^{-N_0} p_{\mathrm{s}}^k \sum_{n=k}^{\infty} ((1 - p_{\mathrm{s}})N_0)^{n-k} \frac{1}{(n-k)!}. \tag{B.3}$$

Introducing the substitution $m = n - k$ leads to

$$P(N = k) = \frac{N_0^k}{k!} e^{-N_0} p_{\mathrm{s}}^k \sum_{m=0}^{\infty} ((1 - p_{\mathrm{s}})N_0)^{m} \frac{1}{m!}. \tag{B.4}$$

Recognizing the sum as the exponential series $e^x = \sum_{i=0}^{\infty} \frac{x^i}{i!}$ simplifies the expression to

$$P(N = k) = \frac{N_0^k}{k!} e^{-N_0} p_{\mathrm{s}}^k e^{(1-p_{\mathrm{s}})N_0} \tag{B.5}$$

This finally yields

$$P(N = k) = \frac{(p_{\mathrm{s}} N_0)^k}{k!} e^{-p_{\mathrm{s}} N_0}, \tag{B.6}$$

which corresponds to a Poisson distribution with mean $p_{\mathrm{s}} N_0$.

B.2 Storage time constant estimators

When performing n pairs of measurements consisting of a short storage time t_{s} and a long storage time t_{l} yielding counts $N_{\mathrm{s},i}$ and $N_{\mathrm{l},i}$ respectively, the following assumptions are made:

- The expected initial number of stored neutrons N_0, i is approximately equal for each pair of measurements.
- The measured counts follow Poisson statistics.

The Poisson distribution is given by

$$P(k) = \frac{\mu^k \exp(-\mu)}{k!}. \tag{B.7}$$

In the present case, $k = N_{\mathrm{s},i}, N_{\mathrm{l},i}$ and $\mu = N_{0,i} \exp(-\lambda_{\mathrm{st}} t)$, where $\lambda_{\mathrm{st}} = 1/\tau_{\mathrm{st}}$ and $t = t_{\mathrm{s}}, t_{\mathrm{l}}$. The corresponding log-likelihood function is given by

$$\begin{aligned} L(\{N_{0,i}\}, \lambda_{\mathrm{st}}) = \sum_{i=1}^{n} &(N_{\mathrm{s},i} + N_{\mathrm{l},i}) \ln(N_{0,i}) - (N_{\mathrm{s},i} + N_{\mathrm{l},i})\lambda_{\mathrm{st}} \\ &- N_{0,i}(\mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{s}}} + \mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{l}}}) - \ln(N_{\mathrm{s},i}!) - \ln(N_{\mathrm{l},i}!). \end{aligned} \tag{B.8}$$

Since only τ_{st} and thus λ_{st} is of interest, the parameters $\{N_{0,i}\}$ can be treated as nuisance parameters. These can be profiled out to obtain the profile likelihood [Paw01, pp. 61-62]

$$L_{\mathrm{p}}(\lambda_{\mathrm{st}}) = \sup_{\{N_{0,i}\}} L(\{N_{0,i}\}, \lambda_{\mathrm{st}}). \tag{B.9}$$

To compute the profile likelihood, the value of each $N_{0,i}$ that maximizes the likelihood as a function of λ_{st} must be determined. This is achieved by evaluating the partial derivatives

$$\frac{\partial L}{\partial N_{0,i}} = \frac{N_{\mathrm{s},i} + N_{\mathrm{l},i}}{N_{0,i}} - \mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{s}}} - \mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{l}}}. \tag{B.10}$$

The maximum with respect to $N_{0,i}$ is obtained by requiring $\frac{\partial L}{\partial N_{0,i}} \stackrel{!}{=} 0$, yielding

$$N_{0,i} = \frac{N_{\mathrm{s},i} + N_{\mathrm{l},i}}{\mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{s}}} + \mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{l}}}}. \tag{B.11}$$

Substituting eq. (B.11) into eq. (B.8) results in the profile log-likelihood

$$\begin{aligned} L_{\mathrm{p}}(\lambda_{\mathrm{st}}) = \sum_{i=1}^{n} &(N_{\mathrm{s},i} + N_{\mathrm{l},i}) \left(\ln(N_{\mathrm{s},i} + N_{\mathrm{l},i}) - \ln\left(\mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{s}}} + \mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{l}}}\right) + 1\right) \\ &- (N_{\mathrm{s},i} + N_{\mathrm{l},i})\lambda_{\mathrm{st}} - \ln(N_{\mathrm{s},i}!) - \ln(N_{\mathrm{l},i}!). \end{aligned} \tag{B.12}$$

The maximum likelihood estimator for λ_{st} is obtained by solving $\frac{\partial L_{\mathrm{p}}}{\partial \lambda_{\mathrm{st}}} \stackrel{!}{=} 0$, which yields

$$\frac{\partial L_{\mathrm{p}}}{\partial \lambda_{\mathrm{st}}} = \sum_{i=1}^{n} \frac{(N_{\mathrm{s},i} + N_{\mathrm{l},i})(t_{\mathrm{s}}\mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{s}}} + t_{\mathrm{l}}\mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{l}}})}{\mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{s}}} + \mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{l}}}} - N_{\mathrm{s},i} t_{\mathrm{s}} - N_{\mathrm{l},i} t_{\mathrm{l}}. \tag{B.13}$$

Rearranging terms gives

$$\sum_{i=1}^{n} (N_{\mathrm{s},i} + N_{\mathrm{l},i})(t_{\mathrm{s}}\mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{s}}} + t_{\mathrm{l}}\mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{l}}}) = \sum_{i=1}^{n} (N_{\mathrm{s},i} t_{\mathrm{s}} + N_{\mathrm{l},i} t_{\mathrm{l}})(\mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{s}}} + \mathrm{e}^{-\lambda_{\mathrm{st}} t_{\mathrm{l}}}). \tag{B.14}$$

Solving this equation for λ_{st} and substituting $\lambda_{\mathrm{st}} = 1/\tau_{st}$ finally yields

$$\tau_{\mathrm{st}} = \frac{t_{\mathrm{l}} - t_{\mathrm{s}}}{\ln\left(\frac{\sum_{i=1}^{n} N_{\mathrm{s},i}}{\sum_{i=1}^{n} N_{\mathrm{l},i}}\right)}. \tag{B.15}$$

B.3 Variance of random variable product

The variance of the product of two independent random variables follows directly from the definition of the variance and basic properties of expectations. For completeness, the derivation is explicitly shown here. In case of two independent random variables X and Y, the variance of the product is given by

$$\mathrm{Var}(XY) = \mathrm{E}(X^2Y^2) - \mathrm{E}(XY)^2, \tag{B.16}$$

where $E(\cdot)$ denotes the expected value. Using the independence of X and Y, this expression can be written as

$$\mathrm{Var}(XY) = \mathrm{E}(X^2)\mathrm{E}(Y^2) - \mathrm{E}(X)^2\mathrm{E}(Y)^2. \tag{B.17}$$

The expression can be further expanded by adding and subtracting $\mathrm{E}(X^2)\mathrm{E}(Y)^2$ yielding

$$\mathrm{Var}(XY) = \mathrm{E}(X^2)(\mathrm{E}(Y^2) - \mathrm{E}(Y)^2) + (\mathrm{E}(X^2) - \mathrm{E}(X)^2)\mathrm{E}(Y)^2. \tag{B.18}$$

Identifying the bracketed terms as variances of the respective variables leads to

$$\mathrm{Var}(XY) = \mathrm{E}(X^2)\mathrm{Var}(Y) + \mathrm{E}(Y)^2\mathrm{Var}(X). \tag{B.19}$$

For the case of normally distributed X and Poisson distributed Y, the relation the variance of the product becomes

$$\mathrm{Var}(XY) = \mathrm{E}(X)^2\mathrm{Var}(Y) + \mathrm{Var}(X)\left(\mathrm{E}(Y)^2 + \mathrm{Var}(Y)\right). \tag{B.20}$$

Appendix C

Light guide simulation

To import geometric shapes, the simulation relies on the binary stereolithography (STL) file standard[1]. The STL format stores a geometric object as a list of triangles, with all triangle parameters encoded as 32-bit floating point numbers. Most 3D design software packages include meshing algorithms that approximate a designed shape by triangular facets and export the result in the STL format. This approach allows different light guide geometries to be tested without modifying the simulation code itself. In addition to the STL file, the refractive index of each shape must be specified to model refraction correctly and a boolean flag must indicate whether a given shape represents a detector. Events, i.e. alpha particles absorbed in the scintillator, are modeled as point like light sources emitting photons uniformly in all directions. Photons are represented as light rays characterized by an origin, a direction vector and a polarization vector stored as 32-bit floating point numbers. Direction and polarization vector are generated using a 32-bit permuted congruential generator (PCG) random number generator [ONe14], which produces uniformly distributed numbers in the interval $[0, 1]$. Generating isotropic direction vectors is equivalent to sampling points uniformly on a unit sphere. This distribution can be expressed in terms of the polar angle $\theta \in [0, \pi]$ and an azimuthal angle $\phi \in [0, 2\pi]$ with probability density function

$$\rho(\phi, \theta) = \frac{\sin(\theta)}{4\pi}. \tag{C.1}$$

The normalization factor 4π ensures that the distribution integrates to unity, while the trigonometric factor accounts for the fact that the unit sphere diameter varies when varying θ. Consequently, values of θ near the equator of the sphere occur with higher probability than those near the poles. To map uniformly distributed random variables from the unit interval onto the spherical coordinate domains, the inverse cumulative distribution function (CDF^{-1}) of both θ and ϕ is required [Bon17, pp.78–79]. The CDFs are given by

$$\mathrm{CDF}(\phi) = \int_0^{\phi} \int_0^{\pi} \frac{\sin(\theta)}{4\pi} d\theta d\tilde{\phi} = \frac{1}{2\pi}\phi \tag{C.2}$$

and

$$\mathrm{CDF}(\theta) = \int_0^{\theta} \int_0^{2\pi} \frac{\sin(\tilde{\theta})}{4\pi} d\phi d\tilde{\theta} = \frac{1}{2}\left(1 - \cos(\theta)\right). \tag{C.3}$$

Assuming $u = \mathrm{CDF}(\phi)$ and $v = \mathrm{CDF}(\theta)$, where u and v are uniformly distributed random variables in the unit interval, the inverted CDFs become

$$\phi = 2\pi\, u \tag{C.4}$$

[1]The original *Stereolithography Interface Specification* by 3D Systems (1989) defines this format, but the official document is not publicly available. The binary structure of the format is e.g. documented on `https://en.wikipedia.org/wiki/STL_(file_format)`

and

$$\theta = \arccos(1 - 2v). \tag{C.5}$$

Sampling θ and ϕ in this way and transforming to Cartesian coordinates yields uniformly distributed direction vectors. To generate a uniformly distributed polarization vector, the direction vector is rotated by $\pi/2$ (adjusting θ accordingly) and then rotated around the direction vector by an angle sampled from a uniform distribution over $[0, 2\pi]$.

Subsequently, the photon trajectory is propagated. To determine the intersection point of a photon with a shape boundary, intersections between the ray and all triangles defined in the STL files are computed using the Möller-Trumbore algorithm [MT97]. The physically relevant intersection point is determined by finding the intersection point closest to the photon origin and this point becomes the new photon origin for the next iteration. To determine the new propagation direction, the refractive index of the adjacent medium must be identified. This is achieved by checking whether any intersection found for other shapes coincides with the new photon origin and, if so, assigning the refractive index of that shape. If no such match is found, the refractive index defaults to 1.0 corresponding to air or vacuum. To decide whether the photon is transmitted or reflected, the Fresnel equations

$$R_\perp = \left| \frac{n_1 \cos\theta_\mathrm{i} - n_2\sqrt{1 - \left(\frac{n_1}{n_2}\sin\theta_\mathrm{i}\right)^2}}{n_1 \cos\theta_\mathrm{i} + n_2\sqrt{1 - \left(\frac{n_1}{n_2}\sin\theta_\mathrm{i}\right)^2}} \right|^2 \tag{C.6}$$

and

$$R_\parallel = \left| \frac{n_1\sqrt{1 - \left(\frac{n_1}{n_2}\sin\theta_\mathrm{i}\right)^2} - n_2\cos\theta_\mathrm{i}}{n_1\sqrt{1 - \left(\frac{n_1}{n_2}\sin\theta_\mathrm{i}\right)^2} + n_2\cos\theta_\mathrm{i}} \right|^2 \tag{C.7}$$

are evaluated. Here, n_1 and n_2 denote the refractive indices of the incident and transmitted media, respectively, θ_i is the angle of incidence and $R_\perp$ and $R_\parallel$ are the reflection probabilities for perpendicular and parallel polarization. Parallel and perpendicular are defined relative to the plane of incidence, which is determined by the intersection point, the interface normal and the photon direction. The polarization state is determined by projecting the photon polarization vector $\vec{P}_\mathrm{ph}$ onto the normal vector of the plane of incidence $\vec{n}_\mathrm{i}$ via their scalar product. This yields the probability

$$p_\perp = \vec{P}_\mathrm{ph} \cdot \vec{n}_\mathrm{i}. \tag{C.8}$$

Subsequently, a random number r_p uniformly distributed in the unit interval is drawn to determine the polarization. If $r_\mathrm{p} < p_\perp$ holds, the photon is assigned perpendicular polarization. Otherwise, the polarization is parallel. Depending on the polarization state, either Eq. (C.6) or Eq. (C.7) is evaluated to obtain the reflection probability $R_{\perp,\parallel}$. This probability is compared to another random number r_R uniformly distributed in the unit interval. If $r_\mathrm{R} < R_{\perp,\parallel}$ holds, the photon is reflected. Otherwise, it is transmitted. In case of reflection, the new direction vector is

$$\vec{d'}_\mathrm{ph} = \vec{d}_\mathrm{ph} - 2\,(\vec{d}_\mathrm{ph} \cdot \vec{n}_\mathrm{s})\,\vec{n}_\mathrm{s}, \tag{C.9}$$

where $\vec{d}_{\text{ph}}$ is the previous direction vector and $\vec{n}_{\text{s}}$ is the interface normal. In case of transmission, the transmission angle θ_{t} is computed first using Snell's law. This yields

$$\theta_{\text{t}} = \arcsin\left(\frac{n_1}{n_2}\sin(\theta_{\text{i}})\right). \tag{C.10}$$

