Improving ultracold neutron statistics for the n2EDM experiment

A thesis submitted to attain the degree of

Doctor of Sciences
(Dr. sc. ETH Zurich)

presented by

Cornelis Bernardus Doorenbos

MSc in Physics
University of Fribourg

born on 4 April 1993

accepted on the recommendation of

Prof. Dr. Klaus Stefan Kirch
Prof. Dr. Kazimierz Bodek
Dr. Bernhard Lauss

2024

ABSTRACT

The n2EDM experiment, currently being commissioned at the UCN source at the Paul Scherrer Institute, aims to search for the neutron electric dipole moment (nEDM) with a sensitivity of $1.1 \times 10^{-27} \, e \cdot \mathrm{cm}$. This requires improving both the statistical and the systematic uncertainties, compared to the previous experiment. The n2EDM experiment is a Ramsey experiment with ultracold neutrons (UCN). The statistical sensitivity depends on the number of UCN counted in the detectors at the end of each Ramsey cycle. Therefore, it is important to maximise the UCN yield of the UCN source and the UCN transport and storage in the n2EDM experiment. In this thesis, I will present work on both of these aspects.

The UCN source at the PSI is a pulsed, spallation-driven, solid-deuterium-moderated, superthermal UCN source. UCN are generated by the downscattering of cold neutrons via the excitation of phonon states in the solid deuterium lattice. Using Raman spectroscopy, we determined that the hydrogen deuteride (HD) concentration in the UCN source solid deuterium is $(0.210 \pm 0.002_{\mathrm{sys}} \pm 0.016_{\mathrm{stat}})$ %. At this concentration, neutron capture in hydrogen is not the dominant loss channel for UCN. It is not recommended to replace the UCN source deuterium with even higher-purity deuterium.

During UCN source operation, rapid temperature cycling causes a rough layer of 'frost' to build up on the solid deuterium surface, which reduces the UCN extraction rate and thus the UCN yield. The UCN yield is recovered by a gradual thermal cycling procedure, called 'conditioning'. We tested a new method of conditioning that relies on radiation-induced heating from proton beam pulses instead of electric heaters. We verified that conditioning with proton beam pulses is reliable enough to be automated, and determined the optimal conditioning frequency.

In 2022, we measured a record UCN yield during production pulses, of 5.71×10^7 normalised to $2.2 \, \mathrm{mA}$ of beam current, with a solid deuterium mass in the moderator vessel of $5.28 \, \mathrm{kg}$. With $5.68 \, \mathrm{kg}$ of solid deuterium in the moderator vessel, we measured a lower UCN yield. We suspect that in the latter case, the deuterium surface was not sufficiently cooled, because the filling level exceeded the height of the cooling channels in the moderator vessel, reducing UCN output. We expect we can achieve even larger UCN yields with a moderator vessel with adapted cooling channels.

We compared the UCN intensities at the UCN source beamports South and West-1. For storable UCN, the UCN intensity ratio at these beamports has a small relative standard deviation of 0.14 %. While the UCN intensity ratio of West-2 to South/West-1 is less stable, we can use a detector on West-2 to monitor the UCN intensity in an experiment on South or West-1 with an accuracy of 0.4 %, when we apply drift correction and limit the number of pulses to 50. We determined that the UCN intensity at beamport South was (81.6 ± 0.2) % of the intensity at West-1.

We measured the UCN transmission of the UCN guides for n2EDM using a direct transmission setup. The best guides had a transmission of more than 95 %, normalised to a guide length of $1 \, \mathrm{m}$. Some UCN guides had bad UCN transmission due to wavy irregularities on the interior of the guides. We made new guides from tubes without these wrinkles to provide more UCN to n2EDM.

The UCN guides that are installed closest to the precession chamber are not allowed to contain magnetic dipoles larger than $15 \, \mathrm{nA \, m^2}$. This would disrupt the magnetic field gradient reconstruction by the caesium

magnetometer array, and cause systematic errors. We searched for dipoles using our magnetic gradiometer and found that the guides satisfy this limit.

We investigated the shift in resonance frequency in n2EDM caused by the effect of the $\frac{\pi}{2}$-pulse for the mercury comagnetometer on the UCN spin state, as well as a possible shift induced by a tilt between the static magnetic field and the rotating magnetic field axis during the Ramsey sequence. We designed window functions to limit these frequency shifts to a value well below the systematics budget.

ZUSAMMENFASSUNG

Das n2EDM-Experiment, das sich derzeit an der Quelle für ultrakalte Neutronen (UCN) am Paul Scherrer Institut im Aufbau befindet, hat das Ziel, das elektrische Dipolmoment des Neutrons (nEDM) zu messen. Dabei soll eine Sensitivität von $1.1 \times 10^{-27}\,e \cdot \mathrm{cm}$ erreicht werden. Dies erfordert eine Verbesserung der statistischen und systematischen Unsicherheiten im Vergleich zu dem vorherigen Experiment. Das n2EDM-Experiment ist ein Ramsey-Experiment mit ultrakalten Neutronen (UCN). Die statistische Sensitivität ist abhängig von der Anzahl der UCN, welche am Ende jedes Ramsey-Zyklus durch Detektoren gezählt werden. Daher müssen Produktion in der Quelle, Transport der UCN und deren Speicherung in n2EDM maximiert werden. In dieser Dissertation präsentiere ich meine Arbeit an den ersten zwei Aspekten.

Die UCN-Quelle am PSI ist eine gepulste, spallationsgetriebene, durch festes Deuterium moderierte, superthermische UCN-Quelle. Die UCN werden durch Herunterstreuen von kalten Neutronen in festem Deuterium erzeugt, wobei Phononzustände im Gitter angeregt werden. Mit Raman-Spektroskopie haben wir festgestellt, dass die Wasserstoffdeuteridkonzentration (HD) im festen Deuterium der UCN-Quelle $(0.210 \pm 0.002_{\mathrm{sys}} \pm 0.016_{\mathrm{stat}})$ % beträgt. Bei dieser Konzentration ist Neutroneneinfang in Wasserstoff kein dominanter Verlustkanal. Wir empfehlen daher nicht, das Deuterium mit noch reinerem Deuterium zu ersetzen.

Während des Betriebs der UCN-Quelle verursachen schnelle Temperaturwechsel den Aufbau einer Frostschicht an der Oberfläche des festen Deuteriums. Dies reduziert die UCN-Extraktion und in Folge die UCN-Intensität an den Strahlausgängen. Die UCN-Ausbeute kann allerdings durch ein Verfahren, in welchem die Temperatur langsam verändert wird, wiederhergestellt werden. Man nennt diesen Prozess 'Aufbereitung'. Es wurde eine neue Aufbereitungsmethode untersucht, welche auf der Erwärmung durch Protonenstrahlpulse beruht, anstatt das Deuterium mit Hilfe elektrischer Heizelemente zu erwärmen. Wir haben verifiziert, dass diese Methode zuverlässig genug ist, um automatisiert zu werden, und haben die optimale Aufbereitungsfrequenz etabliert.

Im Jahr 2022 wurde an der PSI UCN-Quelle ein Rekord-UCN-Ertrag, normalisiert auf einen Strahlstrom von 2.2 mA, von 5.71×10^7 UCN gemessen. Dabei befanden sich 5.28 kg an festem Deuterium im Moderatorbehälter. Mit 5.68 kg an festem Deuterium war der UCN-Ertrag geringer. Wir vermuten, dass im zweiten Fall die Deuteriumoberfläche nicht ausreichend gekühlt wurde, da der Füllstand die Höhe der Kühlkanäle im Moderatorbehälter übertroffen hat. Wir erwarten, dass mit einem neuen Moderatorbehälter und mit angepassten Kühlkanälen noch größere UCN-Erträge erzielt werden können.

Es wurden ebenfalls die UCN-Intensitäten an den Strahlausgängen West-1 und Süd verglichen. Für speicherbare UCN hat das UCN-Intensitätsverhältnis eine kleine relative Standardabweichung von 0.14 %. Das Intensitätsverhältnis der Strahlausgänge West-2 und West-1/Süd ist weniger stabil. Um die Intensität an Experimenten der anderen Strahlausgänge mit einer Genauigkeit von 0.4 % zu überwachen, kann ein Detektor an West-2 verwendet werden, wobei eine Steigungskorrektur anwendet und die Anzahl der Pulse auf 50 begrenzt werden muss. Wir haben festgestellt, dass die UCN-Intensität am Strahlausgang Süd (81.6 ± 0.2) % der Intensität an West-1 beträgt.

Die UCN-Transmission der UCN-Leiter für n2EDM wurde mit Hilfe eines direkten Transmissionsaufbaus gemessen. Die besten Leiter hatten eine Transmission von über 95 %. Einige UCN-Leiter zeigten eine schlechte Transmission aufgrund von wellenförmigen Unebenheiten im Inneren der Leiter. Wir haben

neue UCN-Leiter aus Rohren ohne sichtbare Unebenheiten hergestellt, um die Transmission von UCN zum n2EDM-Experiment zu maximieren.

Die UCN-Leiter, die am nächsten zur Präzessionskammer sind, dürfen keine magnetische Dipole grö-ßer als $15\,\text{nA}\,\text{m}^2$ enthalten. Das würde die Rekonstruktion des Magnetfeldgradienten durch das Cäsium-Magnetometerarray stören und systematische Unsicherheiten verursachen. Die magnetische Kontamination in den Leitern wurde durch Messungen in einem Gradiometer gemessen. Es wurde festgestellt, dass die Leiter diese Grenze einhalten.

Wir haben mögliche Verschiebungen in der Resonanzfrequenz des Ramsey-Zyklus, verursacht durch die Quecksilbercomagnometer-$\frac{\pi}{2}$-Pulse oder durch eine Neigung zwischen statischen Magnetfeld und der Ach-se des rotierenden Magnetfelds, untersucht. Wir haben Fensterfunktionen entworfen, um diese Frequenzver-schiebungen weit unter dem systematischen Fehlerbudget zu halten.

TABLE OF CONTENTS

List of Figures

LIST OF TABLES

CHAPTER 1

INTRODUCTION

1.1 Overview

In this introduction, the search for the neutron electric dipole moment (nEDM) is motivated. The principles of nEDM experiments, which have not changed since 1950, are explained. A short introduction to ultracold neutron (UCN) physics is given to motivate their suitability for nEDM experiments. After this general introduction, the reader is acquainted with the UCN source and the n2EDM experiment at PSI, where the work detailed in this thesis took place.

1.2 The search for the neutron electric dipole moment

1.2.1 Baryogenesis

One of the great unsolved puzzles in physics and cosmology is the baryon asymmetry of the Universe: the predominance of matter over anti-matter. Without baryon asymmetry, almost all matter would have annihilated shortly after the Big Bang. The Universe would be filled with photons. No physicists would be around to ponder the absence of planets and stars.

In 1967, A.D. Sakharov postulated three necessary conditions [1], [2] to generate baryon asymmetry in the Universe:

1. the existence of processes that violate the baryon number,

2. violation of charge (C) and charge-parity (CP) symmetry to ensure an asymmetric state can develop from a symmetric initial state,

3. departure from thermal equilibrium, so that charge-parity-time (CPT) symmetry does not cause processes increasing and decreasing the baryon number to cancel each other out.

The amount of CP violation in the Standard Model of Particle Physics (SM) is insufficient to explain the observed amount of baryonic matter in the Universe [3]. This conclusion prompted the search for new sources of CP violation.

1.2.2 CP violation and the neutron electric dipole moment

Classically, electric dipole moments (EDM, $\mathbf{d}$) and electric fields ($\mathbf{E}$) are represented by vectors, whereas magnetic dipole moments (MDM, $\boldsymbol{\mu}$) and magnetic fields ($\mathbf{B}$) are represented by pseudovectors (axial vectors). Vectors change sign under parity inversion $(t, x, y, z,) \rightarrow (t, -x, -y, -z)$ but are invariant under time reversal $(t, x, y, z) \rightarrow (-t, x, y, z)$, whereas pseudovectors change sign under time inversion but are invariant under parity inversion. EDMs and MDMs are not physically observable, but their interactions with electric and magnetic fields are observable. Classically, the scalar products $\mathbf{d} \cdot \mathbf{E}$ and $\boldsymbol{\mu} \cdot \mathbf{B}$ are not inverted or twice inverted under time reversal and parity inversion and are thus invariant under these operations.

Quantum mechanically, the spin-axis is the only preferred direction in the system of a particle and both $\mathbf{d}$ and $\boldsymbol{\mu}$ must be proportional to it.[1] *Ipso facto*, the neutron EDM transforms as a pseudovector, together with the spin and the MDM. This means that the sign of $\mathbf{d} \cdot \mathbf{E}$ is reversed under parity inversion and time reversal, while it is invariant to charge inversion. The existence of a permanent EDM as a property of a fundamental particle is CP-violating. CPT-symmetry is conserved.

Sources of CP violation in the SM and in physics beyond the SM predict different contributions to the electric dipole moments of hadrons or leptons, including the neutron. The search for EDMs is thus a sensitive probe for new physics [5], [6]. Searches for permanent EDMs of nuclei, hadrons, and leptons have not resulted in discoveries of non-zero EDMs, but have provided stringent upper limits [7]. In particular, the current most accurate measurement of the neutron electric dipole moment (nEDM), conducted by the nEDM collaboration at PSI, is compatible with zero: $d_n = (0.0 \pm 1.1_{\text{stat}} \pm 0.2_{\text{sys}}) \times 10^{-26} e \cdot \text{cm}$ [8].

The Standard Model

In quantum chromodynamics (QCD), the theory of the strong interaction, the topological vacuum structure generates a CP-violating term $\theta G_{\mu\nu} \tilde{G}^{\mu\nu}$ in the Lagrangian, where θ is a free parameter of the theory, and G is the gluon field strength tensor. The parameter θ is naturally expected to be of order 1 rad, which leads to a predicted nEDM of order $5 \times 10^{-16} e \cdot \text{cm}$ [9]. The current experimental limit on the nEDM ($1.8 \times 10^{-26} e \cdot \text{cm}$, 90 % C.L. [8]) imposes a stringent limit: $\theta \lesssim 10^{-10}$. This fine-tuning is called the Strong CP Problem [10]. Since θ is so small, it is natural to assume that it vanishes due to a symmetry or a dynamic mechanism [11], [12].

Within the SM, mixing of quark mass states and weak interaction states, represented by the Cabibbo-Kobayashi-Maskawa (CKM) matrix [13], provides a contribution to the nEDM of up to $1 \times 10^{-31} e \cdot \text{cm}$ [12]. This will be beyond the reach of nEDM experiments for the foreseeable decades.

Beyond Standard Model physics

If the θ parameter in QCD is promoted to a field [11], this gives rise to a light, pseudoscalar boson, the axion [14]. The axion field would naturally relax to zero, solving the strong CP-problem. In addition, axions could be a constituent of dark matter. The axion field would produce a periodic oscillation in the nEDM, which could be detected in an nEDM experiment. The most sensitive experiments with free neutrons to date have not found evidence of axion-like particles [15].

1.2.3 Ramsey's separated oscillating fields method

The search for the nEDM started in the 1950s. In 1950, Norman Ramsey Jr. developed the separated oscillating fields method [16] to measure the MDM of molecules using molecular beams. This method is a

---

[1] According to the Wigner-Eckart theorem [4], a spin-0 elementary particle cannot possess an EDM, since the only preferred direction in an elementary particle is the spin axis.

refinement of Rabi's method [17]. It is a nuclear resonance technique to find the Larmor precession frequency ω_0 of particles in a magnetic field B, given by

$$\hbar\omega_0 = 2\mu B, \tag{1.1}$$

for a particle with a magnetic dipole moment μ.

In the separated oscillating fields method, particles are prepared such that their spins are polarised along the axis of a static magnetic field. An oscillating (or rotating) magnetic field pulse is applied in the plane perpendicular to the static field. The frequency of this pulse, ω, should be close to the precession frequency of the particles about the static field, ω_0. The field strength and duration of the pulse are chosen such that the spin polarisation of the particles is rotated by 90°, so that it precesses in the plane perpendicular to the static field. After an interval of free precession, a second pulse is applied. Both pulses have an identical intensity and frequency, and they are in phase with each other. In the resonance condition ($\omega = \omega_0$), the spin polarisation has been rotated by 180° relative to its initial direction. The resonance frequency is found by varying the frequency of the oscillating (or rotating) pulse.

To measure EDMs, a static electric field is required, collinear to the static magnetic field. The EDM contributes to the Larmor frequency:

$$\hbar\omega_0^\pm = 2\mu B \pm 2dE. \tag{1.2}$$

The sign of the EDM contribution depends on the sign of d and whether the static electric and magnetic fields are parallel or antiparallel. The Larmor frequency is measured in both of these cases. Then the EDM is calculated by subtracting the frequencies [18]:

$$d = \frac{\hbar(\omega_0^+ - \omega_0^-)}{4E}. \tag{1.3}$$

The statistical sensitivity for a single cycle of a Ramsey experiment is [19], [20]

$$\sigma(d_n) = \frac{\hbar}{2\alpha E T \sqrt{N}}. \tag{1.4}$$

The product of the electric field strength, E, and the precession duration, T, is proportional to the phase shift built up during the Ramsey cycle. The amplitude of the observed signal, α, is proportional to the polarisation of the neutrons at the end of the Ramsey cycle, and the efficiency of the spin analysis. N is the number of detected neutrons. It follows that we can increase the statistical sensitivity of a Ramsey experiment by increasing these four parameters.