The new direction vector is then given by

$$\vec{d'}_{\text{ph}} = \frac{n_1}{n_2}\vec{d}_{\text{ph}} - \left(\frac{n_1}{n_2}\cos(\theta_{\text{i}}) - \sqrt{1-\sin^2(\theta_{\text{t}})}\right)\vec{n}_{\text{s}}. \tag{C.11}$$

Finally, the polarization vector is updated. For perpendicular polarization, the new polarization vector is given by

$$\vec{P'}_{\text{ph}} = \vec{n}_{\text{i}}. \tag{C.12}$$

Otherwise,

$$\vec{P'}_{\text{ph}} = \vec{n}_{\text{i}} \times \vec{d'}_{\text{ph}} \tag{C.13}$$

holds. If a photon is transmitted into a shape designated as a detector, it is considered absorbed and the photon counter of the corresponding detector is incremented. Otherwise, the propagation procedure is repeated with the updated origin, direction and polarization until the photon is either absorbed or no longer intersects any surface.

Because 32-bit floating point numbers provide only about seven significant decimal digits, numerical operations introduce finite precision errors. As a result, computed intersection points may not lie exactly on the respective surface. This can lead to unphysical repeated interactions with the same interface in subsequent iterations. To mitigate such effects, a tolerance of $\epsilon = 3 \times 10^{-4}$ mm is introduced. Numerical values agreeing within this margin are treated as equal. Compared to typical photon path lengths per iteration, which are on the order of centimeters to millimeters this tolerance is negligible. Additionally, the computed intersection point is projected onto the plane of the intersected triangle in each iteration to further reduce accumulated numerical error. The entire simulation was implemented in C++. A dedicated library[2] was used for all vector operations. Furthermore, event evaluation was parallelized[3] to reduce the required computation time.

[2] C++ library: Eigen3 (v. 3.4.0)
[3] C++ library: OpenMP (v. 4.5)

Appendix D

SiPM test setup

D.1 Temperature readout

The test setup incorporates two thermistor temperature sensors[1]. Electrically, each sensor behaves as an Ohmic resistor whose resistance exhibits a strong temperature dependence. Fig. D.1 shows the corresponding readout circuit which is implemented on the test setup PCB to read out the sensors. The temperature sensor with resistance R_T is is highlighted in orange. The analog-to-digital converter (ADC) module[2] provides four 16-bit input channels over a range from $V_\mathrm{ADC,\,min} = 0\,\mathrm{V}$ to $V_\mathrm{ADC,\,max} = 3.3\,\mathrm{V}$ and communicates with the Raspberry Pi of the test setup via the $\mathrm{I^2C}$ protocol. To convert the temperature dependent resistance into a measurable voltage, the sensor is biased with a dedicated voltage V_B and a bias resistor R_B. Together, bias resistor and temperature sensor form a voltage divider, from which the sensor voltage V_T is obtained. Following the manufacturer's recommendation [Tex23], a bias resistor of $10\,\mathrm{k\Omega}$ was selected. To bias the circuit, the 3.3 V output of the Raspberry Pi was used. The design temperature range was chosen to span from $0\,°\mathrm{C}$ to $50\,°\mathrm{C}$. Using the lookup table provided the manufacturer, the sensor resistance can be converted into a temperature and vice versa. The tabulated values were interpolated using cubic spline interpolation[3]. Over the selected temperature range, this results in a resistance interval from approximately $R_\mathrm{T,min} = 8.5\,\mathrm{k\Omega}$ to $R_\mathrm{T,max} = 11.6\,\mathrm{k\Omega}$. For the voltage across the temperature sensor,

$$V_\mathrm{T} = \frac{R_\mathrm{T}}{R_\mathrm{B} + R_\mathrm{T}} \tag{D.1}$$

holds. Within the selected temperature interval, this corresponds to a voltage range from $V_\mathrm{T,\,min} = 1.47\,\mathrm{V}$ to $V_\mathrm{T,\,max} = 1.77\,\mathrm{V}$. Although this lies well within the ADC input range, it utilizes only a fraction of the available dynamic range and therefore limits the achievable resolution. To improve this, the signal is amplified and shifted using an instrumentation amplifier[4] in combination with a second voltage divider (see Fig. D.1). The amplifier multiplies the difference between its positive (+) and negative (−) inputs by an amplification factor A, determined by an external gain resistor (not shown in the simplified circuit). The output voltage is therefore defined as

$$V_\mathrm{out} = A\,(V_+ - V_-). \tag{D.2}$$

To map the lower end of the design temperature range to the minimum ADC input voltage, the offset resistor R_0 in the second voltage divider should be equal or close

[1] Texas Instruments TMP61
[2] DFRobot I2C ADS1115 16-bit ADC module
[3] Python package scipy.interpolate.CubicSpline
[4] Analog Devices AD8226ARZ

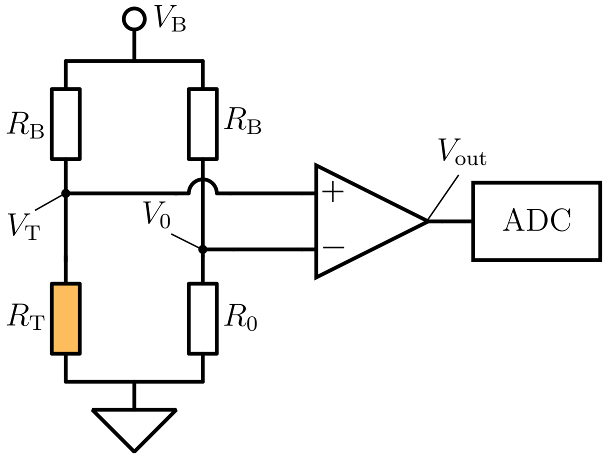

FIGURE D.1: Readout circuit for the temperature sensors used in the SiPM test setup. The orange colored resistor represents the temperature dependent resistance R_T of the thermistor.

to $R_{T,\,min}$. A value of $R_0 = 8.0\,\mathrm{k\Omega}$ was selected as the closest available nominal resistor value. To simultaneously map the upper end of the temperature range to $V_{ADC,\,max}$, A must satisfy $V_{ADC,\,max} = A\,(V_{T,\,max} - V_{T,\,min})$. Solving for A yields a design amplification factor of $A_d = 11.0$. For the gain resistor

$$R_G = \frac{49.4\,\mathrm{k\Omega}}{A - 1} \tag{D.3}$$

holds [Ana08]. Substituting A_d for A yields $R_G = 4.94\,\mathrm{k\Omega}$. Consequently, the nearest standard value of $R_G = 5.0\,\mathrm{k\Omega}$ was used. To reconstruct R_T from measured ADC value, the voltages indicated in Fig. D.1 must be considered. For the output voltage

$$V_{out} = A\,(V_T - V_0) \tag{D.4}$$

holds. The voltage across the offset resistor R_0 is given by

$$V_0 = \frac{R_0}{R_B + R_0}. \tag{D.5}$$

Substituting Eq. (D.1) and Eq. (D.5) into Eq. (D.4) and solving for R_T yields

$$R_T = \left[\frac{1}{R_B\,\frac{V_{out}}{A\,V_B} + \frac{R_0}{R_0+R_B}} - \frac{1}{R_B}\right]^{-1}. \tag{D.6}$$

The resulting resistance can then be converted into a temperature using the interpolated lookup table. To ensure that the actual circuit closely matches the design parameters, precision resistors with a tolerance of 0.1 % and a temperature coefficient of 25 ppm/°C were used. Additionally, the bias voltage is read out on a separate ADC channel and inserted into Eq. (D.6) together with the measured amplifier output voltages. After PCB assembly, the actual amplification factors of both amplifiers were measured using a multimeter and found to be 10.9 for both channels. This value was subsequently used in the readout software.

D.2 Breakdown voltage extraction procedure

This section describes the procedure used to extract the breakdown voltage of an SiPM from its IV-curve by determining the maximum of the second derivative of the logarithm of the current. To compute the first derivative, the measured data points were interpolated using a cubic spline[5]. The interpolation provides a continuous representation of the logarithmic IV-curve, from which the derivative can be obtained via the spline's piecewise derivative function provided by the interpolation package. An example of the resulting first derivative is shown in Fig. D.2. Due to statistical fluctuations between adjacent data points, the derivative exhibits an oscillatory behavior across the entire voltage range. These high frequency components obscure the underlying physical structure relevant for breakdown voltage determination. To suppress these oscillation, the derivative was resampled on an equidistant grid of 1000 points covering the bias voltage interval from 23.0 V to 26.5 V. A digital Butterworth low pass filter[6] was then applied to the resampled data. This type of filter was chosen because it provides a smooth, monotonic frequency response without passband ripple [Hol21, pp. 504-511]. To avoid phase distortions of the underlying derivative, zero-phase forward-backward filtering [Gus96] was employed as recommended in the package documentation. This procedure effectively doubles the filter order. Empirically, a 4th-order Butterworth filter (corresponding to 8th-order after forward-backward filtering) with a cutoff of $3.6\,\mathrm{V}^{-1}$ was found to sufficiently suppress the oscillations while preserving the relevant structure of the derivative. The filtered derivative is shown in Fig. D.2 together with the unfiltered result. To obtain the second derivative of the logarithmic IV-curve, the same spline based differentiation procedure as described above was applied to the filtered first derivative. The breakdown voltage was then determined by numerically computing the maximum of the second derivative using

[5] Python package: scipy.interpolate.CubicSpline (v. 1.15.2)

[6] Python packages: scipy.signal.butter and scipy.signal.sosfiltfilt (v. 1.15.2)

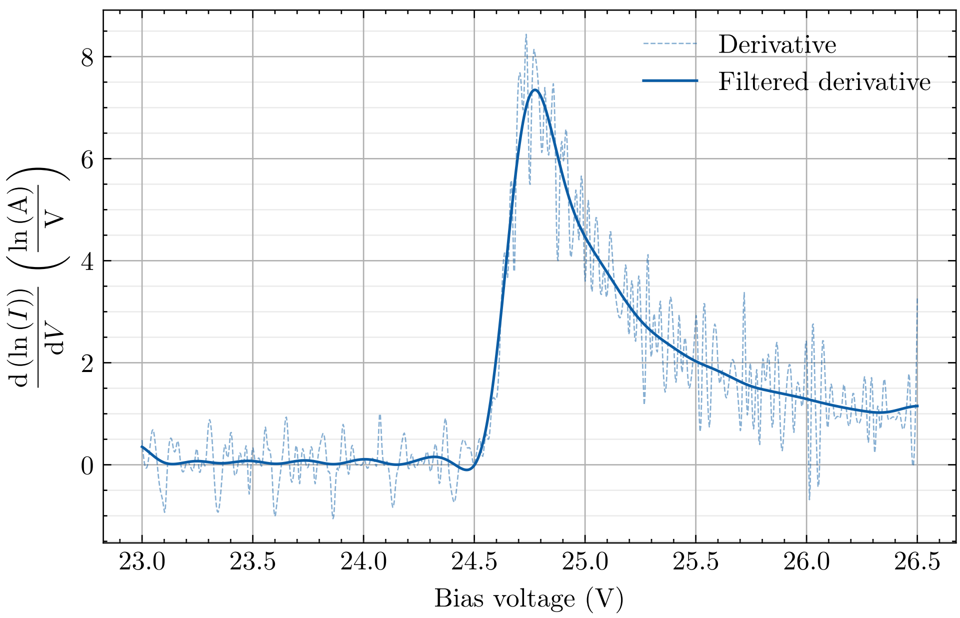

FIGURE D.2: First derivative of the logarithmic IV-curve of SiPM sample 0. The raw derivative is shown together with the result after digital low pass filtering using a Butterworth filter.

a standard optimization routine[7]. Filtering was applied to the first derivative rather than the raw data or the second derivative. Filtering the raw data did not sufficiently suppress high frequency components that are amplified during differentiation, while filtering the second derivative required strong smoothening due to noise amplification during differentiation. Therefore, filtering the first derivative provides a compromise that suppresses high frequency components before further amplification due to the second differentiation, while preserving the structure relevant for breakdown voltage determination.