1.2.4   A brief history of nEDM experiments

Ramsey's separated oscillating fields method was first used to search for the nEDM in Oak Ridge National Laboratories (ORNL) in 1951. The experiment was done using a beam of neutrons and imposed a limit of $|d_n| < 5 \times 10^{-20}\, e \cdot \mathrm{cm}$ [21]. For the next two decades, increasingly sensitive nEDM measurements were done using neutron beams, culminating in a limit of $|d_n| < 3 \times 10^{-24}\, e \cdot \mathrm{cm}$ in 1974 [22]. At this point, the velocity-dependent motional magnetic field, $\mathbf{B}_m = \mathbf{v} \times \mathbf{E}/c^2$, became an important systematic effect. To suppress this effect, ultracold neutrons (UCN) were used to measure the nEDM in the Leningrad Nuclear Physics Institute, [2] in 1980. Apart from suppressing the motional magnetic field, the use of stored UCN enhances the statistical sensitivity by increasing the free precession time. This experiment brought the limit down to $|d_n| < 1.6 \times 10^{-24}\, e \cdot \mathrm{cm}$ [23]. All subsequent improvements of the nEDM limit were done using UCN.

Presently, the most accurate measurement of the nEDM was done at the Paul Scherrer Institut (PSI) in Switzerland, by the nEDM collaboration at PSI. The experiment found an nEDM of $d_n = (0.0 \pm 1.1_\text{stat} \pm 0.2_\text{sys}) \times 10^{-26} e \cdot \mathrm{cm}$ [8], corresponding to a 90 % confidence limit of $|d_n| < 1.8 \times 10^{-26}\, e \cdot \mathrm{cm}$.

---

[2]Now Petersburg Nuclear Physics Institute (PNPI).

Several groups are working to provide more accurate measurements of the nEDM.

- At PSI, we are working on the n2EDM experiment [24], with a projected sensitivity (1σ) of $10^{-27} e \cdot$ cm in the baseline configuration.[3]

- At the Institut Laue-Langevin in Grenoble, a new nEDM experiment, PanEDM [25], is under construction at the new SuperSUN [26] superfluid helium-4 UCN source. The projected sensitivity is in the mid $10^{-27} e \cdot$ cm range.

- At Canada's particle accelerator centre, TRIUMF, the TRIUMF Ultracold Advanced Neutron Collaboration (TUCAN) is building an nEDM experiment [27] with a projected sensitivity of $10^{-27} e \cdot$ cm.

- In Bern, work is ongoing on the development of a new beam experiment [28], [29]. This experiment aims to use high-intensity cold neutron pulses to directly measure the $\mathbf{v} \times \mathbf{E}$ systematic effect. The sensitivity is expected to be competitive with UCN experiments. The experiment is proposed to be run in the European Spallation Source in Lund, currently under construction.

1.2.5   Ultracold neutrons, and their interactions

Ultracold neutrons (UCN) are neutrons with kinetic energies below 300 neV ($v < 7.57\,\mathrm{m\,s^{-1}}$). This corresponds to temperatures below the 4 mK range [30], [31]. UCN have a de Broglie wavelength $\lambda > 50$ nm. UCN can be confined in material containers or magnetic traps for several minutes. This makes them uniquely suitable for experiments where the sensitivity depends on the neutron storage duration, such as nEDM experiments.

Neutrons are affected by all four fundamental interactions. On the timescales, length scales and magnetic-field scales used in UCN experiments, the four interactions produce effects of similar orders of magnitude.

Weak interaction

Free neutrons are unstable particles. They decay via beta decay, mediated by the weak interaction, producing a proton, an electron, and an electron antineutrino:

$$n \to p^+ + e^- + \bar{\nu}_e. \tag{1.5}$$

The free neutron lifetime measured using UCN traps is $(877.75 \pm 0.28_{\mathrm{sys}} \pm 0.22_{\mathrm{stat}})$ s. [32]. Using neutron beam experiments, a lifetime of $(887.7 \pm 1.2_{\mathrm{sys}} \pm 1.9_{\mathrm{stat}})$ s has been found [33]. The 4σ discrepancy between these values is known as the neutron lifetime puzzle [34]. A neutron lifetime experiment, τSPECT [35], is currently in operation at the PSI UCN source.

The neutron lifetime provides a natural limit for the storage duration of UCN.

Strong interaction

Neutrons interact with nuclei via the residual strong interaction. This interaction is mediated by the exchange of virtual pions. Since pions are massive (in contrast to photons) the interaction range is approximately 1 fm. Nonetheless, the strong interaction has consequences on the macroscopic scale.

A simple but effective model for the strong interaction between a neutron and a proton is a spherical potential well, with a depth of 40 MeV and a radius of $R = 2$ fm. For larger nuclei, $R \approx R_0 A^{1/3}$ [36], where A

---

[3]By setting the static magnetic field to $10.5\,\mu$T (the so-called magic-field configuration) the false EDM due to phantom modes (higher order harmonics of the magnetic field that do not produce a top-bottom gradient in the precession chambers) will be suppressed by at least a factor of 30. This will be done after running the experiment in baseline configuration with a static field of $1.0\,\mu$T [24].

is the atomic mass number (R_0 is smaller than 2 fm). Since $\lambda \gg R$, UCN are primarily scattered by S-wave scattering (zero orbital angular momentum).

The wave function outside the nucleus associated with S-wave scattering is

$$\psi \propto e^{-i\mathbf{k}\cdot\mathbf{r}} - \frac{a}{r}e^{-ikr}, \tag{1.6}$$

where the first term corresponds to the incoming plane wave and the second term to the spherical scattered wave. The scattering length a can be interpreted as the radius of a classical hard sphere that would produce the same scattered wave function. The interaction is repulsive if a is positive, and attractive if a is negative. The interaction force is repulsive for the vast majority of nuclei. The scattering length could be determined for every nucleus if the shape and effective radius of the neutron-nucleus interaction potential were accurately known. In practice, scattering lengths are determined experimentally.

Under the Born approximation (every nucleus scatters only the incident wave, not the scattered waves from other nuclei), the potential associated with the interaction of a neutron with a single, bound nucleus can be approximated by

$$V(\mathbf{r}) = \frac{2\pi\hbar^2 a}{m_n}\delta^{(3)}(\mathbf{r}), \tag{1.7}$$

where m_n is the neutron mass and $\delta^{(3)}$ is the three-dimensional Dirac delta function [36].

For an ensemble of nuclei, separated by distances much smaller than the neutron wavelength, an effective average potential can be calculated:

$$V_F = \frac{2\pi\hbar^2}{m_n}\sum_i a_i N_i, \tag{1.8}$$

where a_i and N_i are respectively the scattering length and number density of the i^{th} isotope. This effective potential is called the neutron optical potential or Fermi potential. When the Fermi potential of a material is larger than the kinetic energy of an impacting neutron, the neutron will be reflected, regardless of the angle of incidence. Materials with high Fermi potentials can be used to store UCN [30], [36].

UCN can be lost during interactions with condensed matter due to inelastic upscattering or neutron capture. This behaviour can be captured in the imaginary component of the interaction potential of a material, often expressed as a fraction of the real Fermi potential,

$$V = V_F(1 - i\eta), \tag{1.9}$$

where the parameter η is called the loss parameter or loss coefficient.

The probability that a UCN is reflected during a wall interaction can be trivially obtained from the solutions of the time-independent Schrödinger equation for the one-dimensional step potential:

$$R = \left|\frac{k_1 - k_2}{k_1 + k_2}\right|^2, \tag{1.10}$$

where $k_1 = \sqrt{2mE/\hbar^2}$ is the neutron wavenumber in vacuum and $k_1 = \sqrt{2m(E - V_F(1 - i\eta))/\hbar^2}$ is the neutron wavenumber inside the wall material.[4] The loss probability during a single wall collision, referred to as the loss-per-bounce rate, equals $1 - R$. The loss-per-bounce rate is shown in Figure 1.1 as a function of energy for different loss coefficients. Larger loss coefficients lead to larger loss-per-bounce rates for UCN with energies below the Fermi potential.

The Fermi potentials and capture loss coefficients for materials relevant in this work are given in Table 1.1.

---

[4]If the UCN is not impacting the wall perpendicularly, k_1 and k_2 are the projections of the wave vector on the normal of the wall plane.

(a)

(b)
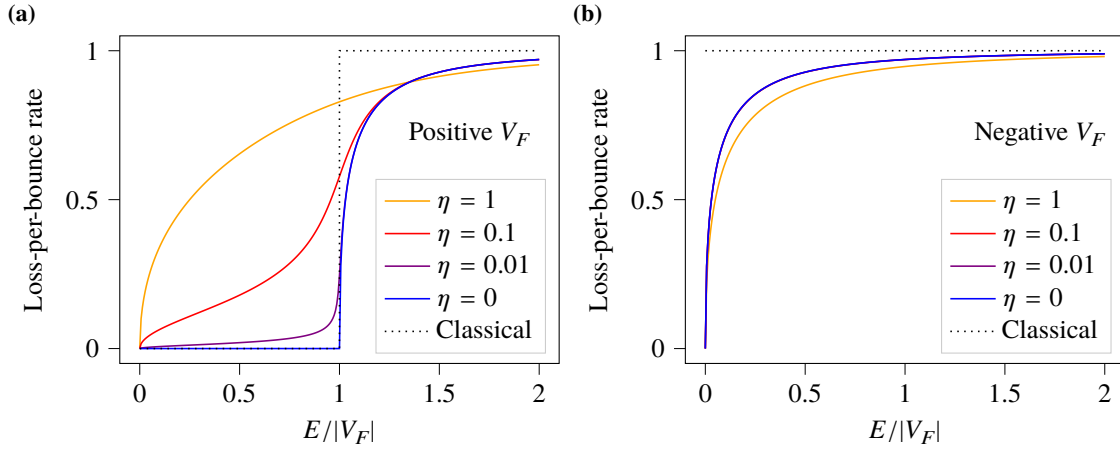
Figure 1.1: The theoretical loss-per-bounce rates for UCN colliding perpendicularly with a wall with a (a) positive or (b) negative Fermi potential V_F and a loss coefficient η, following the calculations in [30]. Classically, UCN would experience no losses if $E < V_F$ and total losses if $E \geq V_F$. Quantum mechanically, non-zero loss coefficients cause UCN losses even if $E < V_F$, while if $E \geq V_F$, the loss-per-bounce rate is not one, but approaches one asymptotically.

Table 1.1: Fermi potentials and loss coefficients of relevant materials. Values taken from [30], unless otherwise indicated.

Material	V_F (neV)	η_{capture} (10^{-5})
Nickel	252	8.9
Diamond-like carbon [37]	235	0.015
NiMo (nickel-molybdenum, 85:15 mass ratio) [37]	220	12
Stainless steel	188	9.3
Copper	165	13.6
Deuterated polystyrene [38]	165	30
Solid D_2 (5 K) [39]	99	0.002
Float glass	96	0.58
Aluminium	54	1.9
Titanium	−50	49

Electromagnetic interaction

Neutrons have no net electric charge but are nonetheless affected by the electromagnetic interaction due to their magnetic dipole moment (MDM). The neutron MDM equals $-60.307\,\text{eV}\,\text{T}^{-1}$ [40]. The potential energy of neutrons in a 5 T field equals $\pm 300\,\text{neV}$. This is of the same order as Fermi potentials of good neutron reflectors such as Nickel ($V_F = 252\,\text{neV}$, see Table 1.1). The sign depends on the spin state: it is negative for high-field seeking neutrons and positive for low-field seeking neutrons. Strong magnetic fields can thus be used to store low-field seekers or to selectively permit the passage of high-field seekers to fill a volume with spin-polarised UCN.

UCN spins precess around magnetic field lines. Using carefully crafted guide fields, the polarisation state of UCN can be preserved during neutron transport. If the guide field direction changes slowly compared to the Larmor frequency of the neutrons around the guide field, the neutron spin axis adiabatically follows the field direction.

During wall collisions with ferromagnetic materials (such as iron or nickel), the UCN can interact with the magnetic domains and lose their polarisation. This is not of concern with non-ferromagnetic materials, or ferromagnetic materials above the Curie temperature. Since the UCN wavelength is three orders of magnitude larger than the interatomic spacing, it interacts with many nuclei at once and the uncorrelated nuclear and electron spins average out to zero.

No electric dipole interactions have been observed to date. The n2EDM experiment [24] aims to either determine a non-zero value for the nEDM for the first time or reduce the limit to an order of magnitude below the current limit of $1.8 \times 10^{-26}\,e \cdot \text{cm}$ (90 % C.L.) [8].

Gravity

With a mass of $939.57\,\text{MeV}\,c^{-2}$ [40], neutrons experience a gravitational force of $102.6\,\text{neV}\,\text{m}^{-1}$ at Earth's surface ($g = 9.81\,\text{m}\,\text{s}^{-2}$). UCN with energies up to 300 neV therefore cannot rise more than 3 m in Earth's gravitational field. The relative height of components of a UCN experiment must be planned carefully. This property can be used for magneto-gravitational UCN confinement without wall collisions. The acceleration of UCN in the gravitational field can also be used to ensure low-energy UCN have enough energy to pass through, for example, a detector entrance foil.

1.3 The UCN source at PSI

The UCN source at PSI is a pulsed, spallation-driven, solid-deuterium-moderated, superthermal UCN source. The UCN source is installed on one of the beamlines of the High Intensity Proton Accelerator facility (HIPA) at PSI. The main science objective driving the design and construction of the UCN source was to search for the nEDM with unprecedented statistical sensitivity. The UCN source started operation in 2011 [31], [41], [42].

An annotated CAD image of the UCN source is shown in Figure 1.2. UCN production can be divided into the following steps:

- the HIPA proton beam induces the production of fast neutrons in the spallation target,

- the fast neutrons are moderated by heavy water at room temperature and become thermal neutrons,

- the thermal neutrons are moderated by solid deuterium at 5 K, producing cold neutrons,

- cold neutrons are converted into UCN by exciting phonon states in the solid deuterium lattice.
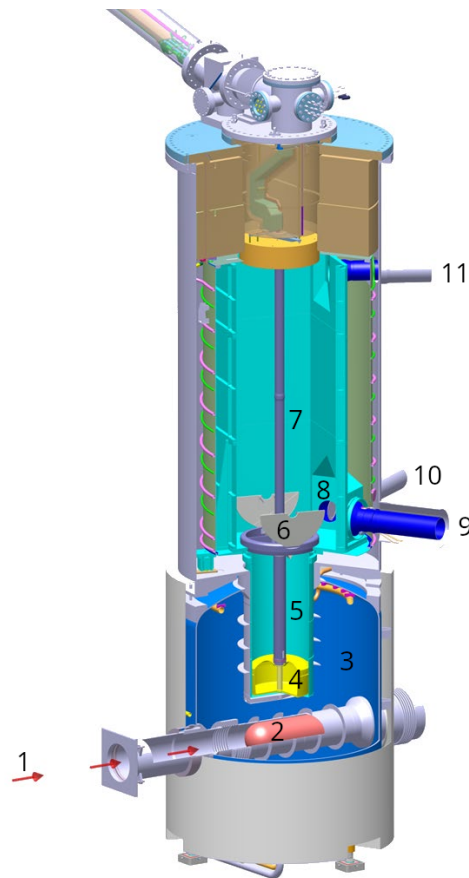
Figure 1.2: An overview of the UCN source at PSI. (1) The HIPA proton beam impacts (2) the spalla-tion target, creating free fast neutrons. In (3) the heavy water tank, fast neutrons are moderated to room temperature. In (4) the solid deuterium moderator, thermal neutrons are moderated to produce cold, and subsequently, ultracold neutrons. The UCN travel up through (5) the vertical beam guide, through (6) the flapper valve, into (7) the UCN storage vessel. After the pulse, the flapper valve closes, storing the UCN in the tank, and ensuring they do not fall back into the moderator. If (8) the UCN shutters (Neutronenleiterklappe, NLK) are open, UCN can be guided to (9) beamline West-1, (10) beamline South, or (11) beamline West-2. (Image from [31].)
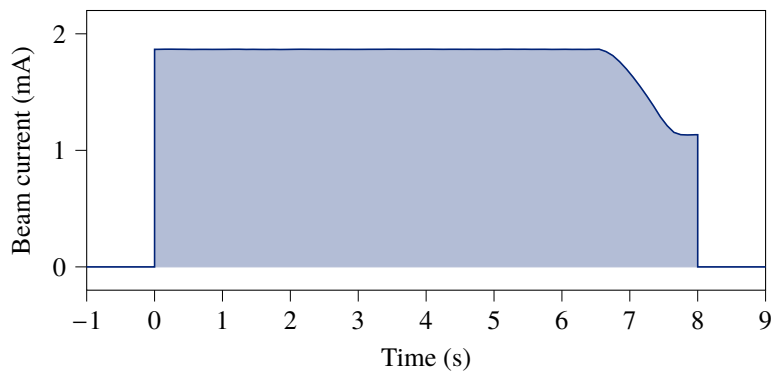
Figure 1.3: The beam current directed to the UCN source spallation target during an 8-second pulse. The beam current is ramped down from 6.5 s after the start of the pulse. This is done to prevent excessive thermal cycling in the SINQ target when the beam is directed back at SINQ.

1.3.1 Proton beam pulses

The HIPA [43] provides protons with kinetic energies of 590 MeV, and beam currents up to 2.4 mA, corresponding to a beam power up to 1.42 MW. Most of the time, the HIPA supplies protons to the muon facilities and the Swiss Spallation Neutron Source (SINQ). During the operation of the UCN source, the proton beam is periodically directed onto the UCN source spallation target in pulses lasting up to 8 s. This is achieved using a fast kicker magnet in the HIPA beamline, designed with a rise time (5 % to 95 %) of 0.5 ms to minimise proton beam losses [44].

After the pulse, the beam is directed back to the SINQ target. If the pulse lasts longer than two seconds, the beam current is ramped down during the last 1.5 s of the pulse (see Figure 1.3). This protects the SINQ target from excessive thermal cycling. As a result, the mean beam current delivered to the UCN spallation target is slightly lower than the nominal beam current.

Two types of pulses were used to generate UCN data in this work. Production pulses are 8 s pulses, intended to deliver the largest possible amount of UCN to an experiment. The flapper valve (described below) at the bottom of the UCN storage vessel closes at the end of the pulse to store the UCN, so they can be guided to the experiment. Norm pulses are 2 s pulses, used to monitor the UCN yield of the UCN source in a standardised way, so that it can be compared over multiple years [45]–[48]. Because they last only two seconds, the proton beam is not ramped down at the end of the pulse. The flapper valve is not operated and kept open, eliminating UCN intensity variation due to the timing of mechanical components.[5]

In this work, UCN data have been used from 2021 to 2023. The HIPA facility can deliver beam currents up to 2.4 mA, but the capacity during these years was temporarily reduced, due to work on cavities of Injector-2 of the HIPA beamline. The mean beam current on target during production pulses was typically close to 1.8 mA (the ramp at the end of the pulse is taken into account). The term 'standard pulse' is used throughout this work as a unit of beam charge on target of 1.8 mA $\cdot$ 8 s, or 14.4 mC.

1.3.2 The spallation target

Neutrons are produced in the UCN source via the process of nuclear spallation [49]. When a heavy nucleus is impacted by an energetic particle, it enters an excited state. This state relaxes through the emission of nucleons, through a process called nuclear evaporation. Up to 30 neutrons can be released, depending on the input energy and the mass of the target nucleus. In contrast, nuclear fission typically releases two or three neutrons per nucleus of fuel. Fission is an exothermic process, the fission fragments deposit a lot of energy

---

[5]Since 2019, one of the flaps remains closed permanently. Thus, during norm pulses, one flap remains open and the other remains closed.
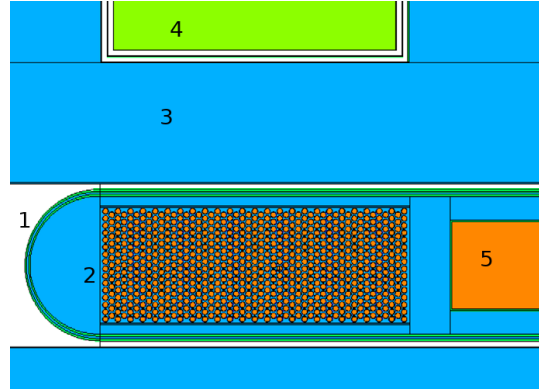
Figure 1.4: A visualisation of the MCNP model used to simulate the UCN source spallation target, adapted from [50]. (1) beam window, (2) array of Zirconium-clad lead tubes, (3) heavy water moderator, (4) solid deuterium moderator, (5) aluminium clad lead.

in the fuel. Spallation is thus a more efficient process for neutron production than fission, releasing more neutrons for a given thermal power in the target/fuel assembly. Spallation is endothermic and therefore self limiting, and does not require nuclear fuel. This makes it simpler to acquire the materials and permissions to build a spallation source.

Lead is the primary spallation material in the UCN source. The spallation target consists of 760 hollow Zircaloy tubes ($\varnothing_{outer} = 10.75$ mm, $\varnothing_{inner} = 9.25$ mm) filled with lead [50]. To accommodate thermal expansion of the lead, the filling factor is approximately 90 %. The spacing between the tubes is 2.00 mm, permitting the circulation of heavy water to cool the assembly. The tubes are oriented horizontally, transversally to the beam direction. The tubes are packed in a hexagonal lattice, which guarantees that there are no gaps for protons to pass through. Using MCNP simulations, it was predicted that approximately 7.3 neutrons are produced per primary proton [51]. A visualisation of the MCNP model of the target is shown in Figure 1.4.

1.3.3   Neutron moderation

The fast neutrons leave the target assembly and are moderated to thermal neutrons ($T \approx 298$ K, $v \approx 2200$ m s^{-1}) in heavy water, in the heavy water tank that surrounds the spallation target (3 in Figure 1.2).

A fraction of the thermal neutrons enter the solid deuterium moderator vessel (4). The moderator vessel is made from the aluminium alloy AlMg4.5Mn. It is cylindrical, with a diameter of about half a metre. The bottom and walls of the vessel are cooled using supercritical helium around 4 bar to maintain a deuterium temperature around 5 K. The volume of the vessel is 44 L. The amount of solid deuterium in the vessel varies, with typical values of around 4.5 kg [48].

The thermal neutrons are moderated inside the solid deuterium (near 5 K), and become cold neutrons. The cold neutron energy spectrum overlaps with the available phonon states in the solid deuterium (see Figure 1.5). The cold neutrons can excite phonons in the lattice, sometimes losing almost all of their energy. The resulting neutrons are called 'ultracold neutrons' (UCN). Upon leaving the solid deuterium, the UCN gain a boost perpendicular to the surface of 105 neV, from the repulsive Fermi potential of solid deuterium.

The lid of the vessel is made of 500 μm thick AlMg3 [37]. The lid was designed to withstand 3 bar of overpressure in case of loss of cooling power and rapid expansion of deuterium gas. The UCN transmission of the lid is about 50 %[53]. We are investigating the option of replacing this lid by a thinner aluminium foil supported by a grid to improve UCN transmission.

(a)

(b)
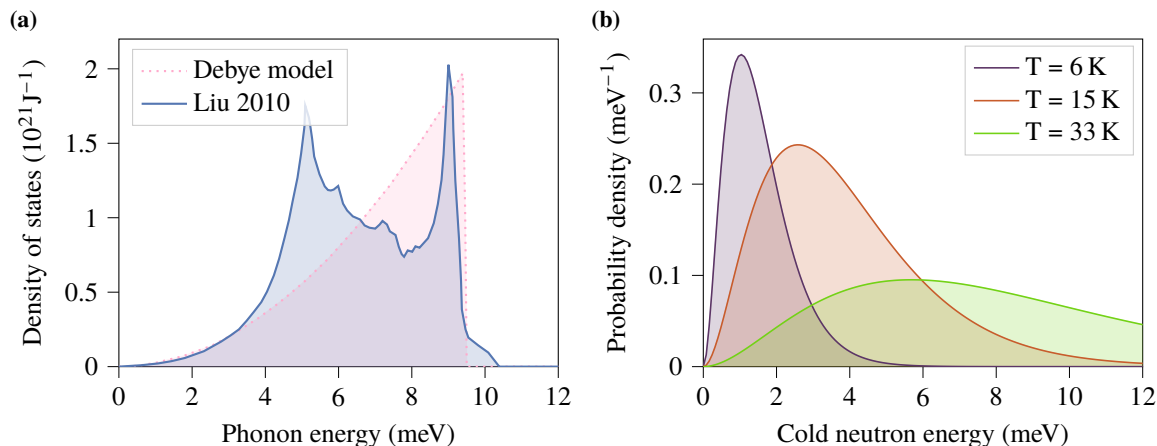
Figure 1.5: (a) The phonon density of states in solid deuterium. The two curves represent the simple Debye model, as well as a more detailed calculation from [52]. (b) The energy distribution (Rayleigh distribution) for cold neutrons at various temperatures. If a cold neutron excites a phonon state in the solid deuterium lattice which is almost exactly equal to its kinetic energy, it becomes ultracold. The UCN production rate depends on the phonon density of states in solid deuterium, the temperature of the cold neutron flux, and the interaction kinematics. UCN production is the highest when the cold neutron flux has a temperature of 33 K.

1.3.4 Storage and delivery of UCN

The UCN exiting the moderator vessel are transported upwards by the 1 m long vertical UCN guide (5 in Figure 1.2), losing the 105 neV they gained when leaving the solid deuterium. They pass through the open flapper valve (6), and enter the storage vessel of the UCN source (7). The storage vessel is made of aluminium plates coated with diamond-like carbon (DLC), which has a Fermi potential of 235 neV [37].

The flapper valve consists of two AlMg3 flaps, named after the two neighbouring villages, Böttstein and Villigen. They were first coated with nickel and afterwards with DLC. The flaps are opened by pressurising a pneumatic cylinder with cold helium, which takes about two seconds. When the flaps are closed, teeth attached to the flap sink into a magnetic field. This eddy current braking ensures smooth closing. Closing the flap takes 1.2 s. Under normal operation, the flapper valve starts closing at the end of the pulse and reopens about 20 second before the next pulse arrives [37]. In 2019 a leak appeared in the pneumatic system of the Böttstein flap. When it is opened, helium flows into the storage vessel. We now keep the Böttstein flap permanently closed and operate the Villigen flap normally [48].

Two beam guides, named West-1 (9) and South (10), are attached to the bottom of the storage vessel. A third beam guide, West-2 (11), is attached two metres higher than the other two [37]. Only UCN with an energy of at least 200 neV, measured at the bottom of the storage vessel, can reach beam guide West-2. The beamlines are around 8 m long. They have a bend in the middle to ensure there is no straight path through the biological shielding. Beam guide shutters (8) (*Neutronenleiterklappe*, NLK) are installed where the guides connect to the storage vessel. The NLKs can close off one beamline to provide a higher UCN intensity in the others.

1.3.5 The deuterium system

The deuterium system has to be able to maintain the solid deuterium in the moderator vessel at a constant temperature. The system has to safely store all 6 kg of deuterium in gaseous form at room temperature when the UCN source refrigeration plant is deactivated.

A simplified schematic drawing of the deuterium system parts most relevant in this work is shown in
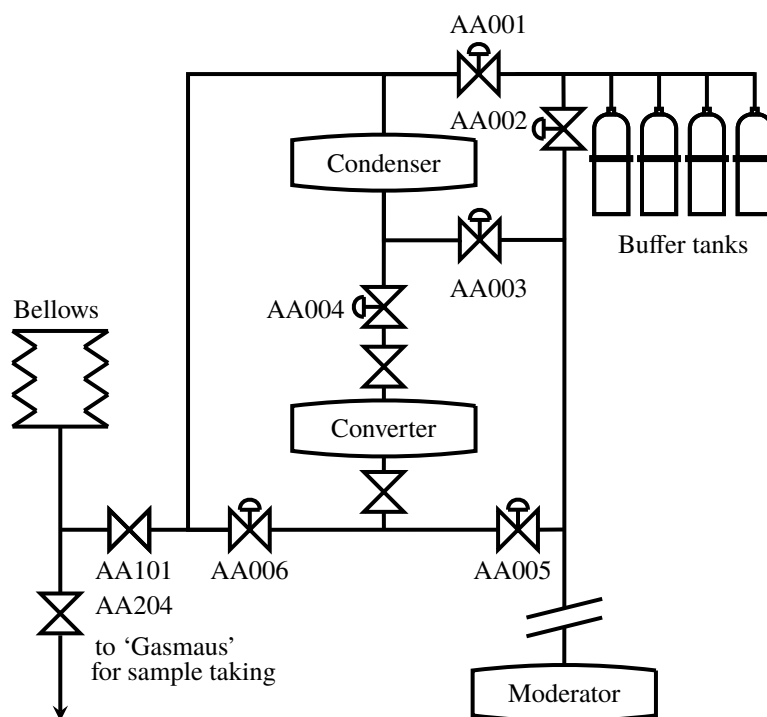
Figure 1.6: A simplified schematic of the deuterium compartments in the UCN source. When not cooled, the deuterium is stored in the storage tanks on the roof of the hall housing the UCN source. The condenser, converter, and moderator are helium-cooled. When the UCN source is started up, the deuterium is first condensed from the storage tanks into the condenser. Afterwards, it is transferred in liquid form to the converter, where para-deuterium is converted into ortho-deuterium. Finally, the deuterium is transferred in liquid form to the moderator vessel and frozen.

Figure 1.6. A detailed schematic is in Appendix B. The condenser, the converter, and the moderator are all helium-cooled using supercritical helium around 4 bar. The helium is initially cooled to 5 K. It can be reheated to 100 K, if higher temperatures are desired. The temperatures of all these compartments can be controlled precisely by controlling the heaters and the deuterium flow.

Deuterium storage

The four deuterium storage tanks are located on the roof of the building. They have a combined volume of approximately $30 \, m^3$. The deuterium is stored at ambient temperature, at approximately 1.25 bar. When the UCN source enters failsafe mode, all valves are opened automatically so that the deuterium evaporates into the tanks. The deuterium system is protected from overpressure by rupture disks.

Condensation and para-to-ortho conversion of deuterium

After the deuterium spent significant time at room temperature, the para-deuterium fraction is too high. Para-deuterium is a spin isomer which is metastable at cryogenic temperatures. It upscatters UCN and decreases the thermal conductivity of the deuterium. A more detailed explanation of the physics involved is given in Chapter 2. To ensure a high UCN yield, the para-deuterium needs to be converted to the stable ortho-deuterium.

The procedure to condense deuterium and convert it to the ortho spin isomer is as follows. First, the deuterium is solidified in the condenser vessel (43.0 L) (Figure 1.6). This is done by cooling the vessel to 5 K. Over about one day, all the deuterium solidifies in the condenser. When the solidification is completed, the valve AA001 is closed, closing off the condenser vessel from the storage tanks.

Next, the deuterium has to be transferred to the converter vessel, where the para-to-ortho conversion takes place. The condenser is heated to about 20 K to melt the deuterium. By opening valve AA004, the deuterium drips down into the converter vessel (47.5 L), where it is kept in the liquid state. The conversion from para-deuterium to ortho-deuterium is catalysed by OXISORB® [31], [54], which consists of paramagnetic chromium trioxide infused to silica gel [55]. This provides inhomogeneous magnetic fields that dephase the deuterium nuclear spins, catalysing the transition from para-deuterium to ortho-deuterium, and vice versa. Over the course of two days, the para-deuterium concentration reaches the equilibrium value (2 % to 3 %).

The converter and condenser vessels are located in the cryobox, on top of the UCN source, outside the biological shielding.

Preparation of the moderator

Finally, the deuterium with a high ortho concentration is transferred to the moderator vessel. Valve AA005 is opened, allowing the deuterium to drip down into the moderator vessel through a long vertical pipe. This process takes about a day. Once the transfer is complete, the helium temperature is decreased to 18 K, just below the triple point of deuterium (18.7 K [56]). This starts the solidification. Once all the deuterium is solid, the helium temperature is gradually reduced to 5 K. Slower solidification and cooling of the deuterium leads to fewer defects in the solid and ensures a higher UCN yield [48].

1.4 The n2EDM experiment

The n2EDM experiment is currently being commissioned at PSI, to search for the nEDM with unprecedented sensitivity [24]. The n2EDM experiment uses Ramsey's separated oscillating field method (see Section 1.2.3) to measure the precession frequency of neutrons in (anti)parallel magnetic and electric fields.

The n2EDM experiment is a double chamber experiment: UCN are stored in two chambers simultaneously, each with an identical static magnetic field, but with opposing electric fields. When the electric field is inverted by changing the polarity of the high voltage (HV) electrode, the change in Larmor frequency (for $d_n \neq 0$) is opposite in both chambers. Several systematic and statistical effects are related to vertical field gradients. For example, differences in UCN energies cause vertical striation of UCN [57] which causes the UCN to sample the magnetic field differently. Therefore, vertical field gradients lead to gravitational depolarisation [58], [59]. In a vertical double chamber experiment, the vertical magnetic field gradient can be estimated and compensated. Since the precession frequencies of UCN in parallel and antiparallel electric and magnetic fields are measured simultaneously, half the number of Ramsey cycles are required to obtain a single nEDM estimate compared to a single chamber experiment. If a proton beam pulse is missed during a Ramsey cycle, fewer data have to be discarded. [23], [24], [60].

A schematic overview of the experiment is shown in Figure 1.7. The central stack, including the precession chambers, is shown in 1.8.

1.4.1   The voyage of UCN through the experiment

As far as UCN transport is concerned, a cycle of the n2EDM experiment can be divided into three phases: filling, storage and counting. UCN are filled from the UCN source into the precession chambers during the filling phase. The filling duration is empirically optimised (50 s in 2023). During the storage phase, the Ramsey pulse sequence is applied to the UCN. This includes a 180 s period of free precession. In the counting phase, UCN leave the precession chambers. They are guided towards the U-shaped spin analysers. Magnetised spin analysing foils selectively transmit or reflect spin up and spin down UCN. The UCN of each spin state, from each chamber, are counted in four UCN detectors.