[7]Python package: scipy.optimize.minimize (v. 1.15.2)

Appendix E

Detector characterization

E.1 Point like light source intensity distribution

This section derives the indefinite integral of the light intensity distribution of an isotropic point like light source, which is used to compute the amount of scintillation light observed by an SiPM in Section 5.4.1. For an isotropic point like source irradiating a plane at distance z_0, a uniform light intensity distribution over a hemisphere centered on the source can be assumed. In spherical coordinates this is given by

$$I(\theta, \phi) = I_0 \frac{\sin(\theta)}{2\pi}, \tag{E.1}$$

where $\theta \in [0, \pi/2]$ and $\phi \in [0, 2\pi]$ denote polar and azimuthal angle, respectively and I_0 is the total intensity. When mapping the spherically uniform distribution onto a plane described by Cartesian coordinates x and y at distance z_0 to the point source, the transformation from spherical to Cartesian coordinates is given by

$$x(\theta, \phi) = z_0 \tan(\theta) \cos(\phi) \tag{E.2}$$

and

$$y(\theta, \phi) = z_0 \tan(\theta) \sin(\phi). \tag{E.3}$$

To express $I(\theta, \phi)$ in terms of the Cartesian coordinates x and y, the absolute value of the determinant of the Jacobian matrix

$$J = \frac{\partial(\theta, \phi)}{\partial(x, y)} \tag{E.4}$$

must be evaluated. According to the inverse function theorem, this determinant is related to the Jacobian of the forward transformation by

$$\det(J) = \left| \det \left(\frac{\partial(\theta, \phi)}{\partial(x, y)} \right) \right| = \frac{1}{\left| \det \left(\frac{\partial(x,y)}{\partial(\theta,\phi)} \right) \right|}. \tag{E.5}$$

Computing the Jacobian of this transformation yields

$$\frac{\partial(x, y)}{\partial(\theta, \phi)} = \begin{bmatrix} z_0 \frac{\cos(\phi)}{\cos^2(\theta)} & -z_0 \tan(\theta) \sin(\phi) \\ z_0 \frac{\sin(\phi)}{\cos^2(\theta)} & z_0 \tan(\theta) \cos(\phi) \end{bmatrix}. \tag{E.6}$$

The determinant of this Jacobian is then given by

$$\det(J) = \frac{\cos^3(\theta)}{z_0^2 \sin(\theta)}. \tag{E.7}$$

Consequently, the light intensity distribution on the plane can be written as

$$I(x,y) = \frac{I_0 \cos^3(\theta(x,y))}{2\pi z_0^2}. \tag{E.8}$$

Using the geometric relation

$$\cos(\theta(x,y)) = \frac{z_0}{\sqrt{x^2+y^2+z_0^2}} \tag{E.9}$$

the planar intensity distribution finally becomes

$$I(x,y) = \frac{I_0 z_0}{2\pi(x^2+y^2+z_0^2)^{3/2}}. \tag{E.10}$$

In order to compute the light intensity observed by an SiPM, the indefinite integral of $I(x,y)$ must be evaluated. This integral is given by

$$H(x,y) = \iint \frac{I_0 z_0}{2\pi(x^2+y^2+z_0^2)^{3/2}} dxdy \tag{E.11}$$

and was solved with the aid of an online integral solver[1], which employs an open-source computer algebra system[2]. To solve the inner integral, the substitution

$$x = \sqrt{y^2+z_0^2}\tan(u) \tag{E.12}$$

can be applied, yielding

$$H(x,y) = \iint \frac{I_0 z_0 \cos(u)}{2\pi(y^2+z_0^2)} dudy. \tag{E.13}$$

Solving the inner integral and re-substituting leads to

$$H(x,y) = \int \frac{I_0 z_0}{2\pi(y^2+z_0^2)} \sin\left(\arctan\left(\frac{x}{\sqrt{y^2+z_0^2}}\right)\right) dy. \tag{E.14}$$

Using the trigonometric identity

$$\sin(\arctan(x)) = \frac{x}{\sqrt{1+x^2}} \tag{E.15}$$

the integrand simplifies to

$$H(x,y) = \int \frac{I_0 z_0 x}{2\pi(y^2+z_0^2)\sqrt{x^2+y^2+z_0^2}} dy. \tag{E.16}$$

To evaluate the remaining integral, the substitution

$$y = \sqrt{x^2+z_0^2}\tan(v) \tag{E.17}$$

is employed, yielding

$$H(x,y) = \int \frac{I_0 z_0}{2\pi} \frac{x}{\cos(v)((x^2+z_0^2)\tan(v)+z_0^2)} dv. \tag{E.18}$$

[1] https://www.integral-calculator.com/
[2] https://maxima.sourceforge.io/

To solve this, a second substitution

$$w = \sin(v) \tag{E.19}$$

can be introduced, which results in

$$H(x,y) = \int \frac{I_0 z_0}{2\pi} \frac{x}{(1-w^2)((x^2+z_0^2)\frac{w^2}{1-w^2}+z_0^2)} dw. \tag{E.20}$$

Simplifying the integrand yields

$$H(x,y) = \int \frac{I_0 z_0}{2\pi} \frac{x}{x^2 w^2 + z_0^2} dw, \tag{E.21}$$

which corresponds to a standard integral. The resulting expression for $H(x, y)$ is therefore

$$H(x,y) = \frac{I_0}{2\pi} \arctan\left(\frac{xw}{z_0}\right). \tag{E.22}$$

Re-substitution yields

$$H(x,y) = \frac{I_0}{2\pi} \arctan\left(\frac{x}{z_0} \sin\left(\arctan\left(\frac{y}{\sqrt{x^2+z_0^2}}\right)\right)\right). \tag{E.23}$$

Applying the identity given in Eq. (E.15) finally results in

$$H(x,y) = \frac{I_0}{2\pi} \arctan\left(\frac{xy}{z_0\sqrt{x^2+y^2+z_0^2}}\right). \tag{E.24}$$

To compute the light intensity observed by an SiPM for an event occurring at a position offset from the origin by (x_0, y_0), x and y in Eq. (E.24) can substituted by $(x - x_0)$ and $(y - y_0)$, respectively.

E.2 Gaussian light source intensity distribution

This section derives the indefinite integral of a two dimensional Gaussian light intensity distribution, which is used to compute the amount of scintillation light observed by an SiPM in Section 5.4.1. For a symmetric Gaussian light intensity distribution, the distribution is described by

$$I_{\mathrm{G}}(x,y) = \frac{I_0}{2\pi\sigma^2} \exp\left(-\frac{x^2+y^2}{2\sigma^2}\right), \tag{E.25}$$

where x and y denote the Cartesian coordinates in the SiPM array plane, I_0 is the total intensity and σ is the Gaussian width. To compute the intensity observed by an SiPM, the indefinite integral of I_{G} must be evaluated. This is given by

$$G(x,y) = \iint \frac{I_0}{2\pi\sigma^2} \exp\left(-\frac{x^2+y^2}{2\sigma^2}\right) dxdy. \tag{E.26}$$

Due to the symmetry of the Gaussian, the integrals can be separated, yielding

$$G(x,y) = \frac{I_0}{2\pi\sigma^2} \int \exp\left(-\frac{x^2}{2\sigma^2}\right) dx \int \exp\left(-\frac{y^2}{2\sigma^2}\right) dy. \tag{E.27}$$

The resulting expression for $G(x, y)$ is

$$G(x, y) = \frac{I_0}{8\sigma^2} \operatorname{erf}\left(\frac{x}{\sqrt{2}\sigma}\right) \operatorname{erf}\left(\frac{y}{\sqrt{2}\sigma}\right). \tag{E.28}$$

As in the case of the isotropic point-like source, the light intensity observed by an SiPM for an event offset from the origin by (x_0, y_0) can be obtained by replacing x and y in Eq. (E.28) by $(x - x_0)$ and $(y - y_0)$, respectively.

E.3 Distribution of bin widths

This section derives the statistical distribution governing the time required to record a fixed number of events at a constant count rate. This result is required to quantify the uncertainty of the bin width discussed in Section 5.6.2. When assuming a constant count rate R_c and a bin width t_bw, the event detection process can be modeled as a Poisson process with probability

$$P(k) = \frac{(R_\mathrm{c} t_\mathrm{bw})^k}{k!} \exp(-R_\mathrm{c} t_\mathrm{bw}), \tag{E.29}$$

where k denotes the number of events occurring within the time interval t_bw. In this framework, the probability for event number N to occur after a time t_bw is given by

$$p(t_N \geq t_\mathrm{bw}) = \sum_{k=0}^{N-1} \frac{(R_\mathrm{c} t_\mathrm{bw})^k}{k!} \exp(-R_\mathrm{c} t_\mathrm{bw}), \tag{E.30}$$

which is equivalent to the probability that fewer than N events to occur within t_bw. Consequently, the probability that event number N occurs within time t_bw is

$$p(t_N \leq t_\mathrm{bw}) = 1 - p(t_N \geq t_\mathrm{bw}), \tag{E.31}$$

which corresponds to the cumulative distribution function (CDF) of t_N. The probability density function (PDF) of t_N is obtained by differentiating with respect to t_bw, yielding

$$\begin{aligned} \mathrm{PDF}(t_\mathrm{bw}) = \frac{d}{dt_\mathrm{bw}} p(t_N \leq t_\mathrm{bw}) = &\sum_{k=0}^{N-1} \frac{(R_\mathrm{c} t_\mathrm{bw})^k}{k!} R_\mathrm{c} \exp(-R_\mathrm{c} t_\mathrm{bw}) \\ &- \sum_{k=1}^{N-1} \frac{R_\mathrm{c}^k t_\mathrm{bw}^{k-1}}{(k-1)!} \exp(-R_\mathrm{c} t_\mathrm{bw}). \end{aligned} \tag{E.32}$$

Introducing the substitution $j = k - 1$ in the second sum allows the expression to be rewritten as

$$\begin{aligned} \mathrm{PDF}(t_\mathrm{bw}) = &\sum_{k=0}^{N-1} \frac{(R_\mathrm{c} t_\mathrm{bw})^k}{k!} R_\mathrm{c} \exp(-R_\mathrm{c} t_\mathrm{bw}) \\ &- \sum_{j=0}^{N-2} \frac{(R_\mathrm{c} t_\mathrm{bw})^j}{j!} R_\mathrm{c} \exp(-R_\mathrm{c} t_\mathrm{bw}). \end{aligned} \tag{E.33}$$

All terms up to $(N-2)$ cancel, leaving

$$\mathrm{PDF}(t_{\mathrm{bw}}) = \frac{R_{\mathrm{c}}^{N} t_{\mathrm{bw}}^{N-1}}{(N-1)!} \exp(-R_{\mathrm{c}} t_{\mathrm{bw}}), \tag{E.34}$$

which can be identified as the Erlang distribution.

E.4 Additional modeled reconstructed event position distributions

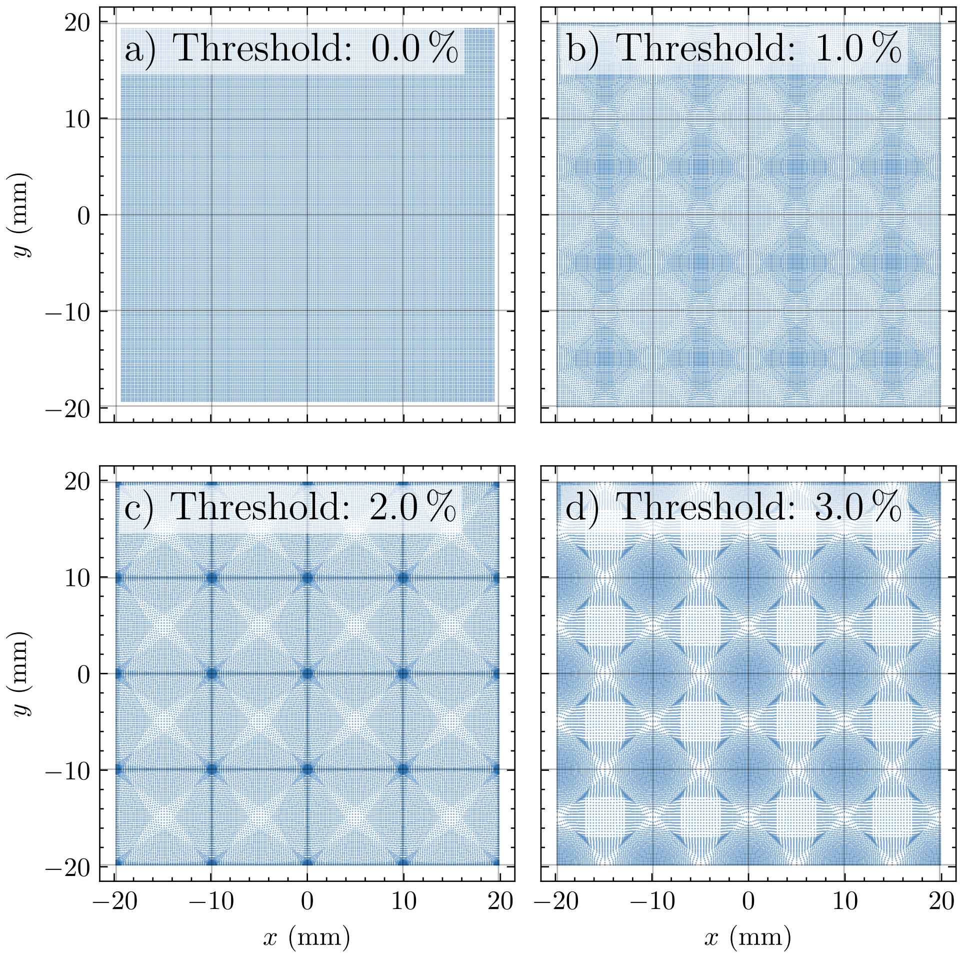

Figure E.1: Additional reconstructed event position distributions generated using the weighted average reconstruction approach and a two dimensional Gaussian signal light distribution with a width of $\sigma = 10\,\mathrm{mm}$. The reconstructed position of each event is represented by a dot. The centers of the SiPMs are located at the center of the depicted grid squares. Through sub-figures (a) to (d), the detection threshold of the SiPM channels is varied. For details on the model calculations, refer to Section 5.4.3.