Filling the chambers with UCN

- UCN are produced in the UCN source (Figure 1.2). NLK West-1 is closed to ensure a high UCN intensity at beamport South, where the n2EDM experiment is situated.

- The superconducting magnet (SCM) provides a 5 T magnetic field. This creates an energy splitting of the neutron spin states, of ± 300 neV. The high-field-seeking UCN pass the magnet, while the low-field-seeking UCN do not pass.[6]

- A Y-shaped UCN guide divides the UCN over two beamlines, which transport the UCN to the top and bottom precession chambers. The internal diameter of the UCN guides before the Y-shaped guide is 18 cm. The internal diameter of the two UCN guides after the Y-shaped guide is 13 cm, with approximately half the cross-sectional area of the 18 cm guides.

- The UCN arrive at the switch. The switch consists of two switchboxes, for the top and bottom beamlines. Each switchbox contains a bent and a straight UCN guide, mounted on a plate. This plate can be moved, inserting either the bent or the straight UCN guide into the beam path. During the filling phase, the straight guide is inserted, creating a path between the SCM and the horizontal UCN guides (Figure 1.9).

- The UCN travel through the horizontal UCN guides. They enter the magnetically shielded room (MSR), and afterwards the vacuum chamber. The UCN shutters, located in the centre of the ground electrodes of both precession chambers (Figure 1.8), are open during the filling phase, allowing UCN to enter the chambers.

---

[6]While low-field-seeking neutrons with energies above 300 neV can pass the SCM, these have energies that greatly exceed the Fermi potentials of the coatings of the UCN guides system in n2EDM. They cannot be stored and will have been lost long before the counting phase.
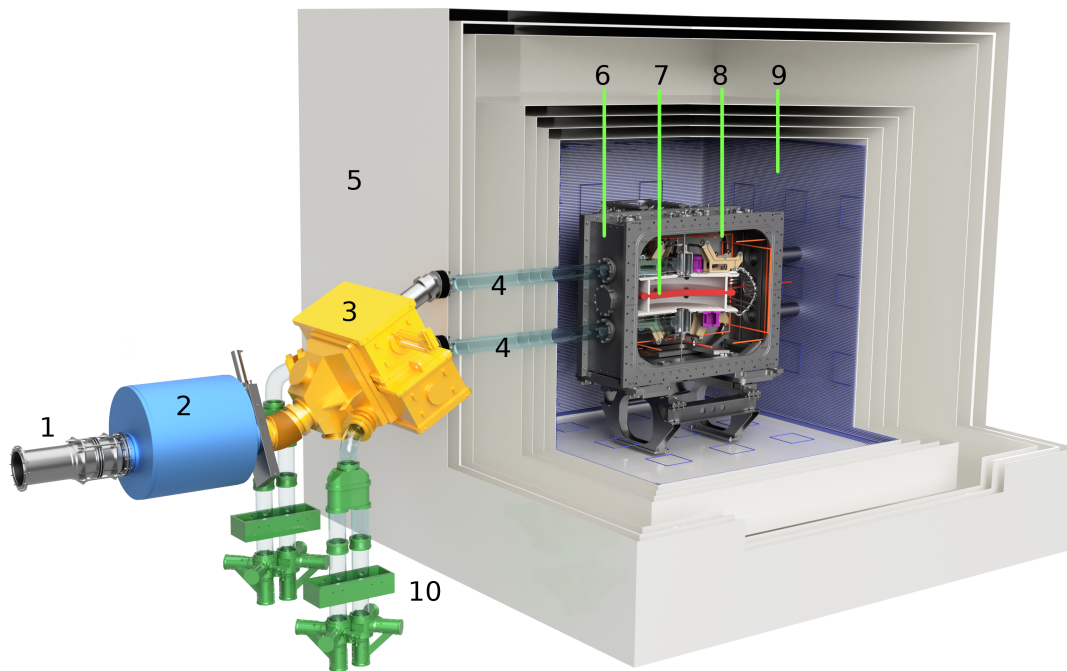
Figure 1.7: A CAD drawing of the n2EDM experiment. Neutrons from the UCN source arrive from (1) the beam guide. (2) The 5 T spin polarising superconducting magnet only permits high-field-seeking UCN to pass. The UCN travel through (3) the switch and then through (4) the horizontal UCN guides. They enter (5) the mu-metal magnetically shielded room, and (6) the vacuum chamber. The precession chambers are inside (7) the central stack, (see Figure 1.8). This is where Ramsey's method of separated oscillating fields is applied. (8) In red: the RF coils, for generating the spin flip pulses for the Ramsey sequence. (9) The B_0 coils generate the static magnetic field. After the Ramsey experiment, the UCN travel back through the horizontal UCN guides and the switch, into (10) the U-shaped spin analysers, with UCN detectors. (Image from [24].)
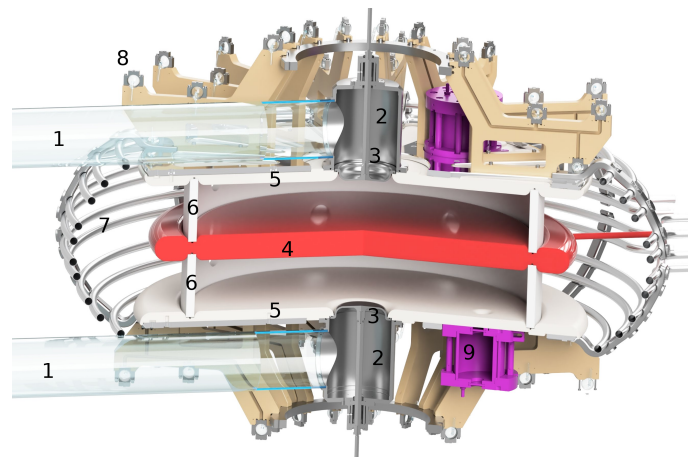
Figure 1.8: A CAD drawing of the n2EDM precession chambers. (1) The last of the four UCN guides (HG4). (2) The UCN shutter housing, containing (3) the UCN shutter. The precession chambers are enclosed by (4) the HV electrode, (5) the ground electrodes and (6) the insulating rings. (7) The ground cage contains the electric field. (8) The caesium magnetometer cells measure the absolute magnetic field strength. From the absolute field strengths at their positions, the field gradients in the precession chambers are estimated. (9) The mercury polarisation cells. (Image from [24].)
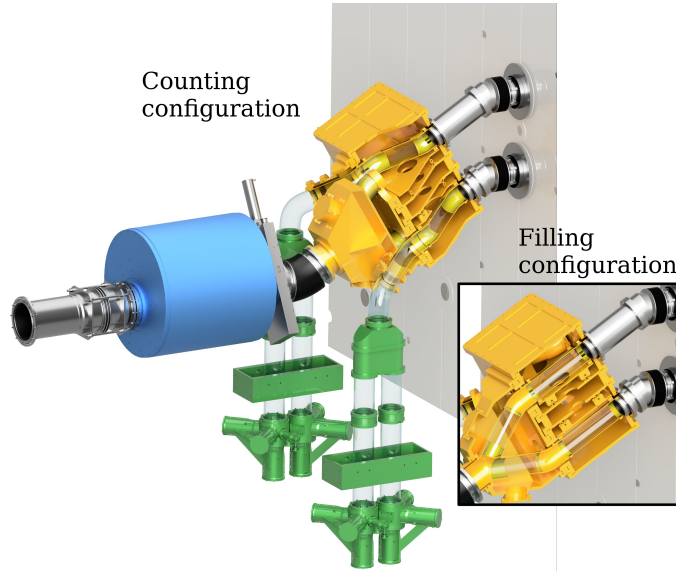
Figure 1.9: A CAD drawing of the switch. It mechanically swaps UCN guides. In the filling configuration, there is a path from the SCM to the precession chambers. In the counting configuration, there is a path from the precession chambers and the spin analysers to the UCN detectors. (Image from [24].)

- When equilibrium is obtained between the ingoing and outgoing UCN, the UCN shutters are pneumatically closed, sealing off the precession chambers. The optimal timing is to be determined empirically but will be on the order of 30 s to 50 s after the start of the proton beam pulse. The storage phase starts.

UCN storage and the Ramsey pulse sequence

The UCN remain stored in the precession chambers (Figure 1.8) during the storage phase. The precession chambers are defined by a central high-voltage electrode, two ground electrodes, and two insulating rings that separate the high-voltage and ground electrodes. The chambers have a diameter of 80 cm, and a height of 12 cm. The high voltage electrode is charged to ± 180 kV, providing homogeneous, vertical electric fields in both precession chambers of ± 15 kV cm^{-1}, with opposite field directions. $\mathbf{B}_0$ coils provide a highly homogeneous, vertical, static magnetic field ($\mathbf{B}_0$) of 1 µT, which is identical in both precession chambers.[7]

- A vapour of ^{199}Hg atoms is polarised in the direction of the static magnetic field using optical pumping, in the Hg polarisation chambers. This vapour is released into the precession chamber. ^{199}Hg has a nuclear spin of $\frac{1}{2}$.

- The RF coils (Figure 1.7) are used to generate a magnetic field pulse, rotating in the horizontal plane, with a frequency identical to the precession frequency of the Hg atoms about $\mathbf{B}_0$ (approximately 8 Hz for $|\mathbf{B}_0| = 1$ µT). The field strength and duration of the pulse are chosen so that the Hg spin state rotates away from the vertical axis by $\frac{\pi}{2}$ rad. Hence, this pulse is called the mercury $\frac{\pi}{2}$-pulse.

- The Hg spins precess in the horizontal plane, ideally for the rest of the storage phase (limited by the Hg spin relaxation time, or T_2 time [62], [63], and the leakage rate of mercury from the chambers). Their spin state is probed to measure and correct for the magnetic field inside the precession chambers. Circularly polarised laser light is shone through the precession chambers. The transmitted light intensity is modulated by Hg spin projection on the travel direction of the light. The intensity of the transmitted light is recorded using a photosensitive diode. The Larmor frequency of the mercury atoms is extracted from this signal. [64]

---

[7]For reference, the geomagnetic field strength in Switzerland is approximately 47 µT [61].

- The UCN $\frac{\pi}{2}$-pulse is applied. The frequency is approximately equal to the UCN precession frequency about the static magnetic field (about 30 Hz). This rotates the UCN spin away from the vertical axis, into the horizontal plane.

- During a free precession interval of $T = 180$ s, the UCN and the Hg spins precess about $\mathbf{B}_0$, in the horizontal plane, accumulating phase.

- A second UCN $\frac{\pi}{2}$-pulse is applied, with identical frequency and in phase with the first UCN $\frac{\pi}{2}$-pulse. In the resonance condition, the rotation frequency of the pulses exactly matches the UCN precession frequency, so that no phase difference builds up between the UCN spin state and the $\frac{\pi}{2}$-pulses. In this case, the second pulse completes the UCN spin flip with perfect efficiency, so that the final UCN spin is rotated by π relative to the initial spin state. If by the time the second pulse starts, the pulse and the UCN spin state are out of phase by $\pm\frac{\pi}{2}$, the pulse will be ineffective at changing the vertical component of the spin state. By varying the amount of detuning, the resonance frequency can be obtained using a fit.

- After the second UCN $\frac{\pi}{2}$-pulse, the UCN shutters open and the storage phase is over.

The counting phase

During the Ramsey cycle, the UCN spin state is manipulated. In contrast to the mercury spin state, the UCN spin state can not be observed directly via the modulation of transmitted laser light. The UCN are separated into different guides for up and down spin states, and counted in UCN detectors.

- The UCN travel back through the UCN shutters, through the horizontal guides, and arrive back at the switch.

- The switch is in counting position: the bent guide is inserted in the switch, creating a path from the precession chambers to the U-shaped spin analysers (Figure 1.9).

- In the U-shaped spin analysers (Figure 1.10), UCN are divided over two legs.

- These legs contain spin analysing foils. These are aluminium foils coated with a thin layer of Iron. The iron is magnetised to saturation (2 T). The potential energy of UCN inside the foil is equal to $V_{F\,\text{iron}} \pm \mu_n B \approx (210 \pm 120)$ neV. UCN with a kinetic energy $E_{\text{kin}} < 90$ neV are not transmitted. UCN with $E_{\text{kin}} > 330$ neV can be transmitted regardless of their spin state. Between these two energies, only spin-down neutrons are transmitted. The insulator rings are coated with deuterated polystyrene, with a Fermi potential of 165 neV. This is the upper limit for the kinetic energy of UCN stored in the precession chamber. To ensure that 90 neV $< E_{\text{kin}} < 330$ neV when the UCN reach the analysing foils, the foils have to be positioned between 90 and 165 cm lower than the precession chambers [24], [65].

- The adiabatic spin flippers are installed above the spin analysing foils. These provide oscillating magnetic fields, turning spin-up into spin-down neutrons and vice versa. In the leg where spin-up neutrons are to be detected, the spin flipper is switched on. By varying which leg is counting spin-up or spin-down neutrons, systematics related to analysing foil or detector efficiencies are cancelled out.

- After passing through the spin analysing foils, UCN are counted in the fast, gaseous UCN detectors (GADGET [66], [67]). UCN pass through a 30 μm thick aluminium entrance foil, into a chamber containing ^{3}He and CF_4 at partial pressures of 15 and 500 mbar, respectively. Neutron capture causes the reaction $^3\text{He} + n \rightarrow {}^1\text{H} + {}^3\text{H}$. The reaction products induce scintillation light in the CF_4, which is picked up in the three photomultiplier tubes in coincidence. UCN signals are distinguished from background events using pulse shape analysis.

The Ramsey fit

In the counting phase, the UCN populations in the spin-up and spin-down states are counted for both the top and bottom precession chambers. For each chamber, the asymmetry is calculated:

$$A = \frac{N_u - N_d}{N_u + N_d}, \tag{1.11}$$

where $N_{u,d}$ represent the counted number of spin-up and spin-down UCN, respectively. Asymmetry is defined on the interval $[-1, +1]$. The $\frac{\pi}{2}$-pulse frequency ω is varied over several cycles. The resonance frequency ω_0 is obtained by fitting the function

$$A = \alpha \cos\left((\omega - \omega_0)T_{\text{eff}}\right) \tag{1.12}$$

for both chambers, where T_{eff} is the effective precession time, equal to the free precession interval duration, plus a function of the UCN $\frac{\pi}{2}$-pulse duration. The visibility or amplitude of the curve, α, depends on the polarisation of the UCN after the Ramsey cycle, and the efficiencies of the spin flippers, spin analysers and UCN detectors. The statistical uncertainty associated with a single cycle equals [24]

$$\sigma(\omega_0) = \frac{1}{\alpha T_{\text{eff}}\sqrt{N}} \left(1 - \left(\frac{A}{\alpha}\right)^2\right)^{-\frac{1}{2}}. \tag{1.13}$$

At least four cycles are required to obtain a single fit for ω_0. In each of these cycles, ω is set to one of the four working points:

$$\omega = \omega_0 + \Delta\omega, \tag{1.14}$$

where ω_0 is the expected precession frequency, with

$$\Delta\omega = \pm\frac{\pi}{2T_{\text{eff}}} \pm \delta\omega. \tag{1.15}$$

The term $\pm\pi/(2T_{\text{eff}})$ ensures that the working points are approximately half a fringe width removed from the resonance, where asymmetry is close to zero, thus minimising σ in Equation 1.13, while $\pm\delta\omega$ is a small additional detuning, required to obtain four points.

1.4.2   Magnetic fields

The vacuum chamber housing the central stack is enclosed in a five-layer, cubic, mu-metal magnetically shielded room (MSR). The shielding factor is 10^5 in all directions for fields oscillating with 0.01 Hz and rises rapidly with frequency, reaching 10^8 at 1 Hz. The internal volume is a cube with sides of 2.93 m [24], [68].

The MSR is situated in the thermohouse, an air-conditioned wooden structure with an internal volume of approximately 1000 m^3. The temperature variation in the thermohouse is kept below 1 K all year round (below 0.1 K for the innermost layer of the MSR). This precludes significant thermal expansion or thermal currents in the MSR, which would alter the magnetic field lines inside.

An active magnetic shield (AMS) [69], [70] is installed on the interior side of the thermohouse walls. The AMS consists of eight dedicated coils, with a combined length of 50 km. The AMS dynamically compensates ambient magnetic fields that permeate the area. These fields come from surrounding machines and experiments. A particular risk is posed by the neighbouring superconducting magnet test facility, SULTAN [71], which operates with magnetic fields up to 11 T. The AMS can compensate fields of $\pm 50\,\mu$T in all directions and field gradients of $\pm 5\,\mu$T m^{-1} in the sub-hertz frequency range.

The magnetic fields in the precession chambers are required to be uniform within 100 pT, 0.01 % of the 1 μT static magnetic field. Trim coils are installed inside the MSR to provide precise control over the magnetic field. In addition, the field is required to be stable within 30 fT over 300 s, the duration of one measurement cycle.
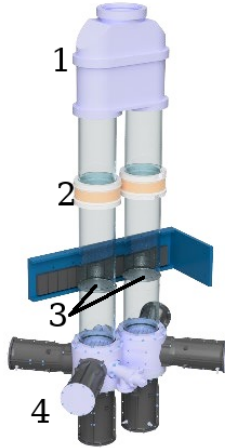
Figure 1.10: A CAD drawing of one of the two simultaneous U-shaped spin analysers (USSA). (1) A NiMo-coated aluminium bifurcation divides the UCN over two vertical arms. (2) Adiabatic spin flippers provide an oscillating magnetic field, rotating the spin of UCN by π rad, if they are switched on. (3) Spin analysing foils only permit spin-down UCN to pass through. In the arm where the spin flipper is engaged, spin-up UCN are converted to spin-down UCN, and are able to pass the foil. UCN that do not pass the foil are reflected upwards. They can enter the other arm of the USSA and be counted. (4) The GADGET UCN detectors count the spin-up and spin-down UCN. (Image from [24].)

1.4.3 The UCN system

To ensure a high statistical sensitivity of the experiment, the UCN guide system must be able to store UCN for a long time, transmit UCN efficiently, and preserve the UCN spin state.

To ensure a long storage time, a tight UCN system is required without gaps between UCN guide segments. The inner surfaces of the UCN system must have a low 'loss-per-bounce rate'. This can be achieved through using surface coatings with a high Fermi potential and a low loss coefficient.

To transmit UCN efficiently it is not only necessary to prevent the loss of UCN but also to preserve their direction of travel. In a guide system with rough surfaces, UCN change direction often and diffuse slowly through the system. If we model the movement of UCN using the diffusion equation, the expected distance travelled of the UCN scales with the square root of time. In a smooth guide system, the UCN momentum along the axis of UCN guides is preserved and the travelled distance scales linearly with time. It follows that UCN guides must be smooth to ensure good UCN transmission.

To prevent UCN depolarisation, the materials should not contain strong local magnetic dipoles that disrupt the uniformity of the UCN guide fields. This becomes even more important near the central stack: local dipoles reduce the effectiveness of the caesium magnetometer array to correct for field gradients, contributing to systematic effects. In the horizontal UCN guides closest to the precession chamber (HG4), no dipoles are allowed larger than $15.5 \, \text{nA} \, \text{m}^2$ [72].

To satisfy these requirements, the straight tubular segments of the UCN guide system are made from borosilicate glass, sputter coated on the inside with nickel-molybdenum (85:15 mass ratio, $V_F = 220 \, \text{neV}$ [37]). Nickel has a high fermi potential (252 neV [30]), but is ferromagnetic ($T_C = 628 \, \text{K}$, saturation magnetisation = $4.0 \, \text{A} \, \text{m}^{-1}$). In NiMo alloys, the Curie temperature and the saturation magnetisation decrease supralinearly with the molybdenum concentration, vanishing between 12 and 13 % (by mass) [73].

Corners or bends in the UCN guide system are often made from machined aluminium, which is hand-polished to decrease the surface roughness. These aluminium pieces are coated with NiMo.

The precession chambers are made of aluminium electrodes, separated by insulating rings. The electrodes are coated with DLC ($V_F = 235 \, \text{neV}$, Table 1.1 [37]). The insulator rings are made of Rexolite (polystyrene),

coated with deuterated polystyrene (V_F = 165 neV [74]). Currently, we are testing new insulator rings made out of quartz, which can be coated with deuterated paraffin (V_F = 180 neV [75]).

1.5 The UCN transport simulation code MCUCN

The UCN group at PSI uses neutron transport simulations to better understand UCN transport through the UCN source and the experiments. The framework used to this end is MCUCN, short for Monte Carlo UCN [76]. MCUCN was developed during the design of the phase space transformer at SINQ[77], and further developed during the design phase of the UCN source at PSI. The results from MCUCN have been cross-checked with other UCN transport simulation codes, such as STARUCN [78], PENTRACK [79], and GEANT4UCN [80]. MCUCN is written in C++. MCUCN derives its speed from the analytical calculation of trajectories between reflection points, as opposed to GEANT4UCN, for example. The UCN Group at PSI uses MCUCN extensively to model and optimise new experimental setups before they are built.

In MCUCN, surfaces are defined by their geometry (approximated by quadrics, three-dimensional generalisations of conic sections) and their surface parameters: their real Fermi potential V_F; the loss coefficient η; and the diffuse scattering fraction d. The loss coefficient is related to the probability a neutron is absorbed by the wall or inelastically scattered instead of reflected [30]. This affects the 'loss-per-bounce' rate (see Figure 1.1). The parameter d specifies the probability for each reflection that a neutron is not scattered specularly, but diffusely. It is an effective parameter, related to the roughness and material properties of the surface.

Diffuse scattering is simulated according to Lambert's cosine law [81]. Under Lambert scattering, the angle of reflection θ is uncorrelated to the angle of incidence. The probability density for scattering with an angle of reflection θ, and an azimuthal angle ϕ is

$$\rho(\theta, \phi) = \frac{1}{\pi \text{ sr}} \cos\theta \sin\theta. \tag{1.16}$$

The Lambert model is a good model for perfect diffuse reflection. The hybrid model of specular reflection and Lambertian reflection is a simplification. Better models exist for reflection that is neither perfectly specular nor perfectly diffuse. Work is currently underway on the implementation of microfacet models. In microfacet models, the surface is simulated as a collection of small surfaces (microfacets) with an orientation following a given distribution. These models were developed originally for computer graphics [82] and have been successfully used in UCN simulations [83].

In this work, MCUCN was used to better understand some measurement results with UCN. The entire UCN system is simulated, from the solid deuterium moderator to the UCN detectors.

PART I

THE UCN SOURCE AT PSI

CHAPTER 2

THE UCN SOURCE DEUTERIUM COMPOSITION

2.1 Overview

Para-deuterium (p-D_2) and hydrogen deuteride (HD) in the UCN source moderator respectively upscatter or capture UCN at a significantly higher rate than ortho-deuterium (o-D_2). Their presence in the deuterium moderator lowers the UCN extraction from the moderator and thus the UCN yield. In this chapter, I explain how we used Raman spectroscopy to study the fractions of para-deuterium and hydrogen deuteride in the UCN source deuterium.

We define the para-deuterium fraction as the fraction of deuterium molecules that are para-deuterium, ignoring other substances. We found that over the year, the para-deuterium fraction decreases from about 2.5 % to about 0.5 %.[1]

We determined that the ortho-to-para conversion rate in the gas bottle used for taking Raman samples corresponds to a lifetime of (157 ± 6) d.

The para-to-ortho conversion rate in the condenser vessel at 6 K corresponds to a lifetime of (100.3 ± 2.5) d.

Raman spectra obtained from the deuterium of the UCN source were compared with a new calibration sample with an HD concentration known to an accuracy of 0.2 %. The molar HD concentration in the UCN source was determined to be $(0.210 \pm 0.002_{\text{sys}} \pm 0.016_{\text{stat}})$ %. This corresponds to an atomic hydrogen concentration of (0.105 ± 0.008) %. At this concentration, HD is responsible for 14 % of UCN absorption and upscattering losses inside the moderator, and is not the dominant loss channel.[2] The gains in UCN yield that could be made by replacing the deuterium by a deuterium with a lower HD concentration are limited and are likely not worth the cost and effort.

The value we obtained does not significantly differ from the previous estimate of 0.18 % in 2019 [84].

---

[1]All concentrations in this chapter are molar concentrations unless stated otherwise.
[2]This assumes the absence of high Z contamination.

2.2   Introduction

2.2.1   UCN loss channels in the solid deuterium moderator

The ideal neutron moderator contains elements with an atomic mass close to the neutron (more energy lost in scattering), high scattering cross-sections, and a low capture cross-section for neutrons. Deuterium is used as a neutron moderator in the UCN source. It is the second-lightest nuclide (2.014 Da), and its neutron capture cross-section (0.519 mb at 2200 m s^{-1}) is 640.8 times lower than that of hydrogen (332.6 mb at 2200 m s^{-1}) [56], [85].

In this chapter, we measured the relative molar concentrations of hydrogen and deuterium. Tritium has a negligible natural abundance (1×10^{-17} [86]), but is produced in the UCN source by neutron capture in deuterium. While the tritium concentration in the UCN source deuterium is so low that it cannot be detected using our Raman spectroscopy setup, the tritium production rate has been estimated using simulations, and activity measurements. The tritium activity in the UCN source was measured to be (6100 ± 1400) kBq g^{-1}, on 2018-08-06 [84]. This corresponds to a tritium molar concentration of $(5.5 \pm 1.3) \times 10^{-9}$. Tritium itself does not absorb UCN, but it decays to ^{3}He, which has a neutron capture cross-section of 5333 b, which is 39 % higher than that of ^{10}B [85]. Fortunately, with a boiling point of 3.2 K, ^{3}He remains in the gaseous phase above the surface of the solid deuterium and is pumped off. In this chapter, the presence of ^{3}H and ^{3}He in the source are neglected. The possible presence of high-Z contamination, which our method of sample taking using Raman spectroscopy is not sensitive to, is also neglected.

Neutron capture

From the microscopic capture cross-section σ_c, the macroscopic cross-section can be found by multiplying with the number density, $\Sigma_c = n\sigma_c$. The capture rate, Γ_c, depends on the amount of material the particle traverses per unit time, $\Gamma_c = vn\sigma_c$. Capture cross-sections are inversely proportional to the velocity. Once measured at a velocity v_0 (usually 2200 m s^{-1}), the cross-section at other velocities can be found using $v\sigma_c(v) = v_0\sigma_c(v_0)$. While the UCN reside in the material, the capture rate is velocity-independent:

$$\Gamma_c = v_0\, n\, \sigma_c(v_0). \tag{2.1}$$

The total capture loss rate depends on the concentration of atomic hydrogen, as well as on contaminations with other substances. At low atomic hydrogen concentrations, hydrogen is mainly present as hydrogen deuteride (HD). We assume that if the HD concentration is small, it does not significantly change the density of the deuterium so that we can use the D$_2$ molar volume (19.93 mL mol^{-1} [56]) to calculate the capture loss rate:

$$\Gamma_c^{HD} = 2156\,\text{s}^{-1}, \tag{2.2}$$

$$\Gamma_c^{D_2} = 6.92\,\text{s}^{-1}. \tag{2.3}$$

Thermal upscattering

When passing through a material, UCN are lost if they gain energy from the material, raising their total energy above the Fermi potential of the storage vessel. This is called upscattering. In thermal upscattering, the neutron interacts with an excited phonon state in the lattice, transferring the phonon energy to the UCN, the reverse process of superthermal UCN moderation. UCN can also incoherently scatter from a single nucleus. The rate depends on the density of excited phonons and thus has a strong temperature dependence [87]. An empirical polynomial fit to the thermal upscattering rate in deuterium was provided in [88], [89], valid from 5 K to 12 K:

$$\Gamma_u^\theta(T) = 3.02 \left(5.931 - 2.168\,\text{K}^{-1}\,T + 0.182\,\text{K}^{-2}\,T^2 + 0.0187\,\text{K}^{-3}\,T^3\right)\,\text{s}^{-1}. \tag{2.4}$$

It is estimated that during an 8 s pulse, the mean temperature of the solid deuterium moderator rises from 5.3 K to 7.3 K. This is based on measurements of the UCN yield as a function of proton beam pulse duration [84] and models based on the heat capacity and thermal conductivity of deuterium [48]. The effective temperature of the moderator during an 8 s pulse is 6.3 K. This means that if the temperature stayed constant at 6.3 K, the UCN source would produce the same UCN yield as with a varying temperature from 5.3 K to 7.3 K. At 6.3 K, the thermal upscattering loss rate is $12.8\,\text{s}^{-1}$.

Spin-flip conversion upscattering in para-deuterium

Deuterium exists in two spin isomers, ortho and para-deuterium, which respectively have even and odd rotational quantum numbers J. Para-deuterium is metastable at 5 K, frozen in the $J = 1$ state. A UCN can interact with it with its magnetic moment, turning it into ortho-deuterium. The energy released in this $J = 1 \rightarrow J = 0$ transition (7.41 meV) is transferred to the UCN. The lifetime of UCN in p-D_2 was measured to be 1.5 ms [90], implying a loss rate of

$$\Gamma_u^{\text{para}} = 667\,\text{s}^{-1}. \tag{2.5}$$

Control of the UCN loss channels

The overview of the UCN loss channels is shown in Figure 2.1. To get the best UCN yield out of the UCN source, the UCN losses must be kept to a minimum. This means that the para-deuterium fraction, the HD concentration, and the temperature of the deuterium in the moderator need to be kept as low as reasonably possible.

The capture losses in deuterium cannot be mitigated and put a hard lower bound on the loss rate.

The thermal upscattering rate depends on the temperature, which in turn depends on the heat deposited by the proton beam pulse, the cooling power, and the thermal conductivity of the deuterium. Reducing the duration or intensity of the proton beam pulse would reduce the neutron production and is not an option. The thermal conductivity depends on the number of defects in the solid and the para-deuterium concentration (Figure 2.2).

For HD and p-D_2, we are satisfied when the loss rates due to HD capture and p-D_2 upscattering are not dominant. If their respective loss rates are below the capture loss rate in deuterium, further reduction of their concentration yields diminishing returns, and would not be worth the investment. In practice, that means that the HD concentration should be below 0.3 %. The HD concentration was previously measured by N. Hild to be 0.18 % [84]. We have confirmed this measurement within 2σ using a new calibration gas mixture.

For para concentrations below 1 %, the para-to-ortho conversion upscattering rate is smaller than the capture rate in D_2. The para-deuterium concentration varies over the year. During the winter shutdown and maintenance of the HIPA accelerator complex, the deuterium in the moderator is evaporated, and stored in storage tanks. The para-deuterium fraction increases due to natural conversion. During the startup of the UCN source in spring, the deuterium is first liquefied in the condenser and then transferred to the converter. The converter is filled with oxisorb®, which catalyses the (naturally slow) conversion between para and ortho-deuterium [54], until it reaches equilibrium at 21.5 K (97.3 % ortho, see Section 2.2.2). Over the year, the ortho fraction increases to about 99.5 %, as a result of natural conversion, and radiation-induced conversion [84]. See Appendix B for a schematic overview of the helium, deuterium and UCN systems of the UCN source.
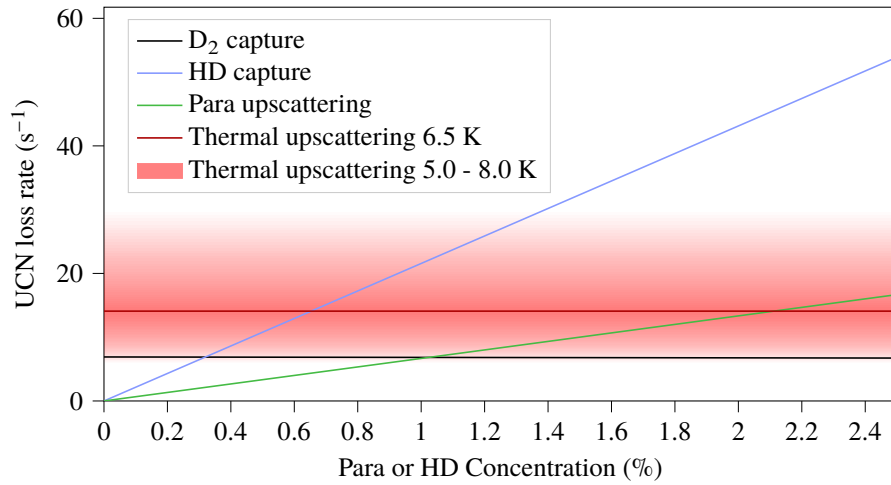
Figure 2.1: Overview of the different UCN loss channels in the solid deuterium moderator of the UCN source. UCN capture in D_2 is inevitable. Thermal losses depend on the beam current, cooling power and the para-deuterium fraction. The neutron capture rates are similar for D_2 and HD if the HD concentration is around 0.3 %. For p-D_2 fractions around 1 %, the para-to-ortho conversion upscattering rate is similar to the D_2 capture rate. Losses due to thermal upscattering are indicated for a deuterium temperature of 6.5, and for the range 5.0 to 8.0 K.
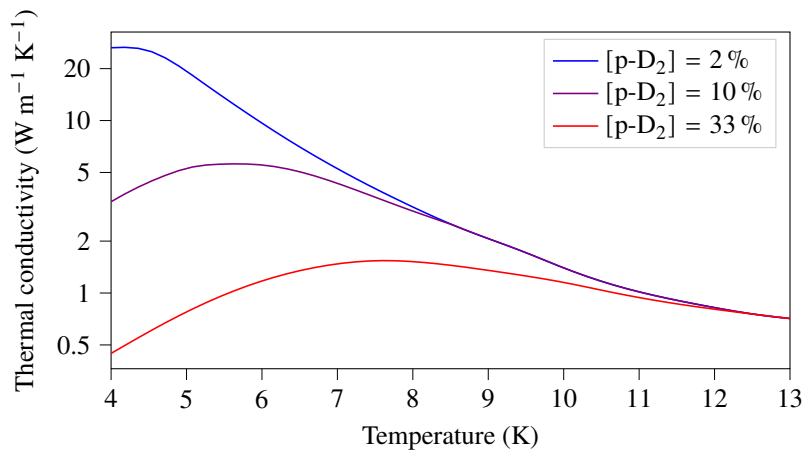
Figure 2.2: The thermal conductivity of deuterium for different para-deuterium fractions. The data are taken from [56]. Above 12 K, the thermal conductivity does not significantly depend on the ortho fraction. There is significant divergence at lower temperatures. The lower thermal conductivity caused by para-deuterium can result in higher temperatures during the pulse, and thus increased thermal upscattering.