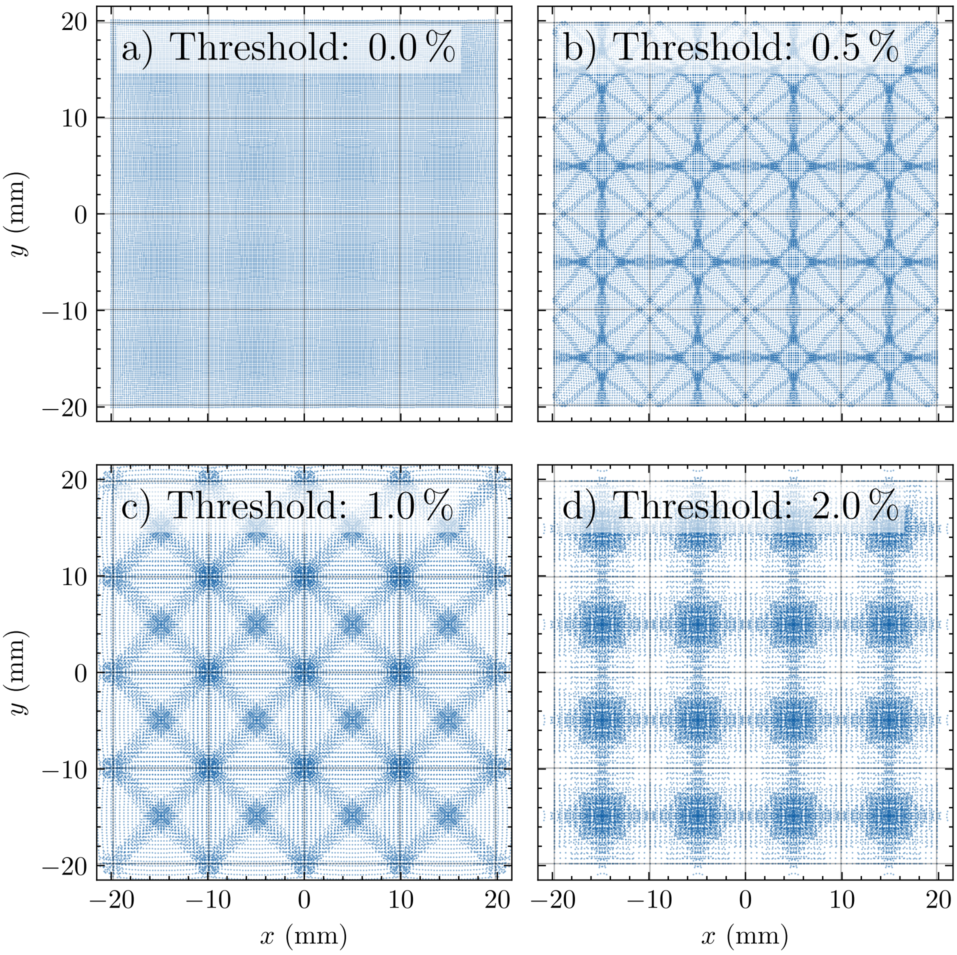

Figure E.2: Additional reconstructed event position distributions generated using the 2D-Gaussian reconstruction approach and a isotropic point light source model. The reconstructed position of each event is represented by a dot. The centers of the SiPMs are located at the center of the depicted grid squares. Through sub-figures (a) to (d), the detection threshold of the SiPM channels is varied. For details on the model calculations, refer to Section 5.4.3.

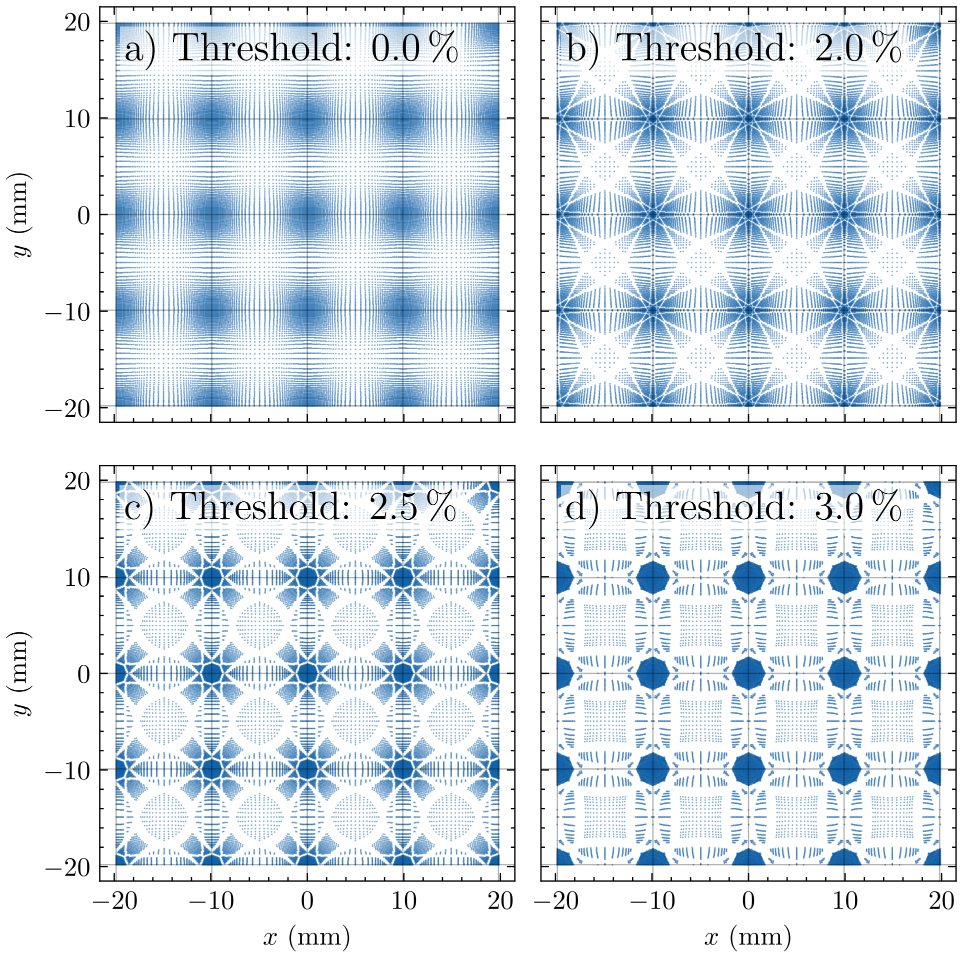

FIGURE E.3: Additional reconstructed event position distributions generated using the isotropic point like light source reconstruction approach and a two dimensional Gaussian signal light distribution with a width of $\sigma = 10\,\text{mm}$. The reconstructed position of each event is represented by a dot. The centers of the SiPMs are located at the center of the depicted grid squares. Through sub-figures (a) to (d), the detection threshold of the SiPM channels is varied. For details on the model calculations, refer to Section 5.4.3.

Appendix F

Gaussian approximation of probability distributions

As part of this thesis Poisson, binomial and Erlang distributions are implicitly approximated by a Gaussian distribution in order to compute and propagate uncertainties of measured quantities. This appendix provides an overview of the conditions under which these approximations are appropriate for each of these distributions. To assess the quality of the Gaussian approximation, the so-called total variation distance [DGL96, p. 32] [LPW08, pp. 47-48] is used. For continuous distributions, it is defined as

$$d_{\mathrm{TV}} = \frac{1}{2} \int |P(x) - G(x)| dx \tag{F.1}$$

and for discrete distributions as

$$d_{\mathrm{TV}} = \frac{1}{2} \sum_i |P(x_i) - G(x_i)|. \tag{F.2}$$

Here $P(x)$ denotes the respective probability distribution and $G(x)$ is its Gaussian approximation given by

$$G(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right), \tag{F.3}$$

where μ and σ^2 are set to the mean and the variance of $P(x)$, respectively. The factor $1/2$ in the definition of d_{TV} is a normalization factor such that $d_{\mathrm{TV}} \in [0, 1]$ holds. This can be understood by considering two normalized distributions with no overlap, which would otherwise yield $d_{\mathrm{TV}} = 2$. The total variation distance can be interpreted as follows. The probability for a value drawn from the true distribution to lie in an interval constructed using the Gaussian approximation differs by at most d_{TV} between true distribution and approximation. While this does not provide a strict bound on the error of the uncertainty estimate, the total variation distance nevertheless serves as a reasonable measure of the overall closeness between the true distribution and its Gaussian approximation. To enable numerical evaluation of the variation distance, a 3σ interval around the mean of the Gaussian approximation is considered. This interval contains about 99.73 % of all values drawn from a Gaussian distribution and therefore provides a good approximation of d_{TV} for the purpose of uncertainty estimation and propagation. The resulting truncated total variation distance is defined as

$$\hat{d}_{\mathrm{TV}} = \frac{1}{2w_3} \int_{\mu-3\sigma}^{\mu+3\sigma} |P(x) - G(x)| dx \tag{F.4}$$

for continuous distributions and

$$\hat{d}_{\mathrm{TV}} = \frac{1}{2w_3} \sum_{x_i \in C_{3\sigma}} |P(x_i) - G(x_i)| \tag{F.5}$$

for discrete distributions. Here, $w_3 \approx 0.9973$ is an additional normalization factor correcting for the truncation and $C_{3\sigma}$ is defined as the interval $[\mu - 3\sigma, \mu + 3\sigma]$. In practice, uncertainties are commonly reported as one sigma intervals with at most one or two digits of numerical precision. Accordingly, a threshold of $\hat{d}_{\mathrm{TV,th}} = 2\,\%$ is used throughout the thesis to assess whether a Gaussian approximation of the respective distribution is adequate for uncertainty estimation and propagation.

F.1 Poisson distribution

The Poisson distribution is defined as

$$P(k) = \frac{\lambda^k e^{-\lambda}}{k!} \tag{F.6}$$

with mean λ and variance λ, where k is typically referred to as the number of counted events. Fig. F.1 shows the truncated total variation distance for a Gaussian approximation of the Poisson distribution as a function of k.

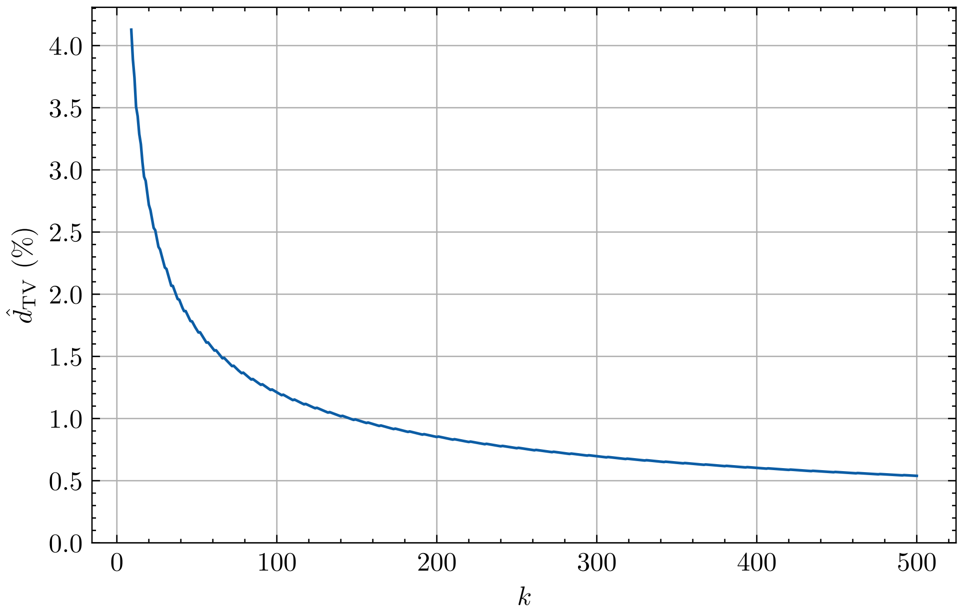

FIGURE F.1: Truncated total variation distance between a Poisson distribution and its Gaussian approximation as a function of the number of counted events k (see Eq. (F.6)).