2.2.2 Rotational and vibrational spectrum of hydrogen species

Rotation

Hydrogen isotopes form diatomic molecules. There are two main internal degrees of freedom over which diatomic molecules can distribute energy: rotation and vibration (i.e. oscillation of the internuclear separation). The interaction between rotational and vibrational state is a higher-order effect we can safely neglect. The Hamiltonian for the internal energy of a diatomic molecule is

$$\hat{H} = \hat{H}_{\mathrm{rot}} + \hat{H}_{\mathrm{vib}}. \tag{2.6}$$

The rotational part of the Hamiltonian depends on the angular momentum L, and the moment of inertia, I:

$$\hat{H}_{\mathrm{rot}} = \frac{\hat{L}^2}{2I}. \tag{2.7}$$

The square of the angular momentum is quantised. Its eigenstates are

$$L^2 = \hbar^2 J(J+1), \tag{2.8}$$

where J is the angular momentum quantum number. This gives us the rotational energy eigenstates

$$E_{\mathrm{rot}}(J) = \frac{\hbar^2 J(J+1)}{2I}. \tag{2.9}$$

Spectroscopists prefer to express this as a frequency:

$$F_{\mathrm{rot}}(J) = BJ(J+1), \tag{2.10}$$

with

$$B = \frac{h}{8\pi^2 cI}, \tag{2.11}$$

where h is Planck's constant and c is the speed of light in vacuum.

For a diatomic molecule, the moment of inertia around its axis is approximately zero, whereas the moment of inertia in the plane perpendicular to its axis equals

$$I = \frac{m_1 m_2}{m_1 + m_2} r^2 = \mu d^2, \tag{2.12}$$

where $m_{1,2}$ are the atomic masses, from which we can calculate the reduced mass μ, and r is the internuclear separation.

Due to centrifugal effects, the internuclear separation is larger at higher rotational states. This increases the moment of inertia and reduces the difference between consecutive energy levels. Instead of writing this as a dependence of I on J, this effect is usually absorbed in Equation 2.10 as follows:

$$F_{rot}(J) = BJ(J+1) - DJ^2(J+1)^2. \tag{2.13}$$

Because the rotational and vibrational states are not completely independent, B is a function of the vibrational quantum number v (Equation 2.11 of [56]):

$$B_v(v) = B_e - \alpha_e \left(v + \tfrac{1}{2} \right) + \gamma_e \left(v + \tfrac{1}{2} \right). \tag{2.14}$$

At room temperature, the fraction of hydrogen molecules in an excited vibrational state is negligible, as is explained in the next section. We therefore used $B_v(v=0)$ in our calculations, for the values of B_e, α_e, and γ_e given in Table 2.2 of [56].

Apart from J, the rotational state has another degree of freedom: the angular momentum projection onto a chosen axis. The associated quantum number, M_J, takes values of $-J, -J + 1, ..., J$, which means each angular momentum state has a degeneracy of

$$g(J) = 2J + 1. \tag{2.15}$$

The rotational constants, B and D, as well as the energy required to access the first rotational state, are tabulated for all hydrogen species in Table 2.1.

Vibration

Diatomic molecules only have a single mode of vibration, which is the oscillation of the internuclear separation, r, around its equilibrium position. The vibrational part of the Hamiltonian can be written as

$$\hat{H}_{\text{vib}} = -\frac{\hbar^2}{2\mu} \left(\frac{\partial}{\partial r} \right)^2 + U(r) = \frac{\hat{p}_r^2}{2\mu} + U(r). \tag{2.16}$$

Here, $\hat{p}_r$ is the radial momentum operator, and $U(r)$ is the (radially symmetric) interatomic potential. Many interatomic potentials have been studied, their exact characteristics are outside the scope of this work. Interatomic potentials approach zero as the separation approaches infinity, and they increase steeply for small separations (as the electronic shells overlap, the screening of the nuclear charges is reduced, leading to strong Coulomb repulsion between the nuclei). Any potential that permits the formation of molecules must have a minimum at a finite separation, r_0. For separations close to r_0, $U(r)$ can be approximated by the first non-constant term of its Taylor expansion around r_0, which is quadratic in r. This reduces the problem to the harmonic oscillator:

$$\hat{H}_{\text{vib}} = \frac{\hat{p}_r^2}{2\mu} + \pi m \nu_e^2 \hat{r}^2, \tag{2.17}$$

with the energy levels

$$E_{\text{vib}} \approx h \nu_e \left(v + \tfrac{1}{2} \right), \tag{2.18}$$

where v is the vibrational quantum number, and h is Planck's constant. To account for the anharmonicity of the oscillator, higher-order correction terms can be added:

$$E_{\text{vib}} \approx h \nu_e \left(\left(v + \tfrac{1}{2} \right) - x_e \left(v + \tfrac{1}{2} \right)^2 + y_e \left(v + \tfrac{1}{2} \right)^3 \right). \tag{2.19}$$

For small v, these corrections are small. For the six hydrogen species, x_e ranges from 0.015 (for T_2) to 0.028 (for H_2), and y_e from 6.2×10^{-5} (T_2) to 1.8×10^{-4} (H_2) [56].

Because of the low atomic mass of hydrogen, the energy required for the molecules to enter the first vibrational state is high ($E \gg 25\,\text{meV}$), see Table 2.1. The fraction of H_2 molecules that are in the first excited vibrational state at room temperature is thus only 6×10^{-10} (5×10^{-6} for T_2). We will therefore disregard the vibrational states in the rest of this chapter.

Spin isomers

Deuterium The nuclear spin of the deuteron (the deuterium nucleus) is 1. The total nuclear spin quantum number of the deuterium molecule, S, can therefore take values of 0, 1 and 2. Since the nuclear spin projection quantum number, M_S, can take values from the set $-S, -S + 1, ..., S$, the spin degeneracy of these spin states is respectively 1, 3 and 5.

The even total spin states are symmetric under the exchange operation, whereas the odd total spin states are antisymmetric. The same is true for the angular momentum states. Deuterons are bosons and must satisfy Bose-Einstein statistics. The total wave function, which is the product of the vibrational state (always

Table 2.1: Rotational and vibrational constants of Hydrogen species. The values were taken from Table 2.2 of [56]. $B_v(v = 0)$ was calculated using Equation 2.11 of [56].

Species	$B_v(v = 0)$ (GHz)	D (GHz)	$E(J = 1, v = 0)$ (meV)	$E(J = 0, v = 1)$ (meV)
H_2	1778.8	1.40	14.69	515.9
HD	1339.2	0.790	11.06	450.3
HT	1192.6	0.624	9.854	425.9
D_2	896.81	0.349	7.412	371.3
DT	749.02	0.250	6.191	340.3
T_2	600.79	0.154	4.967	306.0

symmetric), the angular momentum state and the spin state, must therefore be symmetric under the exchange operation. It follows that $J + S$ must be even. This gives rise to two species, or spin isomers, of deuterium:

Ortho-deuterium: $S \in \{0, 2\}$, $J \in \{0, 2, 4, ...\}$, $g(J) = 6(2J + 1)$
Para-deuterium: $S = 1$, $J \in \{1, 3, 5, ...\}$, $g(J) = 3(2J + 1)$

From the degeneracies $g(J)$, it follows that in the high-temperature limit, the ortho-deuterium fraction is $\frac{2}{3}$, and the para fraction $\frac{1}{3}$. In thermal equilibrium at 0 K, all molecules are in the ground state, $J = 0$. However, without the use of a catalyst, the time constant for spin relaxation is long: between 79 d (solid at 4.2 K) and 129 d (liquid at 20.4 K, table 24.1 of [56]). When the deuterium is cooled, a significant fraction of para-deuterium stays in the metastable state $J = 1$, with an energy of 7.4 meV above the ground state.

Hydrogen and tritium   The total nuclear spin of hydrogen, S, can take the values 0 and 1, since the proton spin is $\frac{1}{2}$. The degeneracies of these states are 1 and 3. The total wave function must be antisymmetric to satisfy Fermi-Dirac statistics. The $S = 1$ spin-triplet, which is antisymmetric for deuterium, is symmetric for hydrogen, as can be seen from the Clebsch-Gordan coefficients. To make the total state antisymmetric, $J + S$ must be even.

Ortho-hydrogen: $S = 1$, $J \in \{1, 3, 5, ...\}$, $g(J) = 3(2J + 1)$
Para-hydrogen: $S = 0$, $J \in \{0, 2, 4, ...\}$, $g(J) = 1(2J + 1)$

In the high-temperature limit, three quarters of hydrogen molecules are ortho-hydrogen. Since the tritium nucleus, the triton, is a spin-$\frac{1}{2}$ particle as well, the same reasoning applies to tritium.

The ortho and para spin isomers refer to different angular momentum states in deuterium and hydrogen/tritium. However, for hydrogen, deuterium and tritium, the ortho spin isomers always constitute a majority in the high-temperature limit, and their spin state is always symmetric.

Combined species   For the species containing different hydrogen isotopes (HD, HT, and DT), exchange symmetry is not applicable. There are no spin isomers. All combinations of S and J are allowed.

Statistical mechanics

For an ensemble of molecules in thermal equilibrium at temperature T, we can write the probability of finding a molecule in a certain rotational state as

$$P(J|T) = \frac{1}{Z(T)} \cdot g(J) \cdot \exp\left(-\frac{E(J)}{k_B T}\right), \tag{2.20}$$

where we used $E = hF$ to obtain the energy levels from Equation 2.13, k_B is the Boltzmann constant, and $Z(T)$ is the partition function:

$$Z(T) = \sum_{J=0}^{\infty} g(J) \cdot \exp\left(-\frac{E(J)}{k_B T}\right). \tag{2.21}$$
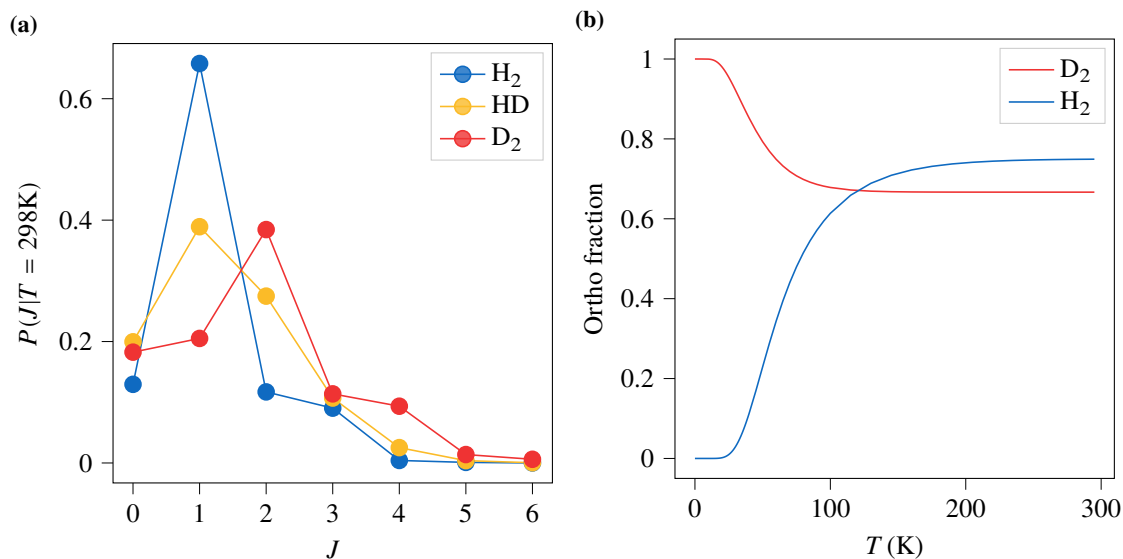
Figure 2.3: (a) The calculated population of the first seven angular momentum states at $T = 298$ K for the three most common hydrogen species. Note the 'oscillating' nature of the plots for hydrogen and deuterium, a result of the degeneracies of their spin states, and how such behaviour is absent in the plot for hydrogen deuteride. (b) The calculated ortho fraction in thermodynamic equilibrium, as a function of temperature, for hydrogen and deuterium. In the high-temperature limit, the ortho fraction approaches 3/4 for hydrogen, and 2/3 for deuterium.

For H_2, HD and D_2, the populations of the angular momentum states are shown in Figure 2.3a.

From $P(J|T)$, we can calculate the fraction of the ensemble in the ortho state, as a function of temperature, assuming thermal equilibrium:

$$[O] = \begin{cases} \sum_{n=0}^{\infty} P(2n|T) & \text{for deuterium} \\ \sum_{n=0}^{\infty} P(2n + 1|T) & \text{for hydrogen and tritium} \end{cases}. \tag{2.22}$$

The resulting ortho fractions are shown for hydrogen and deuterium, in Figure 2.3b. For deuterium in thermal equilibrium, the ortho fraction is 66.67 % at 298 K, 98.00 % at 20 K, and 99.97 % at 10 K.

2.2.3 Raman spectroscopy

In Raman spectroscopy, the composition of a substance is investigated by using laser light to probe its rotational and vibrational spectrum.

Raman transitions

When a photon interacts with a molecule, the electric field component of the photonic field induces a transient dipole moment in the electronic shell, with which the photon can interact. The photon can be absorbed by the molecule, whose rotational or vibrational state can be modified, whereafter a new photon is emitted. This is called the Raman effect, named after the discoverer [91]. In Stokes Raman scattering, the molecule absorbs energy, and the outgoing photon has less energy than the incoming photon. In anti-Stokes Raman scattering, the molecule loses energy, the outgoing photon has more energy than the incoming photon. If the molecular state and photon energy are unchanged, the process is called Rayleigh scattering.

Table 2.2: The first Stokes Raman transitions of hydrogen, hydrogen deuteride and deuterium. The values were calculated using rotational and vibrational constants from [56].

Species	Transition	Spin isomer	ΔE (meV)	Raman shift (cm^{-1}).
H_2	$J = 0 \rightarrow J = 2$	para	43.93	354.3
	$J = 1 \rightarrow J = 3$	ortho	72.76	586.8
	$J = 2 \rightarrow J = 4$	para	100.9	813.7
HD	$J = 0 \rightarrow J = 2$		33.11	267.1
	$J = 1 \rightarrow J = 3$		54.93	443.0
	$J = 2 \rightarrow J = 4$		76.35	615.8
	$J = 3 \rightarrow J = 5$		97.22	784.2
	$J = 4 \rightarrow J = 6$		117.4	946.8
D_2	$J = 0 \rightarrow J = 2$	ortho	22.20	179.1
	$J = 1 \rightarrow J = 3$	para	36.89	297.5
	$J = 2 \rightarrow J = 4$	ortho	51.40	414.6
	$J = 3 \rightarrow J = 5$	para	65.67	529.7
	$J = 4 \rightarrow J = 6$	ortho	79.63	642.2
	$J = 5 \rightarrow J = 7$	para	93.20	751.7
	$J = 6 \rightarrow J = 8$	ortho	106.3	857.6
	$J = 7 \rightarrow J = 9$	para	118.9	959.3

For diatomic molecules, the allowed transitions satisfy

$$\Delta J \in \{-2, 0, 2\} \tag{2.23}$$

and

$$\Delta \nu \in \{-1, 0, 1\}. \tag{2.24}$$

Raman transitions preserve the ortho and para populations because ΔJ is even,

For vibrational transitions, specific selection rules apply to different molecules, which depend on the symmetry group of the modes of vibration.

Given the energy levels described in the previous section, we can calculate the Raman shifts associated with several of the $\Delta v = 0$ transitions in the Stokes branch ($\Delta J = 2$). The change in energy can be expressed as a change in wavenumber, using

$$\Delta E = hc\Delta \tilde{\nu}, \tag{2.25}$$

where $\tilde{\nu}$ is the wavenumber. The Raman shifts under $1000 \, \text{cm}^{-1}$ for the three most common hydrogen species are tabulated in Table 2.2, and were calculated from values provided by [56].

In a typical Raman setup, the outgoing light is spectrally analysed using a diffraction grating and collected by a CCD camera. For a given transition, the measured intensity of the corresponding spectral line has the form [84], [92]

$$I_{obs} = \left(\frac{\pi}{\varepsilon_0}\right)^2 A_\nu \tilde{\nu}_0 \tilde{\nu}_s^3 N \Phi \mathscr{I}. \tag{2.26}$$

Here ε_0 is the vacuum permittivity, A_ν is the (wavelength dependent) detection efficiency of the setup, $\tilde{\nu}_0$ is the wavenumber of the laser light, $\tilde{\nu}_s$ is the wavenumber of the outgoing light, N is the number density of molecules in the initial state of the transition (the quantity we want to infer), $\mathscr{I}$ is the irradiance. Φ contains information about the cross-section of this transition, which among other factors depends on the polarisability of the transition and the orientation of the molecules relative to the polarisation of the incoming light.

The relative line strength method

The observed number of photons ostensibly depends on many factors, which are themselves susceptible to many sources of variation. As the laser output degrades over time, the irradiance is reduced. The cross-section for a given transition depends on its polarisability. The detection efficiency of the CCD may not be constant over its whole range. This complicates the extraction of an absolute N from I_{obs}. However, we are usually not interested in absolute number densities, but in ratios of different species within a sample.