F.2 Binomial distribution

The binomial distribution is given by

$$B(k) = \binom{N}{k} p^k (1-p)^{N-k}, \tag{F.7}$$

where N is the number of trials, p is the success probability and k is the number of successful trials. In general, the binomial distribution can be approximated by a Gaussian distribution for $N \gg 1$ and $0 \ll p \ll 1$. Fig. F.2 shows the truncated total variation distance for a Gaussian approximation of the binomial distribution as a function of p and for different values of N.

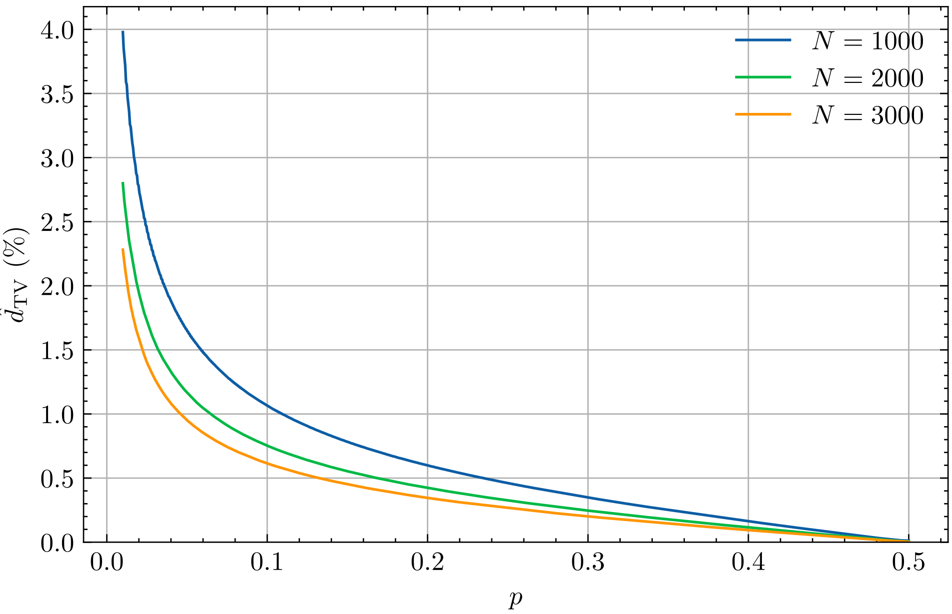

FIGURE F.2: Truncated total variation distance between a binomial distribution and its Gaussian approximation as a function of the success probability p for different number of trials N (see Eq. (F.7)).

F.3 Erlang distribution

The Erlang distribution is defined as

$$E(t) = \frac{\lambda^N t^{N-1} e^{-\lambda t}}{(N-1)!}, \tag{F.8}$$

where N is the number of events, λ is the event rate and t is the time required to observe N events at that rate. For $N \gg 1$, the Erlang distribution can be approximated by a Gaussian distribution. Fig. F.3 shows the truncated total variation distance for a Gaussian approximation of the Erlang distribution as a function of N.

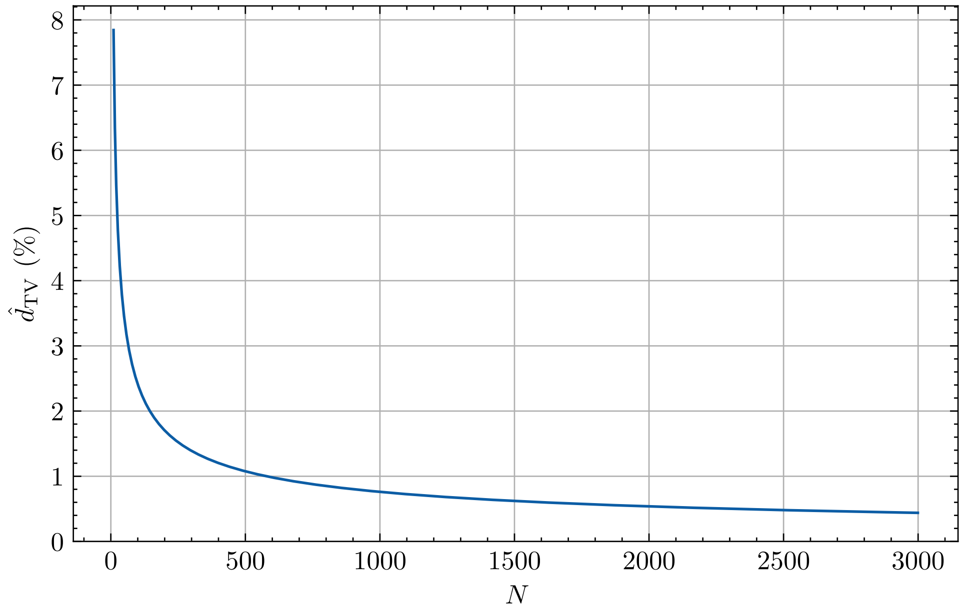

FIGURE F.3: Truncated total variation distance between an Erlang distribution and its Gaussian approximation as a function of the number of events N (see Eq. (F.8)).

Appendix G

Detector PCB pinout

Table G.1: Detector PCB pinout

Detector channel	ASIC channel	Quadrant	Cable line (out)	Cable line (in)
1	CH3	B	JB15	JB6
2	CH10	B	JB16	JB5
3	CH0	B	JB17	JB4
4	CH1	B	JB18	JB3
5	CH56	D	JD6	JD15
6	CH55	D	JD5	JD16
7	CH60	D	JD4	JD17
8	CH63	D	JD3	JD18
9	CH7	B	JB11	JB10
10	CH14	B	JB12	JB9
11	CH5	B	JB13	JB8
12	CH12	B	JB14	JB7
13	CH52	D	JD10	JD11
14	CH51	D	JD9	JD12
15	CH54	D	JD8	JD13
16	CH53	D	JD7	JD14
17	CH4	B	JB7	JB14
18	CH6	B	JB8	JB13
19	CH15	B	JB9	JB12
20	CH8	B	JB10	JB11
21	CH61	D	JD14	JD7
22	CH59	D	JD13	JD8
23	CH50	D	JD12	JD9
24	CH57	D	JD11	JD10
25	CH27	B	JB3	JB18
26	CH11	B	JB4	JB17
27	CH2	B	JB5	JB16
28	CH13	B	JB6	JB15
29	CH38	D	JD18	JD3
30	CH9	D	JD17	JD4
31	CH62	D	JD16	JD5
32	CH58	D	JD15	JD6
33	CH29	A	JA15	JA6
34	CH30	A	JA16	JA5
35	CH31	A	JA17	JA4
36	CH22	A	JA18	JA3

Detector channel	ASIC channel	Quadrant	Cable line (out)	Cable line (in)
37	CH35	C	JC6	JC15
38	CH36	C	JC5	JC16
39	CH33	C	JC4	JC17
40	CH34	C	JC3	JC18
41	CH25	A	JA11	JA10
42	CH24	A	JA12	JA9
43	CH26	A	JA13	JA8
44	CH28	A	JA14	JA7
45	CH40	C	JC10	JC11
46	CH41	C	JC9	JC12
47	CH39	C	JC8	JC13
48	CH37	C	JC7	JC14
49	CH21	A	JA7	JA14
50	CH20	A	JA8	JA13
51	CH23	A	JA9	JA12
52	CH22	A	JA10	JA11
53	CH46	C	JC14	JC7
54	CH44	C	JC13	JC8
55	CH42	C	JC12	JC9
56	CH43	C	JC11	JC10
57	CH16	A	JA3	JA18
58	CH18	A	JA4	JA17
59	CH17	A	JA5	JA16
60	CH19	A	JA6	JA15
61	CH49	C	JC18	JC3
62	CH48	C	JC17	JC4
63	CH47	C	JC16	JC5
64	CH45	C	JC15	JC6

Appendix H

Nutzung KI-Tools

Table H.1: Nutzung von KI-Tools

KI-Tool	Genutzt für	Warum	Wann bzw. Wo
Perplexity, ChatGPT	Literaturrecherche	Finden relevanter Literatur	Gesamte Arbeit
ChatGPT	Korrekturvorschläge Rechtschreibung und Grammatik	Verbesserung der Lesbarkeit	Gesamte Arbeit
ChatGPT	Vorschläge zur Verbesserung von Formulierungen	Verbesserung des Leseflusses	Gesamte Arbeit
ChatGPT	Vorschläge für Gliederung der Unterkapitel	Verbesserung der Lesbarkeit	Gesamte Arbeit
ChatGPT	Erstellen und Erklären von Python, C++ und Bash Codebeispielen	Verbesserung der Codequalität, Beschleunigung von Entwicklungszyklen	Gesamte Arbeit
ChatGPT	Zusammenfassen relevanter Software-dokumentation	Beschleunigung von Entwicklungszyklen	Gesamte Arbeit
ChatGPT	Vorschläge für Analysemethoden zur Auswertung von Daten	Methodenrecherche	Gesamte Arbeit
ChatGPT	Erklären von Zwischenschritten in mathematischen Herleitungen in Literatur	Besseres Verständnis	Gesamte Arbeit
ChatGPT	Konsistenzcheck ausgewählter mathematischer Herleitungen	Identifizieren von Fehlern	Gesamte Arbeit
ChatGPT	Erklären von Software Fehlermeldungen	Beschleunigte Problemlösung	Gesamte Arbeit

KI-Tool	Genutzt für	Warum	Wann bzw. Wo
DeepL	Überprüfen der Übersetzung der Englischen Zusammenfassung	Überprüfen der Übersetzung	Zusammenfassung
ChatGPT	Zusammenfassen von üblichen Zitierstandards für Datenblätter und Datenbanken	Finden einer angemessenen Zitierweise für Datenblätter und Datenbanken	Gesamte Arbeit
ChatGPT	Erklären von Ausschnitten elektrischer Schaltpläne	Besseres Verständnis	Gesamte Arbeit
ChatGPT	Erstellen von LaTeX Beispielen	Beschleunigung des Schreibprozesses	Gesamte Arbeit

Bibliography

[Abe+20] C. Abel et al. "Measurement of the Permanent Electric Dipole Moment of the Neutron". In: *Phys. Rev. Lett.* 124.8 (2020), p. 081803. DOI: 10.1103/PhysRevLett.124.081803.

[ABG48] R. A. Alpher, H. Bethe, and G. Gamow. "The Origin of Chemical Elements". In: *Phys. Rev.* 73.7 (1948), pp. 803–804. DOI: 10.1103/PhysRev.73.803.

[Ana08] Analog Devices. *Wide Supply Range, Rail-to-Rail Output Instrumentation Amplifier AD8226.* datasheet. Version Rev. D. AD8226. 2008.

[And+17] V. Andreev et al. "Ultracold Neutron Detector for Neutron Lifetime Measurements". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment.* Proceedings of the Vienna Conference on Instrumentation 2016 845 (2017), pp. 548–551. ISSN: 0168-9002. DOI: 10.1016/j.nima.2016.05.005.

[Ang+18] A. Anghel et al. "Solid Deuterium Surface Degradation at Ultracold Neutron Sources". In: *Eur. Phys. J. A* 54.9 (2018), p. 148. ISSN: 1434-601X. DOI: 10.1140/epja/i2018-12594-2.

[Arz+00] S. Arzumanov et al. "Neutron Life Time Value Measured by Storing Ultracold Neutrons with Detection of Inelastically Scattered Neutrons". In: *Physics Letters B* 483.1 (2000), pp. 15–22. ISSN: 0370-2693. DOI: 10.1016/S0370-2693(00)00579-7.

[Arz+12] S. S. Arzumanov et al. "Analysis and Correction of the Measurement of the Neutron Lifetime". In: *Jetp Lett.* 95.5 (2012), pp. 224–228. ISSN: 1090-6487. DOI: 10.1134/S0021364012050025.

[Arz+15] S. Arzumanov et al. "A Measurement of the Neutron Lifetime Using the Method of Storage of Ultracold Neutrons and Detection of Inelastically Up-Scattered Neutrons". In: *Physics Letters B* 745 (2015), pp. 79–89. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2015.04.021.

[Aul+24] J. Auler et al. "τSPECT: A Spin-Flip Loaded Magnetic Ultracold Neutron Trap for a Determination of the Neutron Lifetime". In: *J. Phys. G: Nucl. Part. Phys.* 51.11 (2024), p. 115103. ISSN: 0954-3899. DOI: 10.1088/1361-6471/ad727b.

[Aul+25] J. Auler et al. "Ultra-Cold Neutron Simulation Framework for the Free Neutron Lifetime Experiment τSPECT". In: *Eur. Phys. J. A* 61.10 (2025), p. 223. ISSN: 1434-601X. DOI: 10.1140/epja/s10050-025-01673-8.

[Ave+21] E. Aver et al. "Improving Helium Abundance Determinations with Leo P as a Case Study". In: *J. Cosmol. Astropart. Phys.* 2021.03 (2021), p. 027. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2021/03/027.

[Bae+08] S. Baeßler et al. “First Measurements with the Neutron Decay Spectrometer aSPECT”. In: *Eur. Phys. J. A* 38.1 (2008), pp. 17–26. ISSN: 1434-601X. DOI: 10.1140/epja/i2008-10660-0.

[Ban+16] G. Ban et al. “Ultracold Neutron Detection with 6Li-doped Glass Scintillators”. In: *Eur. Phys. J. A* 52.10 (2016), p. 326. ISSN: 1434-601X. DOI: 10.1140/epja/i2016-16326-4.

[Bas06] M. S. Basunia. “Nuclear Data Sheets for A = 237”. In: *Nuclear Data Sheets* 107.8 (2006), pp. 2323–2422. ISSN: 0090-3752. DOI: 10.1016/j.nds.2006.07.001.

[BB30] W. Bothe and H. Becker. “Künstliche Erregung von Kern-γ-Strahlen”. In: *Z. Physik* 66.5 (1930), pp. 289–306. ISSN: 0044-3328. DOI: 10.1007/BF01390908.

[BBB20] B. Belfatto, R. Beradze, and Z. Berezhiani. “The CKM Unitarity Problem: A Trace of New Physics at the TeV Scale?” In: *Eur. Phys. J. C* 80.2 (2020), p. 149. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-020-7691-6.

[Bec+20] M. Beck et al. “Improved determination of the β–$\overline{\nu}_e$ angular correlation coefficient a in free neutron decay with the aSPECT spectrometer”. In: *Phys. Rev. C* 101.5 (2020), p. 055506. DOI: 10.1103/PhysRevC.101.055506.

[Ber+05] M. Berger et al. *ESTAR, PSTAR, and ASTAR: Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions.* National Institute of Standards and Technology (NIST). 2005.

[Bis+20] G. Bison et al. “Neutron Optics of the PSI Ultracold-Neutron Source: Characterization and Simulation”. In: *Eur. Phys. J. A* 56.2 (2020), p. 33. ISSN: 1434-601X. DOI: 10.1140/epja/s10050-020-00027-w.

[Bis+22] G. Bison et al. “Ultracold Neutron Storage and Transport at the PSI UCN Source”. In: *Eur. Phys. J. A* 58.6 (2022), p. 103. ISSN: 1434-601X. DOI: 10.1140/epja/s10050-022-00747-1.

[Bon17] M. Bonamente. *Statistics and Analysis of Scientific Data.* Graduate Texts in Physics. Springer: New York, NY, 2017. ISBN: 978-1-4939-6570-0 978-1-4939-6572-4. DOI: 10.1007/978-1-4939-6572-4.

[Boo20] A. T. Boothroyd. *Principles of Neutron Scattering from Condensed Matter.* Oxford University Press: 2020. ISBN: 978-0-19-886231-4. DOI: 10.1093/oso/9780198862314.001.0001.

[Bro+18a] D. A. Brown et al. “ENDF/B-VIII.0: The 8th Major Release of the Nuclear Reaction Data Library with CIELO-project Cross Sections, New Standards and Thermal Scattering Data”. In: *Nuclear Data Sheets.* Special Issue on Nuclear Reaction Data 148 (2018), pp. 1–142. ISSN: 0090-3752. DOI: 10.1016/j.nds.2018.02.001.

[Bro+18b] M. A.-P. Brown et al. “New result for the neutron β-asymmetry parameter A_0 from UCNA”. In: *Phys. Rev. C* 97.3 (2018), p. 035505. DOI: 10.1103/PhysRevC.97.035505.

[BW11] M. Berglund and M. E. Wieser. “Isotopic compositions of the elements 2009 (IUPAC Technical Report)”. In: *Pure and Applied Chemistry* 83.2 (2011), pp. 397–410. ISSN: 1365-3075. DOI: 10.1351/PAC-REP-10-06-02.

[Byr+90] J. Byrne et al. “Measurement of the Neutron Lifetime by Counting Trapped Protons”. In: *Phys. Rev. Lett.* 65.3 (1990), pp. 289–292. DOI: 10.1103/PhysRevLett.65.289.

[Byr+96] J. Byrne et al. “A Revised Value for the Neutron Lifetime Measured Using a Penning Trap”. In: *EPL* 33.3 (1996), p. 187. ISSN: 0295-5075. DOI: 10.1209/epl/i1996-00319-x.

[Cab63] N. Cabibbo. “Unitary Symmetry and Leptonic Decays”. In: *Phys. Rev. Lett.* 10.12 (1963), pp. 531–533. DOI: 10.1103/PhysRevLett.10.531.

[Cat+14] P. W. Cattaneo et al. “Development of High Precision Timing Counter Based on Plastic Scintillator with SiPM Readout”. In: *IEEE Transactions on Nuclear Science* 61.5 (2014), pp. 2657–2666. ISSN: 1558-1578. DOI: 10.1109/TNS.2014.2347576.

[Cay+25] J. Caylor et al. “Detection of Molecular Hydrogen in a Neutron Beam Lifetime Experiment”. In: *Phys. Rev. C* 112.6 (2025), p. 065501. DOI: 10.1103/nr3b-3dtl.

[CDT18] CDT GmbH. *CASCADE-U*. datasheet. 2018.

[Cha32] J. Chadwick. “Possible Existence of a Neutron”. In: *Nature* 129.3252 (1932), pp. 312–312. ISSN: 1476-4687. DOI: 10.1038/129312a0.

[Chr+72] C. J. Christensen et al. “Free-Neutron Beta-Decay Half-Life”. In: *Phys. Rev. D* 5.7 (1972), pp. 1628–1640. DOI: 10.1103/PhysRevD.5.1628.

[Cir+23a] V. Cirigliano et al. “Effective Field Theory for Radiative Corrections to Charged-Current Processes: Vector Coupling”. In: *Phys. Rev. D* 108.5 (2023), p. 053003. DOI: 10.1103/PhysRevD.108.053003.

[Cir+23b] V. Cirigliano et al. “Scrutinizing CKM Unitarity with a New Measurement of the $K\mu3/K\mu2$ Branching Fraction”. In: *Physics Letters B* 838 (2023), p. 137748. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2023.137748.

[CJ32] I. Curie and F. Joliot. “Émission de Protons de Grande Vitesse Par Les Substances Hydrogénées Sous l’influence Des Rayons γ Très Pénétrants”. In: *C. R. Acad. Sci. Paris* 194 (1932), pp. 273–275.

[Cle+19] B. Clement et al. “Boron-10 Conversion Layer for Ultra-Cold Neutron Detection”. In: *J. Inst.* 14.09 (2019), P09003. ISSN: 1748-0221. DOI: 10.1088/1748-0221/14/09/P09003.

[Clé+22] B. Clément et al. “Spatial Resolution Determination of a Position Sensitive Ultra-Cold Neutron Detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1040 (2022), p. 167212. ISSN: 0168-9002. DOI: 10.1016/j.nima.2022.167212.

[CMS04] A. Czarnecki, W. J. Marciano, and A. Sirlin. “Precision Measurements and CKM Unitarity”. In: *Phys. Rev. D* 70.9 (2004), p. 093006. DOI: 10.1103/PhysRevD.70.093006.

[Col+11] G. Collazuol et al. “Studies of Silicon Photomultipliers at Cryogenic Temperatures”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. VCI 2010 628.1 (2011), pp. 389–392. ISSN: 0168-9002. DOI: 10.1016/j.nima.2010.07.008.

[Cry] Cryogenic LTD. *Cryogen-Free Magnet System Cryogenic Job No. 1992.* datasheet.

[Cud+22] C. Cude-Woods et al. "Fill and Dump Measurement of the Neutron Lifetime Using an Asymmetric Magneto-Gravitational Trap". In: *Phys. Rev. C* 106.6 (2022), p. 065506. DOI: 10.1103/PhysRevC.106.065506.

[Cyb+16] R. H. Cyburt et al. "Big Bang Nucleosynthesis: Present Status". In: *Rev. Mod. Phys.* 88.1 (2016), p. 015004. DOI: 10.1103/RevModPhys.88.015004.

[Dav05] David R. Lide. *CRC Handbook of Chemistry and Physics, Internet Version 2005.* CRC Press: Boca Raton, FL, 2005.

[DGL96] L. Devroye, L. Györfi, and G. Lugosi. *A Probabilistic Theory of Pattern Recognition.* Red. by I. Karatzas and M. Yor. Vol. 31. Stochastic Modelling and Applied Probability. Springer: New York, NY, 1996. ISBN: 978-1-4612-6877-2 978-1-4612-0711-5. DOI: 10.1007/978-1-4612-0711-5.

[DM21] D. Dubbers and B. Märkisch. "Precise Measurements of the Decay of Free Neutrons". In: *Annual Review of Nuclear and Particle Science* 71 (Volume 71, 2021 2021), pp. 139–163. ISSN: 0163-8998, 1545-4134. DOI: 10.1146/annurev-nucl-102419-043156.

[Doo25] C. B. Doorenbos. "Improving Ultracold Neutron Statistics for the n2EDM Experiment". PhD thesis. ETH Zürich, 2025.

[DV98] D. J. Daley and D. Vere-Jones. "Basic Properties of the Poisson Process". In: *An Introduction to the Theory of Point Processes*. Ed. by D. J. Daley and D. Vere-Jones. Springer: New York, NY, 1998, pp. 18–37. ISBN: 978-1-4757-2001-3. DOI: 10.1007/978-1-4757-2001-3_2.

[EK00] K. Eberhardt and A. Kronenberg. "The research reactor TRIGA Mainz - A neutron source for versatile applications in research and education / Der Forschungsreaktor TRIGA Mainz — Eine Neutronenquelle für vielseitige Anwendungen in Forschung und Lehre". In: *Kerntechnik* 65.5–6 (2000), pp. 269–274. ISSN: 2195-8580. DOI: 10.1515/kern-2000-655-621.

[ELJ21a] ELJEN TECHNOLOGY. *Alpha Detection EJ-440, EJ.442.* datasheet. 2021.

[ELJ21b] ELJEN TECHNOLOGY. *ZnS:Ag PHOSPHOR POWDER EJ-600.* datasheet. 2021.

[Eng25] M. C. Engler. "Normalization and Filling Optimization of the τSPECT Experiment". PhD thesis. Johannes Gutenberg-Universität Mainz, 2025.

[Erm22] V. Ermuth. "Development of a Compact NMR Magnetometer for the τSPECT Experiment". Bachelor thesis. Johannes Gutenberg-Universität Mainz, 2022.

[Esp+99] S. Esposito et al. "Precision rates for nucleon weak interactions in primordial nucleosynthesis and ^{4}He abundance". In: *Nuclear Physics B* 540.1 (1999), pp. 3–36. ISSN: 0550-3213. DOI: 10.1016/S0550-3213(98)00757-3.

[Ezh+18] V. F. Ezhov et al. "Measurement of the Neutron Lifetime with Ultracold Neutrons Stored in a Magneto-Gravitational Trap". In: *Jetp Lett.* 107.11 (2018), pp. 671–675. ISSN: 1090-6487. DOI: 10.1134/S0021364018110024.

[FG18] B. Fornal and B. Grinstein. “Dark Matter Interpretation of the Neutron Decay Anomaly”. In: *Phys. Rev. Lett.* 120.19 (2018), p. 191801. DOI: 10.1103/PhysRevLett.120.191801.

[Fie+20] B. D. Fields et al. “Big-Bang Nucleosynthesis after Planck”. In: *J. Cosmol. Astropart. Phys.* 2020.03 (2020), p. 010. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2020/03/010.

[Fuw+24] Y. Fuwa et al. *Improved Measurements of Neutron Lifetime with Cold Neutron Beam at J-PARC*. 2024. DOI: 10.48550/arXiv.2412.19519.

[FZ46] E. Fermi and W. H. Zinn. *Reflection of Neutrons on Mirrors*. MDDC-56. Oak Ridge, Tennessee: U.S. Atomic Energy Commission, 1946.

[GB83] R. Golub and K. Böning. “New Type of Low Temperature Source of Ultra-cold Neutrons and Production of Continous Beams of UCN”. In: *Z. Physik B - Condensed Matter* 51.2 (1983), pp. 95–98. ISSN: 1431-584X. DOI: 10.1007/BF01308763.

[GH20] S. Gundacker and A. Heering. “The Silicon Photomultiplier: Fundamentals and Applications of a Modern Solid-State Photon Detector”. In: *Phys. Med. Biol.* 65.17 (2020), 17TR01. ISSN: 0031-9155. DOI: 10.1088/1361-6560/ab7b2d.

[Gmb22] G. S. GmbH. *KONTAMINATIONSMONITOR CoMo-170/CoMo-170 F*. datasheet. 2022.

[Gon+21] F. M. Gonzalez et al. “Improved Neutron Lifetime Measurement with UCNτ”. In: *Phys. Rev. Lett.* 127.16 (2021), p. 162501. DOI: 10.1103/PhysRevLett.127.162501.

[GP77] R. Golub and J. M. Pendlebury. “The Interaction of Ultra-Cold Neutrons (UCN) with Liquid Helium and a Superthermal UCN Source”. In: *Physics Letters A* 62.5 (1977), pp. 337–339. ISSN: 0375-9601. DOI: 10.1016/0375-9601(77)90434-0.

[GRL91] R. Golub, D. Richardson, and S. K. Lamoreaux. *Ultra-Cold Neutrons*. CRC Press: 1991. ISBN: 978-0-7503-0115-2.

[GS23] M. Gorchtein and C.-Y. Seng. “The Standard Model Theory of Neutron Beta Decay”. In: *Universe* 9.9 (2023), p. 422. ISSN: 2218-1997. DOI: 10.3390/universe9090422.