Suppose we want to determine the ratios of the species A and B in the sample S. We choose two transitions belonging to A and B and use the ratio of the populations in the initial states of those transitions (N_{SA}/N_{SB}) as a proxy for their ratio. We compute the ratio of two spectral lines in the sample

$$\frac{I_{SA}}{I_{SB}} = \frac{(\pi/\varepsilon_0)^2 \, \tilde{\nu}_0 \, (A_\nu \tilde{\nu}_0 \tilde{\nu}_s^3 \Phi)_A \, \mathscr{I}_S N_{SA}}{(\pi/\varepsilon_0)^2 \, \tilde{\nu}_0 \, (A_\nu \tilde{\nu}_0 \tilde{\nu}_s^3 \Phi)_B \, \mathscr{I}_S N_{SB}} = \frac{(A_\nu \tilde{\nu}_s^3 \Phi)_A}{(A_\nu \tilde{\nu}_s^3 \Phi)_B} \frac{N_{SA}}{N_{SB}}. \tag{2.27}$$

The constants and the irradiance drop out. When we take a calibration measurement C, with a known ratio N_{CA}/N_{CB}, we obtain

$$\frac{I_{CA}}{I_{CB}} = \frac{(A_\nu \tilde{\nu}_s^3 \Phi)_A}{(A_\nu \tilde{\nu}_s^3 \Phi)_B} \frac{N_{CA}}{N_{CB}}. \tag{2.28}$$

both equations can be combined to give

$$\frac{N_{SA}}{N_{SB}} = \frac{N_{CA}}{N_{CB}} \left(\frac{I_{SA}}{I_{SB}} \cdot \frac{I_{CB}}{I_{CA}} \right). \tag{2.29}$$

This is called the relative line strength method [84]. When we determine the ortho-deuterium fraction in the UCN source deuterium, we use deuterium which was stored in a bottle at room temperature as our calibration sample.

There are some caveats. The populations in the initial states are temperature-dependent. In addition, the parameters A_ν and Φ may depend on temperature. Thermal expansion and contraction of components inside the Raman setup may slightly offset the position of the diffraction grating. Therefore, it is important to keep the temperatures of the samples, the laser, and the CCD as stable as possible, between experiments.

The Raman spectroscopy setup in the UCN group

A detailed description of the setup is given in N. Hild's PhD thesis [84]. Here, I will address the basics. The system used is the RamanRxn1 analyser, which was supplied by Kaiser Optical Systems[3]. The analyser is a single, integrated device, which contains the laser, the spectrograph and the CCD. Optic fibres are used to transport the laser light from the analyser to a sample cell. The sample cell contains the gas, this is where the Raman transitions take place. The collected light is guided back to the analyser. A diffraction grating diffracts the light with Raman shifts ranging from 148 to 1148 cm^{-1} onto a 1024 × 256 pixel CCD camera, which is cooled to 233 K to reduce noise. There are no gaps between the pixels, therefore, the measured spectra should be treated as binned, not sampled. The readout noise is on the order of 1.1 counts. The gain of the camera is 10 electrons per count.

2.3   Procedure for taking deuterium samples from the UCN source

A simplified schematic of the compartments and valves of the deuterium system that are relevant in this chapter is shown in the introduction, in Figure 1.6. We can take samples from the moderator, the condenser,

---

[3]Acquired by Endress+Hauser AG, Kägenstrasse 2, 4153 Reinach BL, Switzerland, https://www.endress.com/

the converter or from the storage tanks. The deuterium in the chosen compartment is heated up to about 20 K to ensure the vapour pressure is sufficiently high to permit sampling. At 20 K, deuterium is liquid.

The sample bottle is connected to the 'Gasmaus' (see Appendix B) and evacuated.

The bellows are compressed to their minimum volume, pumped, and expanded to its maximal volume. Valves in the UCNsource are opened until there is a path between the deuterium compartment and the compressible bellows.

The bellows are now filled with deuterium gas, equal to the vapour pressure of the liquid, which can range from 300 to 500 mbar. To isolate the bellows from the rest of the UCN source gas system, valve AA101 is closed.

The bellows are compressed, increasing the pressure of its contents to about 1 bar. Valve AA204 is opened to fill the deuterium into the sample bottle. This reduces the pressure. To prevent accidental overpressure in the UCN source, the bellows can only be expanded or compressed when isolated from other components by valves AA101 and AA204. Valve AA204 is closed, the bellows are compressed to 1 bar, and valve AA204 is opened, until the desired pressure in the sample bottle (1 bar) has been reached.

2.4   Analysis of Raman spectra

The analysis consists of several steps. First, cosmic ray events are removed from the spectra, using sigma clipping. Then, the line strengths corresponding to the Raman transitions are calculated. By comparing the line strengths with those calculated from a reference spectrum, the concentrations of the various species in the sample are determined. Lastly, the concentration in the liquid deuterium the sample was taken from is determined, by correcting for ortho-to-para conversion in the Raman cell, and distillation effects due to the different vapour pressures.

The analysis is largely based on the analysis described in [84], with some improvements. Instead of fitting the background to three large regions containing multiple transition lines, we fitted the background in the region around every individual line. We also used more transition lines, and calculated the concentrations from all transition line intensities simultaneously, weighted by their uncertainties.

2.4.1   Cosmic rays

Radiation consisting of charged particles induces electron-hole pairs in silicon, and can therefore produce counts in CCD pixels. These can be muons, resulting from the decay of mesons produced in hadronic showers, or recoil electrons from Compton scattering of gamma rays [93]. Since we set the readout time of the CCD in the RamanRxn1 to 100 seconds, there is a high probability that a cosmic ray event is recorded in the spectrum. We have observed that a single cosmic ray event can cause a response in one or multiple pixels, which can be orders of magnitude stronger than the signal we are interested in.

Quantification of the effect of cosmic rays

To make a quantitative statement about the number of cosmic ray events, and their probability distribution, I analysed an empty cell spectrum recorded on 2022-10-19. We switched on the RamanRxn1, and reached thermal equilibrium. Spectra were recorded every 100 seconds for almost a week, resulting in a series of 5941 spectra.

For every spectrum j, the number of counts $C_{i,j}$ in every pixel j is recorded.

We want to establish the probability distribution of the deviation of the counts recorded in each pixel,

relative to their expected counts. We define Δc as

$$\Delta c_{i,j} = c_{i,j} - E_{i,j}, \tag{2.30}$$

where $E(i,j)$ is the number of counts we would expect for spectrum i and pixel j.

If there were no drifts in the spectra over time, $E_{i,j}$ could simply be defined as some measure of central tendency over i, for example as a mean:

$$E_j = \frac{1}{n} \sum_{i=0}^{n} c_{i,j}. \tag{2.31}$$

The expectation values for the counts in every pixel drift over time, largely due to fluctuations in temperature. Due to the thermal expansion and contraction of the setup, the projection of the spectrum on the CCD can change by one pixel. Therefore, this model would overestimate the pixel noise. We can compensate for drifts by comparing $c_{i,j}$ with a weighted mean of the counts observed by the same pixel, in a range of time centred on i:

$$E_{i,j} = \sum_{k=-K}^{+K} w_k c_{i+k,j}; \quad \sum_{k=-K}^{+K} w_k = 1. \tag{2.32}$$

This is a discrete convolution. We chose the weights, w_k, from the discrete Gaussian kernel. In order not to compare $c_{i,j}$ with itself, which would underestimate the deviation, the central weight (w_0) was set to zero, and the weights were renormalised to add up to 1.

With the expected value, we could calculate the deviation, $\Delta c_{i,j}$, using Equation 2.30. The probability distribution of *Deltac* looked like the sum of a Gaussian function and a background. Around *Deltac* = 30, the contribution from the Gaussian function was negligible. We classified counts for which $|\Delta c > 25|$ as outliers (0.14 % of the total number of counts). We introduced a masking variable:

$$m_{i,j} = \begin{cases} 1 & \Leftarrow |\Delta c_{i,j}| \le 30 \\ 0 & \Leftarrow |\Delta c_{i,j}| > 30 \end{cases} \tag{2.33}$$

that covers all observations which are not outliers. This mask was used to recompute the expectation values, ignoring contributions from outliers:

$$E_{i,j} = \frac{\sum_{k=-K}^{+K} w_k m_{i,j} c_{i+k,j}}{\sum_{k=-K}^{+K} w_k m_{i,j}}. \tag{2.34}$$

The resulting outlier-free expectation values were used to recompute the deviations using Equation 2.30. The calculation of Δc is visualised in Figure 2.4. An empirical probability distribution of the deviations of the counts from their expectation values is shown in Figure 2.5. The figure shows two processes at work: most deviations follow a normal distribution with a standard deviation of 4.51 counts, and there is a heavy tail of outliers, mostly caused by cosmic rays. To enhance the visibility of outliers (only 0.14 % of the total) we chose a double logarithmic scale for the plot.

Charged particle radiation can create electron-hole pairs in neighbouring pixels. The masking variable, $m_{i,j}$ was used to identify cosmic ray events, defined as clusters of outliers. In the 5941 spectra, 3631 cosmic ray events were identified. This means that every 36 measurements (1 hour), we detect on average 22 cosmic rays. These cosmic ray events involve on average 2.3 pixels (standard deviation 1.5 pixels).

In this particular set of empty cell spectra, 99.02 % of the total variance of the pixels could be explained due to cosmic ray events, 0.74 % due to low-frequency drifts, and 0.24 % due to statistical noise. While the relative importance of these contributions varies between sets of spectra, this shows the importance of cleaning cosmic ray events from the spectra.

Cosmic ray rejection by sigma clipping

In the first step of analysis, we recombine n spectra to one single spectrum, while removing outliers, using the sigma clipping method. This improves the signal-to-noise ratio in the spectra, and it allows us to eliminate
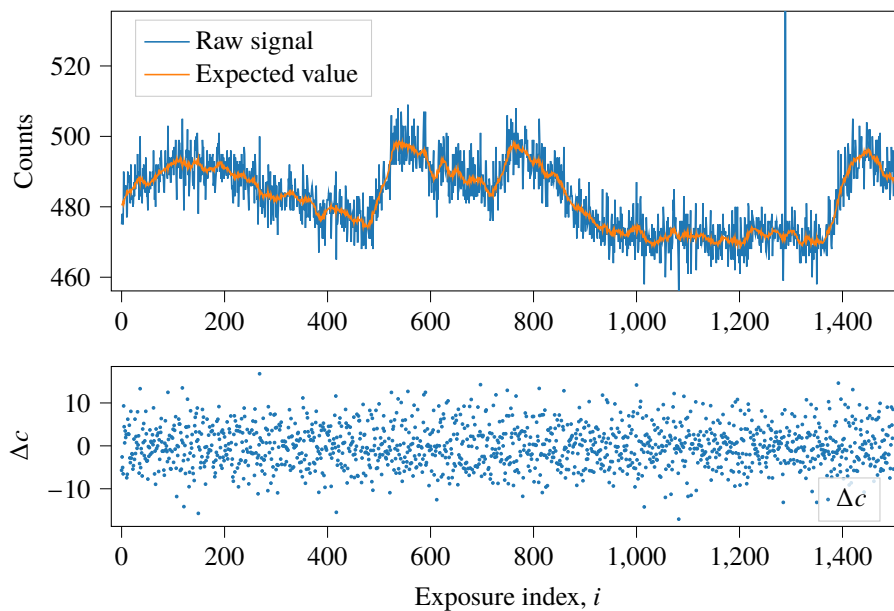
Figure 2.4: Visualisation of the calculation of Δc. The number of counts in pixel 600 is shown for 1500 exposures. Note the cosmic ray at $i = 1289$. The orange curve represents the calculated expected value for pixel 600 at every exposure. Since the outliers were masked in the calculation, the cosmic ray does not affect the expected value. Δc is normally distributed and does not have any structure.
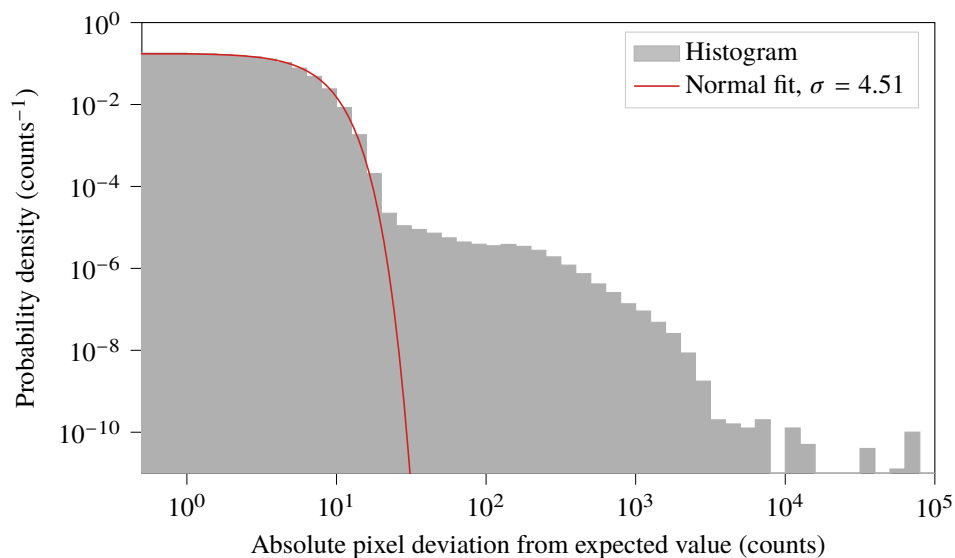
Figure 2.5: An empirical probability distribution of the deviation of the measured number of counts during an empty cell Raman measurement from their expected values. Most deviations follow a normal distribution ($\sigma = 4,51$), but there is a heavy tail of outliers, mostly caused by cosmic ray events. The double logarithmic scale exaggerates the visibility of outliers (0.14 % of the total). The probability density is not transformed into logarithmic space. The probability density around 10^2 counts is 4×10^{-6} counts^{-1}, therefore, the probability for a pixel to have a value between 99.5 and 100.5 counts above the expected value is 4 ppm.

outliers at the cost of temporal resolution. For a set of n exposures, we have the following raw counts: $c_{i,j}$, where i is the exposure index, and j is the pixel index. We assume that the counts follow the Poisson distribution. This is an approximation, it overestimates the uncertainties of the empty cell spectra, where the counts are 300-600, and the standard deviation is 4.5.

For every pixel j, the median number of counts, $\tilde{c}_j$, is determined over the group of exposures. Then a mask is defined:

$$m_{i,j} = \begin{cases} 1 & \Leftarrow c_{i,j} \leq \tilde{c}_j + 2\sqrt{\tilde{c}_j} \\ 0 & \Leftarrow c_{i,j} > \tilde{c}_j + 2\sqrt{\tilde{c}_j} \end{cases}. \tag{2.35}$$

The combined value of each pixel is calculated as

$$c_j = \frac{\sum_{i=0}^{n} m_{i,j} c_{i,j}}{\sum_{i=0}^{n} m_{i,j}}. \tag{2.36}$$

In other words, in each group, the mean of all pixel values for a pixel is taken, excluding all values higher than the median value, plus $2\sqrt{2}$ times the median.

The uncertainties, σ_{c_j}, are calculated as

$$\sigma_{c_j} = \sqrt{\frac{c_j}{G \sum_{i=0}^{n} m_{i,j}}}, \tag{2.37}$$

where G is the gain of the CCD camera, which is 10 photoelectrons per count [94].

Using sigma clipping, if the CCD is hit by a cosmic ray during one exposure, the values from all unaffected pixels can be used, and only the values affected by the cosmic ray are discarded.

The group size, n, was normally set to 20. If the total number of exposures is not a multiple of n, the remaining spectra are included in the last full group of n, this is to avoid having to do sigma clipping on a group which is too small.

2.4.2   Background subtraction

After the removal of cosmic rays from the spectra, the spectra still contain background, which includes the counts that are observed even if there is no laser light ('dark counts'), and Raman scattered light from the optic fibre or the sapphire window between the optic fibre and the sample cell. To characterise the background, Raman spectra are recorded with an empty sample cell and are recombined into a single spectrum by sigma clipping.

Over time, the background is not entirely stable. The absolute background count rate can change. The shape of the background also varies. Before subtracting the background, we need to rescale it, so that it fits the regions in the sample spectra that do not include Raman shift lines. We used two different methods of background subtraction. The difference between both methods was small and irrelevant for determining the para-to-ortho deuterium ratio. However, the use of both methods was relevant to assess how much the estimation of the HD line strength depended on the background subtraction.

Line-based background subtraction

The regions of all HD and D_2 lines were identified and are referred to as the line regions. The complement of the line regions are the background regions. The spectra were divided into sections by cutting them in the regions between the lines, such that every section contained a single line region, surrounded by two parts of the background region.

To match the sample spectrum in the background regions as closely as possible, the background spectrum was scaled with factor A:

$$B_j = Ab_j, \quad \text{with } j \in \text{background region.} \tag{2.38}$$

Then the background was subtracted from the entire section (line and background regions).

$$C'_j = c_j - B_j. \tag{2.39}$$

Since the shape of the background varied, a slope correction was done by fitting a line to the background region and subtracting it from the entire section to obtain the background subtracted counts, C_j. This method is based on the method in [84]. Since the slope and the scale factor are taken without error, the error on C_j is

$$\sigma^2_{C_j} = \sigma^2_{c_j} + A^2 \sigma^2_{b_j}. \tag{2.40}$$

Spline-based background subtraction

We defined the line regions as in the line-based background subtraction, and the background region as the complement. However, instead of cutting the spectrum, we rescaled the background to match all the background regions at once.