[Gua11] F. A. Guardia. “Characterization of Electromagnetic Fields in the aSPECT Spectrometer and Reduction of Systematic Errors”. PhD thesis. Johannes Gutenberg-Universität Mainz, 2011.

[Gus96] F. Gustafsson. “Determining the Initial States in Forward-Backward Filtering”. In: *IEEE Transactions on Signal Processing* 44.4 (1996), pp. 988–992. ISSN: 1941-0476. DOI: 10.1109/78.492552.

[Haa+17] R. Haas et al. “Development and Characterization of a Drop-on-Demand Inkjet Printing System for Nuclear Target Fabrication”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 874 (2017), pp. 43–49. ISSN: 0168-9002. DOI: 10.1016/j.nima.2017.08.027.

[Hal80] K. Halbach. “Design of Permanent Multipole Magnets with Oriented Rare Earth Cobalt Material”. In: *Nuclear Instruments and Methods* 169.1 (1980), pp. 1–10. ISSN: 0029-554X. DOI: 10.1016/0029-554X(80)90094-4.

[Han+25] O. Hanski et al. *Performance of Silicon Photomultipliers at Low Temperatures.* 2025. DOI: 10.48550/arXiv.2501.03812.

[Has+21] M. T. Hassan et al. "Measurement of the Neutron Decay Electron-Antineutrino Angular Correlation by the aCORN Experiment". In: *Phys. Rev. C* 103.4 (2021), p. 045502. DOI: 10.1103/PhysRevC.103.045502.

[Hay24] L. Hayen. "Opportunities and Open Questions in Modern Beta Decay". In: *Annual Review of Nuclear and Particle Science* 74 (Volume 74, 2024 2024), pp. 497–528. ISSN: 0163-8998, 1545-4134. DOI: 10.1146/annurev-nucl-121423-100730.

[Hol+12] A. T. Holley et al. "A High-Field Adiabatic Fast Passage Ultracold Neutron Spin Flipper for the UCNA Experiment". In: *Review of Scientific Instruments* 83.7 (2012), p. 073505. ISSN: 0034-6748. DOI: 10.1063/1.4732822.

[Hol21] T. Holton. *Digital Signal Processing: Principles and Applications.* Cambridge University Press: Cambridge, 2021. ISBN: 978-1-108-29005-0. DOI: 10.1017/9781108290050.

[Hol23] J. T. Holzbeck. "Characterization of 3D Printed Optical Waveguides for a Neutron Detector". Bachelor thesis. Johannes Gutenberg-Universität Mainz, 2023.

[Hoo+19] S. F. Hoogerheide et al. "Progress on the BL2 Beam Measurement of the Neutron Lifetime". In: *EPJ Web Conf.* 219 (2019), p. 03002. ISSN: 2100-014X. DOI: 10.1051/epjconf/201921903002.

[HT20] J. C. Hardy and I. S. Towner. "Superallowed $0^+ \to 0^+$ nuclear β decays: 2020 critical survey, with implications for V_{ud} and CKM unitarity". In: *Phys. Rev. C* 102.4 (2020), p. 045501. DOI: 10.1103/PhysRevC.102.045501.

[ITG14] Y. I. Izotov, T. X. Thuan, and N. G. Guseva. "A New Determination of the Primordial He Abundance Using the He I λ10830 Å Emission Line: Cosmological Implications". In: *Mon Not R Astron Soc* 445.1 (2014), pp. 778–793. ISSN: 0035-8711. DOI: 10.1093/mnras/stu1771.

[Jak+09] J. Jakubek et al. "A Coated Pixel Device TimePix with Micron Spatial Resolution for UCN Detection". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 600.3 (2009), pp. 651–656. ISSN: 0168-9002. DOI: 10.1016/j.nima.2008.12.078.

[Jen+13] T. Jenke et al. "Ultracold Neutron Detectors Based on 10B Converters Used in the qBounce Experiments". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment.* Vienna Conference on Instrumentation 2013 732 (2013), pp. 1–8. ISSN: 0168-9002. DOI: 10.1016/j.nima.2013.06.024.

[Jha+13] A. K. Jha et al. "Simulating Silicon Photomultiplier Response to Scintillation Light". In: *IEEE Transactions on Nuclear Science* 60.1 (2013), pp. 336–351. ISSN: 1558-1578. DOI: 10.1109/TNS.2012.2234135.

[Kah20] J. Kahlenberg. "First Full-Magnetic Storage of Ultracold Neutrons in the τSPECT Experiment for Measuring the Neutron Lifetime". PhD thesis. Johannes Gutenberg-Universität Mainz, 2020.

[Kaw+10] S. Kawasaki et al. “Development of a Pixel Detector for Ultra-Cold Neutrons”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 615.1 (2010), pp. 42–47. ISSN: 0168-9002. DOI: 10.1016/j.nima.2010.01.011.

[KBB22] Z. Kolka, P. Barcik, and V. Biolkova. “Statistical Analysis of Silicon Photomultiplier Output Signals”. In: *Sensors* 22.23 (2022), p. 9134. ISSN: 1424-8220. DOI: 10.3390/s22239134.

[Kha+89] A. G. Kharitonov et al. “Preliminary Results of Neutron Lifetime Measurements with Gravitational UCN Trap”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 284.1 (1989), pp. 98–100. ISSN: 0168-9002. DOI: 10.1016/0168-9002(89)90257-X.

[Kla19] R. Klanner. “Characterisation of SiPMs”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. Silicon Photomultipliers: Technology, Characterisation and Applications 926 (2019), pp. 36–56. ISSN: 0168-9002. DOI: 10.1016/j.nima.2018.11.083.

[Klo+19] M. Klopf et al. “Constraints on the Dark Matter Interpretation $n \to \chi + e^+e^-$ of the Neutron Decay Anomaly with the PERKEO II Experiment”. In: *Phys. Rev. Lett.* 122.22 (2019), p. 222503. DOI: 10.1103/PhysRevLett.122.222503.

[KM73] M. Kobayashi and T. Maskawa. “CP-Violation in the Renormalizable Theory of Weak Interaction”. In: *Prog Theor Phys* 49.2 (1973), pp. 652–657. ISSN: 0033-068X. DOI: 10.1143/PTP.49.652.

[Kra88] K. S. Krane. *Introductory Nuclear Physics*. John Wiley & Sons: New York, 1988.

[Kri+26] M. Krivos et al. *Detection of Ultracold Neutrons with Powdered Scintillator Screens*. 2026. DOI: 10.48550/arXiv.2509.04332.

[Kri25] M. Krivos. “UCNProBe a Beam-Type Experiment Using Ultracold Neutrons”. Conference talk. Satellite Workshop "Neutron Lifetime Puzzle - Measurements, Systematics, and Theoretical Implications at PSI2025" (Villigen, Switzerland). 2025.

[Kuk+21] K. Kuk et al. “Projection Imaging with Ultracold Neutrons”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1003 (2021), p. 165306. ISSN: 0168-9002. DOI: 10.1016/j.nima.2021.165306.

[Leo94] W. R. Leo. *Techniques for Nuclear and Particle Physics Experiments: A How-to Approach*. Springer: Berlin, Heidelberg, 1994. ISBN: 978-3-540-57280-0 978-3-642-57920-2. DOI: 10.1007/978-3-642-57920-2.

[Lin+23] S. Lin et al. “Demonstration of Sub-Micron UCN Position Resolution Using Room-Temperature CMOS Sensor”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1057 (2023), p. 168769. ISSN: 0168-9002. DOI: 10.1016/j.nima.2023.168769.

[LPW08] D. Levin, Y. Peres, and E. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society: Providence, Rhode Island, 2008. ISBN: 978-0-8218-4739-8 978-1-4704-1204-3. DOI: 10.1090/mbk/058.

[LT84] V. I. Luschikov and Yu. V. Taran. "On the Calculation of the Neutron Adiabatic Spin-Flipper". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 228.1 (1984), pp. 159–160. ISSN: 0168-9002. DOI: 10.1016/0168-9002(84)90025-1.

[Mam+89] W. Mampe et al. "Neutron Lifetime Measured with Stored Ultracold Neutrons". In: *Phys. Rev. Lett.* 63.6 (1989), pp. 593–596. DOI: 10.1103/PhysRevLett.63.593.

[Man+23] S. E. Mann et al. "Timing Properties of Radioluminescence in Nanoparticle ZnS:Ag Scintillators". In: *Optical Materials: X* 17 (2023), p. 100226. ISSN: 2590-1478. DOI: 10.1016/j.omx.2022.100226.

[Mar+13] D. Marano et al. "Improved SPICE Electrical Model of Silicon Photomultipliers". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 726 (2013), pp. 1–7. ISSN: 0168-9002. DOI: 10.1016/j.nima.2013.05.127.

[Mär+19] B. Märkisch et al. "Measurement of the Weak Axial-Vector Coupling Constant in the Decay of Free Neutrons Using a Pulsed Cold Neutron Beam". In: *Phys. Rev. Lett.* 122.24 (2019), p. 242501. DOI: 10.1103/PhysRevLett.122.242501.

[MKS05] G. J. Mathews, T. Kajino, and T. Shima. "Big Bang Nucleosynthesis with a New Neutron Lifetime". In: *Phys. Rev. D* 71.2 (2005), p. 021302. DOI: 10.1103/PhysRevD.71.021302.

[Moh+25] P. J. Mohr et al. "CODATA Recommended Values of the Fundamental Physical Constants: 2022". In: *Rev. Mod. Phys.* 97.2 (2025), p. 025002. DOI: 10.1103/RevModPhys.97.025002.

[Mor+17] C. L. Morris et al. "A New Method for Measuring the Neutron Lifetime Using an in Situ Neutron Detector". In: *Review of Scientific Instruments* 88.5 (2017), p. 053508. ISSN: 0034-6748. DOI: 10.1063/1.4983578.

[MS36] H. D. Megaw and F. Simon. "Density and Compressibility of Solid Hydrogen and Deuterium at 4·2° K". In: *Nature* 138.3484 (1936), pp. 244–244. ISSN: 1476-4687. DOI: 10.1038/138244a0.

[MT97] T. Möller and B. Trumbore. "Fast, Minimum Storage Ray-Triangle Intersection". In: *Journal of Graphics Tools* 2.1 (1997), pp. 21–28. ISSN: 1086-7651. DOI: 10.1080/10867651.1997.10487468.

[MuL+11] MuLan Collaboration et al. "Measurement of the Positive Muon Lifetime and Determination of the Fermi Constant to Part-per-Million Precision". In: *Phys. Rev. Lett.* 106.4 (2011), p. 041803. DOI: 10.1103/PhysRevLett.106.041803.

[Mus+25] R. Musedinovic et al. "Measurement of the Free Neutron Lifetime in a Magneto-Gravitational Trap with in Situ Detection". In: *Phys. Rev. C* 111.4 (2025), p. 045501. DOI: 10.1103/PhysRevC.111.045501.

[Mut+22] N. Muto et al. "A Novel Nuclear Emulsion Detector for Measurement of Quantum States of Ultracold Neutrons in the Earth's Gravitational Field". In: *J. Inst.* 17.07 (2022), P07014. ISSN: 1748-0221. DOI: 10.1088/1748-0221/17/07/P07014.

[Nag+17] F. Nagy et al. "A Model Based DC Analysis of SiPM Breakdown Voltages". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 849 (2017), pp. 55–59. ISSN: 0168-9002. DOI: `10.1016/j.nima.2017.01.002`.

[Nag+18] N. Naganawa et al. "A Cold/Ultracold Neutron Detector Using Fine-Grained Nuclear Emulsion with Spatial Resolution Less than 100 Nm". In: *Eur. Phys. J. C* 78.11 (2018), p. 959. ISSN: 1434-6052. DOI: `10.1140/epjc/s10052-018-6395-7`.

[Nav+24] S. Navas et al. "Review of Particle Physics". In: *Phys. Rev. D* 110.3 (2024), p. 030001. DOI: `10.1103/PhysRevD.110.030001`.

[Nes+02] V. V. Nesvizhevsky et al. "Quantum States of Neutrons in the Earth's Gravitational Field". In: *Nature* 415.6869 (2002), pp. 297–299. ISSN: 1476-4687. DOI: `10.1038/415297a`.

[Nes+05] V. V. Nesvizhevsky et al. "Study of the Neutron Quantum States in the Gravity Field". In: *Eur. Phys. J. C* 40.4 (2005), pp. 479–491. ISSN: 1434-6052. DOI: `10.1140/epjc/s2005-02135-y`.

[Nic+05] J. S. Nico et al. "Measurement of the Neutron Lifetime by Counting Trapped Protons in a Cold Neutron Beam". In: *Phys. Rev. C* 71.5 (2005), p. 055502. DOI: `10.1103/PhysRevC.71.055502`.

[Nic09] J. S. Nico. "Neutron Beta Decay". In: *J. Phys. G: Nucl. Part. Phys.* 36.10 (2009), p. 104001. ISSN: 0954-3899. DOI: `10.1088/0954-3899/36/10/104001`.

[Nor] Norland Products Inc. *Norland Optical Adhesive 68*. datasheet. Version REV.0.

[NV17] V. Nesvizhevsky and J. Villain. "The Discovery of the Neutron and Its Consequences (1930–1940)". In: *Comptes Rendus Physique*. Science in the Making: The Comptes Rendus de l'Académie Des Sciences throughout History 18.9 (2017), pp. 592–600. ISSN: 1631-0705. DOI: `10.1016/j.crhy.2017.11.001`.

[ONe14] M. E. O'Neill. *PCG: A Family of Simple Fast Space-Efficient Statistically Good Algorithms for Random Number Generation*. HMC-CS-2014-0905. Harvey Mudd College, Claremont, CA, 2014.

[Ott+17] A. N. Otte et al. "Characterization of Three High Efficiency and Blue Sensitive Silicon Photomultipliers". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 846 (2017), pp. 106–125. ISSN: 0168-9002. DOI: `10.1016/j.nima.2016.09.053`.

[Pat+18] R. W. Pattie et al. "Measurement of the Neutron Lifetime Using a Magneto-Gravitational Trap and in Situ Detection". In: *Science* 360.6389 (2018), pp. 627–632. DOI: `10.1126/science.aan8895`.

[Paw01] Y. Pawitan. *In All Likelihood: Statistical Modelling and Inference Using Likelihood*. Oxford University Press: 2001. ISBN: 978-0-19-850765-9. DOI: `10.1093/oso/9780198507659.001.0001`.

[PET24] PETsys Electronics. *TOFPET 2C/D SiPM Readout ASIC*. datasheet. Version REV 18. 2024.

[Pic+10] A. Pichlmaier et al. "Neutron Lifetime Measurement with the UCN Trap-in-Trap MAMBO II". In: *Physics Letters B* 693.3 (2010), pp. 221–226. ISSN: 0370-2693. DOI: `10.1016/j.physletb.2010.08.032`.

[Pie+20] A. Pietropaolo et al. “Neutron Detection Techniques from µeV to GeV”. In: *Physics Reports*. Neutron Detection Techniques from µeV to GeV 875 (2020), pp. 1–65. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2020.06.003.

[Pov+14] B. Povh et al. *Teilchen und Kerne: Eine Einführung in die physikalischen Konzepte*. Springer-Lehrbuch. Springer: Berlin, Heidelberg, 2014. ISBN: 978-3-642-37821-8 978-3-642-37822-5. DOI: 10.1007/978-3-642-37822-5.

[Pro+22] T. Prohaska et al. “Standard Atomic Weights of the Elements 2021 (IUPAC Technical Report)”. In: *Pure and Applied Chemistry* 94.5 (2022), pp. 573–600. ISSN: 1365-3075. DOI: 10.1515/pac-2019-0603.

[Qai12] S. M. Qaim. “The Present and Future of Medical Radionuclide Production”. In: *Radiochimica Acta* 100.8–9 (2012), pp. 635–651. ISSN: 2193-3405. DOI: 10.1524/ract.2012.1966.

[Rae24] Raesch Quarz GmbH. *Material Specification RQ 500*. datasheet. Version Rev. 1. SPEC_M_80_Eng. 2024.

[Rob51] J. M. Robson. “The Radioactive Decay of the Neutron”. In: *Phys. Rev.* 83.2 (1951), pp. 349–358. DOI: 10.1103/PhysRev.83.349.

[Roß21] K. Roß. “Towards a High Precision Measurement of the Free Neutron Lifetime with τSPECT”. PhD thesis. Johannes Gutenberg-Universität Mainz, 2021.

[Sch+89] K. Schreckenbach et al. “Neutron Decay Measurements with a Helium-Filled Time Projection Chamber”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 284.1 (1989), pp. 120–122. ISSN: 0168-9002. DOI: 10.1016/0168-9002(89)90262-3.

[Sea92] V. F. Sears. “Neutron Scattering Lengths and Cross Sections”. In: *Neutron News* 3.3 (1992), pp. 26–37. ISSN: 1044-8632. DOI: 10.1080/10448639208218770.

[Sem22] Semiconductor Components Industries LLC. *Silicon Photomultipliers (SiPM), Low-Noise, Blue-Sensitive C-Series SiPM Sensors*. datasheet. Version Rev. 9. MICROC-SERIES/D. 2022.

[Sem23] Semiconductor Components Industries LLC. *Introduction to the Silicon Photomultiplier (SiPM)*. Technical paper. Version Rev. 9. AND9770/D. 2023.

[Ser+05] A. Serebrov et al. “Measurement of the Neutron Lifetime Using a Gravitational Trap and a Low-Temperature Fomblin Coating”. In: *Physics Letters B* 605.1 (2005), pp. 72–78. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2004.11.013.

[Ser+18] A. P. Serebrov et al. “Neutron Lifetime Measurements with a Large Gravitational Trap for Ultracold Neutrons”. In: *Phys. Rev. C* 97.5 (2018), p. 055503. DOI: 10.1103/PhysRevC.97.055503.

[SG20] A. Steyerl and E. Gutsmiedl. “Spin-Flip Loss in Magnetic Traps for Ultracold Neutrons”. In: *Phys. Rev. C* 102.4 (2020), p. 045203. DOI: 10.1103/PhysRevC.102.045203.

[Sig] Siglent Technologies Co. LTD. *SDS5000X Series Digital Storage Oscilloscope*. datasheet. Version EN02E.

[Ste+12] A. Steyerl et al. "Quasielastic Scattering in the Interaction of Ultracold Neutrons with a Liquid Wall and Application in a Reanalysis of the Mambo I Neutron-Lifetime Experiment". In: *Phys. Rev. C* 85.6 (2012), p. 065503. DOI: 10.1103/PhysRevC.85.065503.

[Ste+17] A. Steyerl et al. "Spin Flip Loss in Magnetic Confinement of Ultracold Neutrons for Neutron Lifetime Experiments". In: *Phys. Rev. C* 95.3 (2017), p. 035502. DOI: 10.1103/PhysRevC.95.035502.

[Sum+21] N. Sumi et al. "Precise Neutron Lifetime Measurement Using Pulsed Neutron Beams at J-PARC". In: *Proceedings of the 3rd J-PARC Symposium (J-PARC2019)*. Vol. 33. JPS Conference Proceedings 33. Journal of the Physical Society of Japan: 2021. DOI: 10.7566/JPSCP.33.011056.

[Sun+18] X. Sun et al. "Search for dark matter decay of the free neutron from the UCNA experiment: $n \to \chi + e^+e^-$". In: *Phys. Rev. C* 97.5 (2018), p. 052501. DOI: 10.1103/PhysRevC.97.052501.

[Tan+18] Z. Tang et al. "Search for the Neutron Decay $n \to X + \gamma$, Where X is a Dark Matter Particle". In: *Phys. Rev. Lett.* 121.2 (2018), p. 022505. DOI: 10.1103/PhysRevLett.121.022505.

[Tek22] Tektronix Inc. *2450 SourceMeter SMU Instrument Datasheet*. datasheet. Version Rev. 02.2022. 2022.

[Tex23] Texas Instruments. *TMP61 ±1% 10-kΩ Linear Thermistor With 0402 and 0603 Package Options*. datasheet. Version DECEMBER 2018 – REVISED NOVEMBER 2023. SBOS921F. 2023.

[The04] The IEEE and The Open Group. *Usleep - Suspend Execution for Microseconds*. The Open Group Base Specifications Issue 6 IEEE Std 1003.1, 2004 Edition. The Open Group, 2004. URL: https://pubs.opengroup.org/onlinepubs/009695399/functions/usleep.html.

[Tho13] M. Thomson. *Modern Particle Physics*. Cambridge University Press: 2013. ISBN: 978-1-139-52536-7. DOI: 10.1017/CBO9781139525367.

[Thy20] Thyssenkrupp Schulte GmbH. *Werkstoffdatenblatt 1.4435*. datasheet. 2020.

[Tsv22] A. Tsvetkov. "Real time experiment control for τSPECT". Master thesis. Johannes Gutenberg-Universität Mainz, 2022.

[Van+25] P. Vander Griend et al. "The Fermi Function and the Neutron's Lifetime". In: *Physics Letters B* 868 (2025), p. 139678. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2025.139678.

[Wan+15] Z. Wang et al. "A Multilayer Surface Detector for Ultracold Neutrons". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 798 (2015), pp. 30–35. ISSN: 0168-9002. DOI: 10.1016/j.nima.2015.07.010.

[Wan+21] M. Wang et al. "The AME 2020 Atomic Mass Evaluation (II). Tables, Graphs and References*". In: *Chinese Phys. C* 45.3 (2021), p. 030003. ISSN: 1674-1137. DOI: 10.1088/1674-1137/abddaf.

[Wan+22] Y. Wang et al. "Production Review of Accelerator-Based Medical Isotopes". In: *Molecules* 27.16 (2022), p. 5294. ISSN: 1420-3049. DOI: 10.3390/molecules27165294.

[Wei+23] Q. Weitzel et al. "3D-printing of Transparent Granulate Materials for Light Guides and Scintillation Detectors". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1046 (2023), p. 167682. ISSN: 0168-9002. DOI: 10.1016/j.nima.2022.167682.

[WG11] F. E. Wietfeldt and G. L. Greene. "Colloquium: The Neutron Lifetime". In: *Rev. Mod. Phys.* 83.4 (2011), pp. 1173–1192. DOI: 10.1103/RevModPhys.83.1173.

[Wie18] F. E. Wietfeldt. "Measurements of the Neutron Lifetime". In: *Atoms* 6.4 (2018), p. 70. ISSN: 2218-2004. DOI: 10.3390/atoms6040070.

[Wie24] F. E. Wietfeldt. "The Neutron Lifetime Discrepancy and Its Implications for Cosmology and Dark Matter". In: *Symmetry* 16.8 (2024), p. 956. ISSN: 2073-8994. DOI: 10.3390/sym16080956.

[Wie25] F. E. Wietfeldt. "The NIST Beam Neutron Lifetime Experiments". Conference talk. Satellite Workshop "Neutron Lifetime Puzzle - Measurements, Systematics, and Theoretical Implications at PSI2025" (Villigen, Switzerland). 2025. URL: https://indico.psi.ch/event/17167/contributions/57992/.

[Wil+20] J. T. Wilson et al. "Space-Based Measurement of the Neutron Lifetime Using Data from the Neutron Spectrometer on NASA's MESSENGER Mission". In: *Phys. Rev. Res.* 2.2 (2020), p. 023316. DOI: 10.1103/PhysRevResearch.2.023316.

[Wil+21] J. T. Wilson et al. "Measurement of the Free Neutron Lifetime Using the Neutron Spectrometer on NASA's Lunar Prospector Mission". In: *Phys. Rev. C* 104.4 (2021), p. 045501. DOI: 10.1103/PhysRevC.104.045501.

[Yaz24] N. Yazdandoost. "Improving Ultracold Neutron Density and Storage: Applications to Neutron Lifetime and Electric Dipole Moment Experiments". PhD thesis. Johannes Gutenberg-Universität Mainz, 2024.

[YMG86] Z.-C. Yu, S. S. Malik, and R. Golub. "A Thin Film Source of Ultra-Cold Neutrons". In: *Z. Physik B - Condensed Matter* 62.2 (1986), pp. 137–142. ISSN: 1431-584X. DOI: 10.1007/BF01323423.

[YOF23] T.-H. Yeh, K. A. Olive, and B. D. Fields. "The Neutron Mean Life and Big Bang Nucleosynthesis". In: *Universe* 9.4 (2023). ISSN: 2218-1997. DOI: 10.3390/universe9040183.

[YT21] S. Yamamoto and H. Tomita. "Comparison of Light Outputs, Decay Times, and Imaging Performance of a ZnS(Ag) Scintillator for Alpha Particles, Beta Particles, and Gamma Photons". In: *Applied Radiation and Isotopes* 168 (2021), p. 109527. ISSN: 0969-8043. DOI: 10.1016/j.apradiso.2020.109527.

[Yue+13] A. T. Yue et al. "Improved Determination of the Neutron Lifetime". In: *Phys. Rev. Lett.* 111.22 (2013), p. 222501. DOI: 10.1103/PhysRevLett.111.222501.

[YWF15] T. Yanagida, K. Watanabe, and Y. Fujimoto. "Comparative Study of Neutron and Gamma-Ray Pulse Shape Discrimination of Anthracene, Stilbene, and p-Terphenyl". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*. Symposium on Radiation Measurements and Applications 2014 (SORMA XV) 784 (2015), pp. 111–114. ISSN: 0168-9002. DOI: 10.1016/j.nima.2014.12.031.

[Zha+22] J. Zhang et al. "Scintillation Detectors with Silicon Photomultiplier Readout in a Dilution Refrigerator at Temperatures down to 0.2 K". In: *J. Inst.* 17.06 (2022), P06024. ISSN: 1748-0221. DOI: 10.1088/1748-0221/17/06/P06024.

[ZZB10] J. F. Ziegler, M. D. Ziegler, and J. P. Biersack. "SRIM – The Stopping and Range of Ions in Matter (2010)". In: *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*. 19th International Conference on Ion Beam Analysis 268.11 (2010), pp. 1818–1823. ISSN: 0168-583X. DOI: 10.1016/j.nimb.2010.02.091.

Diese Seite wurde aus Datenschutzgründen entfernt.