The scaled background has the shape

$$B_j = S(j|\{J\}, \{A\})b_j. \tag{2.41}$$

Here b_j are the background counts, $S(j|\{J\}, \{A\})$ is a cubic spline, that cubically interpolates between the points (J_k, A_k). The set $\{J\}$ is predefined. The points in $\{A\}$ are fitted, such that the scaled background B_j matches the sample spectra in the background regions as closely as possible. Afterwards, the scaled background, B_j, was subtracted from the sample spectra:

$$C_j = c_j - B_j. \tag{2.42}$$

The fitted spline is taken without error, so the error from the background is propagated into the subtracted counts as follows:

$$\sigma^2_{C_j} = \sigma^2_{c_j} + S(j|\{J\}, \{A\})^2 \sigma^2_{b_j}. \tag{2.43}$$

2.4.3 Line strength calculation

Since there are no gaps between the CCD pixels, the spectra should be treated as binned data. It is not necessary to use a higher-order numerical integration scheme, such as Simpson's rule or the trapezoidal rule. To calculate the line strengths, one simply sums over the counts in the line region:

$$I = \sum_{j \in \text{l.r.}} C_j, \tag{2.44}$$

with the error

$$\sigma^2_I = \sum_{j \in \text{l.r.}} \sigma^2_{C_j}. \tag{2.45}$$

The *HD* lines were rather weak in the spectra obtained from UCN source deuterium. We assume that the distribution of photons from the HD lines, over the pixels in CCD, follows a Gaussian distribution. Because the pixels have non-zero size, each pixel would not sample this distribution at a given point, but integrate over the distribution. The 'binned' Gaussian function that we fitted to the HD lines was

$$g_j(I, \mu, \sigma) = \frac{I}{\sqrt{2\pi}} \int_{j-\frac{1}{2}}^{j+\frac{1}{2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx, \tag{2.46}$$

where I is the line strength, μ is the location of the peak, and σ the standard deviation. Solving the integral gives

$$g_j(I, \mu, \sigma) = I\left[F\left(j + \tfrac{1}{2}|\mu, \sigma\right) - F\left(j - \tfrac{1}{2}|\mu, \sigma\right)\right],\qquad(2.47)$$

where $F(x|\mu, \sigma)$ is the cumulative distribution function for the normal distribution with mean μ and standard deviation σ.

To illustrate why the integration step is necessary, consider a Gaussian PDF with $\mu = 0.5$ and $\sigma = 0.2$, and a total area of 1. If this function is evaluated at integer numbers, its value will be 0.088 at $x = 0$ and $x = 1$, and even lower for all other x. The sum of all the values would be 0.175. Conversely, if we take a Gaussian with the same width, and $\mu = 0$, its value would be 1.994 at $x = 0$. The value of the function at integer points will not add up to 1. If we used this function to fit the data, we would risk drastically under- or overestimating the area under the curve. The binned Gaussian described above mitigates this problem, as it guarantees the area under the curve is I.

2.4.4   Determination of concentrations in the sample

In Section 2.2.3, we obtained the following relation for the ratio of two species in the sample spectrum:

$$\frac{N_{SA}}{N_{SB}} = \frac{N_{CA}}{N_{CB}}\left(\frac{I_{SA}}{I_{SB}} \cdot \frac{I_{CB}}{I_{CA}}\right).\qquad(2.48)$$

Here, I is the line strength and N is the quantity of one of the species. The subscripts A and B refer to species A and B, or a specified Raman line belonging to this species. The subscripts S and C refer to respectively the spectra obtained from the sample and those from the calibration compound with known concentrations.

By introducing the line ratio $R = I_S/I_C$ (the ratio of a line strength in the sample spectrum to the line strength in the calibration spectrum), Equation 2.48 can be rewritten as

$$\frac{N_{SA}}{N_{SB}} = \frac{N_{CA}}{N_{CB}}\frac{R_A}{R_B}.\qquad(2.49)$$

Since we have more than one Raman shift line for each species, we can generalise this to

$$\frac{N_{SA}}{N_{SB}} = \frac{N_{CA}}{N_{CB}}\frac{\langle R_A\rangle}{\langle R_B\rangle}\qquad(2.50)$$

using the uncertainty-weighted means of the line ratios

$$\langle R\rangle = \frac{\sum_{l=1}^{L} R_l w_l}{\sum_{l=1}^{L} w_l},\quad w_l = \frac{1}{\sigma_{R_l}^2}.\qquad(2.51)$$

The deuterium in the calibration spectra always comes from a gas bottle filled with commercial deuterium, which has been stored at room temperature. We can safely assume that the ortho-deuterium fraction in the calibration compound equals two thirds ($N_{CO}/N_{CP} = 2$).

To estimate the HD concentration in the UCN source, we had to prepare a special mixture of HD and D_2. This will later be discussed in detail.

2.4.5   Inferring concentrations in the liquid/solid from the gas samples

Since the deuterium sample cell is at room temperature, the ortho-deuterium fraction in the sample will approach $\frac{2}{3}$ slowly. To get the correct ratio at the time the sample was taken, an exponential curve of the shape

$$f_o(t;f_{oi}, \tau) = \frac{2}{3} + \left(f_{o_i} - \frac{2}{3}\right)e^{-\frac{t}{\tau}}\qquad(2.52)$$

is fitted to the ortho fractions. Here f_{oi} is the initial ortho fraction, and τ is the time constant. To determine the HD concentration, no such correction is required.

We have established the procedure to estimate concentrations of o-D_2 , p-D_2 and HD in the gaseous sample at the time the Raman measurement is started. If the sample was taken from vapour of liquid deuterium, as is usually the case, the concentrations in the gaseous sample are not equal to those in the liquid. When a sample is taken from the deuterium in the moderator, the converter, or the condenser, the deuterium is allowed to warm up above the triple point (18 K). The surface of the deuterium melts, and there is a decent vapour pressure from which to sample.

The vapour pressures of ortho-deuterium and para-deuterium are not equal. In the temperature ranges from 18 to 28 K, empirical formulae have been determined for the vapour pressures for hydrogen deuteride (p_h), natural deuterium (p_n, ortho fraction $f_o = \frac{2}{3}$) and deuterium at thermodynamic equilibrium (p_e, ortho fraction $f_{oe}(T)$) [56]. These depend on the ortho and para-deuterium vapour pressures, respectively $p[\text{o-D}_2]$ and $p[\text{p-D}_2]$. According to Raoult's law [95], the partial pressures of different components of a liquid are equal to the vapour pressures of their pure components, multiplied by their molar fraction in the liquid. This gives us the following relations:

$$\begin{cases} p_e(T) = \frac{2}{3}p[\text{o-D}_2] + \frac{1}{3}p[\text{p-D}_2] \\ p_n(T) = f_{oe}(T)p[\text{o-D}_2] + (1 - f_{oe}(T))p[\text{p-D}_2] \end{cases} . \tag{2.53}$$

As in [84], we can solve this system of equations for $p[\text{o-D}_2]$ and $p[\text{p-D}_2]$ to find:

$$\begin{cases} p[\text{p-D}_2](T) = \dfrac{2p_e - 3p_n f_{oe}(T)}{2 - 3f_{oe}(T)} \\ p[\text{o-D}_2](T) = \dfrac{3}{2}p_n - \dfrac{2p_e - 3p_n f_{oe}(T)}{2 - 3f_{oe}(T)} \end{cases} . \tag{2.54}$$

The resulting vapour pressures are displayed in Table 2.3 and Figure 2.6.

During sampling, the pressure in the relevant deuterium compartment is noted down. If ortho-deuterium is the dominant contribution, the temperature can be estimated by solving $p[\text{o-D}_2](T) = p_{\text{measured}}$ for T, and used to estimate $p[\text{p-D}_2]$ and $p[\text{HD}]$. This can then be used to calculate the molar fractions of the different species in the liquid deuterium, from the molar fractions in the sample:

$$f_{\text{LA}} = \frac{f_{\text{SA}}p_{\text{A}}}{\sum_B f_{\text{SB}}p_{\text{B}}}. \tag{2.55}$$

When taking samples from condensed deuterium, the measured pressure was always between 350 mbar and 480 mbar. The difference in the vapour pressure ratios in this range is so small that for measurements with unknown pressure, the pressure was assumed to be 415 mbar, introducing a 0.1 % relative error on the estimation of the para-deuterium fraction, and a 0.02 relative error on the estimation of the HD fraction.

2.5 Reproducibility of bottle measurements

To estimate the systematic uncertainty caused by variations in the conditions between experiments. We compared five spectra obtained from the D_2 bottle, obtained from 2020 to 2023. One spectrum was used to estimate the ortho-deuterium fraction in the other spectra, as if they were samples with an unknown fraction. The standard deviation, 0.31 %, can be taken as the uncertainty on a single measurement.

Concentrations and ratios can be converted into each other using

$$R = \frac{C}{1 - C}, \quad C = \frac{R}{1 + R}, \tag{2.56}$$

Table 2.3: Vapour pressures of deuterium and hydrogen deuteride at 20.65 K [56], [84].

Temperature (K)	$p[\text{o-D}_2]$ (mbar)	$p[\text{p-D}_2]$ (mbar)	$p[\text{HD}]$ (mbar)	$p[\text{p-D}_2]/p[\text{o-D}_2]$	$p[\text{HD}]/p[\text{o-D}_2]$
19.0	196	188	352	0.958	1.798
19.5	242	233	426	0.961	1.756
20.0	297	286	509	0.964	1.717
20.5	360	348	605	0.966	1.681
21.0	433	419	713	0.967	1.648
21.5	516	500	835	0.968	1.617
22.0	611	593	971	0.969	1.589
22.5	719	697	1123	0.969	1.563
23.0	840	815	1292	0.970	1.538
23.5	975	946	1478	0.970	1.515
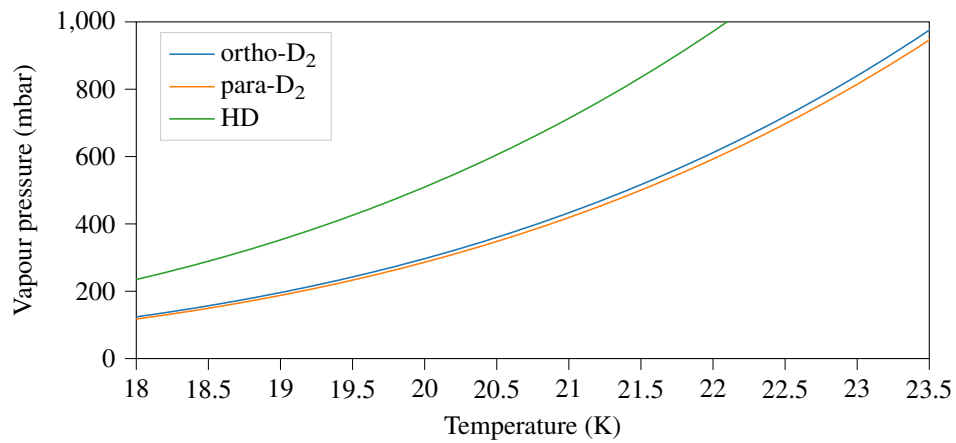
Figure 2.6: Vapour pressure of liquid ortho-deuterium, para-deuterium and hydrogen deuteride, as a function of temperature.

Table 2.4: A summary of Raman samples with the measured para-deuterium fraction.

Sample No.	Sample date	Compartment	Para-deuterium fraction in UCN source (%)
1	2020-07-31	converter	4.19 ± 0.06
2	2020-08-17	moderator	1.25 ± 0.02
3	2020-09-30	moderator	0.60 ± 0.03
4	2020-10-13	moderator	0.59 ± 0.02
5	2021-05-04	moderator	2.52 ± 0.03
6	2021-07-06	condenser	3.34 ± 0.05
7	2021-12-07	moderator	0.45 ± 0.02
8	2022-09-12	condenser	12.41 ± 0.15
9	2022-10-27	condenser	7.90 ± 0.10
10	2022-12-15	moderator	0.57 ± 0.02

with uncertainties

$$\sigma_R = \frac{\sigma_C}{(C-1)^2}, \quad \sigma_C = \frac{\sigma_R}{(R+1)^2}. \tag{2.57}$$

The ortho-to-para ratio, at room temperature equilibrium, equals 2, the relative uncertainty on the measured value equals 1.4 %. This can then be used to determine the minimum error on a concentration:

$$\sigma_C = 0.014C(1-C). \tag{2.58}$$

2.6 Summary of deuterium samples from the UCN source

A summary of Raman samples taken from 2020 to 2022 is shown in Table 2.4.

2020

In September and October 2020, tests were done to investigate the impact of the cooling power during the freezing of deuterium in the moderator vessel on the UCN output. During these tests, several UCN source samples were taken to ensure the o-D_2 fraction was still adequate after evaporating the deuterium.

In August (sample 2 in Table 2.4), the moderator p-D_2 fraction was 1.25 %, and this improved further to 0.60 % in September and October (samples 3 and 4, respectively).

2021

In 2021, the mirror neutron measurement campaign was conducted [96]. We filled the moderator with the largest amount of deuterium up to that point [48] (see also Chapter 4). First, we transferred deuterium from the moderator to the condenser, with a temperature just above the melting point to ensure the highest possible ortho fraction. We then gradually warmed up the condenser to 33 K to extract part of the deuterium adsorbed to the oxɪsoʀʙ®. The deuterium that was released, was then frozen inside the moderator. This increases the para-deuterium fraction slightly, so we measured the para fraction and found 2.52 %(sample 5). This is 0.5 % higher than the equilibrium at 20 K, and will be converted to ortho-deuterium by radiation conversion.

End of June, a failure of one of the UCN source computers caused the UCN source to go into failsafe mode. All valves were opened, and cooling power was lost. All deuterium boiled off into the storage tanks. Afterwards, we liquefied the deuterium in the condenser. We determined the para-deuterium fraction in the

condenser (see Figure 1.6 or Appendix B) to be 3.34 % (sample 6) and decided to transfer the deuterium to the converter before refreezing it in the moderator.

At the end of the year, the para-deuterium fraction was 0.45 % (sample 7). This is the lowest observed para fraction observed from 2020 to 2022. Due to the mirror neutron campaign, 2021 was a year of intense operation of the UCN source, resulting in a lot of radiative conversion (Chapter 8 of [84]).

2022

In September and October, samples 8 and 9 were taken from the condenser vessel to determine the conversion rate in the condenser (Section 2.8), and to find out whether the para-deuterium fraction is low enough to add the deuterium in the condenser to the moderator.

In December, the deuterium in the condenser was transferred to the moderator, filling the moderator with the largest amount of deuterium to date. We investigated whether and how much this increases the UCN output (see Chapter 4). Since the UCN yield is strongly dependent on the para-deuterium fraction, the para-deuterium fraction was measured to allow a comparison of the UCN yield with other measurements. It was found to be 0.57 %.

2.7 The ortho-to-para conversion rate in the sample bottle

To find out how much conversion takes place in the sample bottle in the time between sampling and filling the Raman cell (around 20 min) we measured the conversion time constant in the sample bottle.

The sample bottle was filled with deuterium with a high ortho fraction. We filled the Raman cell from the sample bottle three times, with one day in between. We obtained a series of exposures after each filling. An exponential curve was fitted to the ortho-deuterium fractions in each series, finding the initial ortho-deuterium fraction, and the conversion time constant in the Raman cell. To find the deuterium fraction at the time the sample was taken and the conversion time constant for the sample bottle, we fitted another exponential curve to the fit results. The results are shown in Figure 2.7.

We obtained a time constant for the conversion in the sample bottle of (157 ± 6) d. This should be seen as an order-of-magnitude estimate. The conversion time constant in the Raman cell normally varies between 60 and 120 days. This depends on how well the system has been pumped: oxygen is a catalyst for conversion between para and ortho-deuterium [97]. Since the sample bottle is larger than the Raman cell, the rate of wall collisions is lower, and it is very plausible that its conversion rate is lower, in the absence of oxygen.

2.8 The natural conversion rate for deuterium in the condenser vessel

In 2022, we filled the moderator with 5.68 kg of deuterium, the highest amount to date. After we filled the moderator with 5.28 kg, we warmed up the converter, releasing 400 g of deuterium that was still adsorbed in the OXISORB® into the storage tanks. We transferred the deuterium to the converter. It remained there for about two months, at a temperature of 6.1 K, so that the ortho-deuterium fraction increased due to natural conversion. Then it was transferred to the moderator. This is explained in detail in Chapter 4.

This provided us the opportunity to measure the natural conversion rate in the condenser. The condenser has the most accurately determined temperature of the UCN source. It is located above the biological shielding of the UCN source. It only receives background radiation and is not sensitive to proton beam pulses.

We took two samples from the deuterium in the condenser, separated by 45 days. The first sample was taken on 2022-09-12 14:00 (sample 8 in Table 2.4), the second on 2022-10-27 17:47 (sample 9). We determined the initial ortho-deuterium fraction in the samples to be 87.98 % and 92.39 %. The corresponding
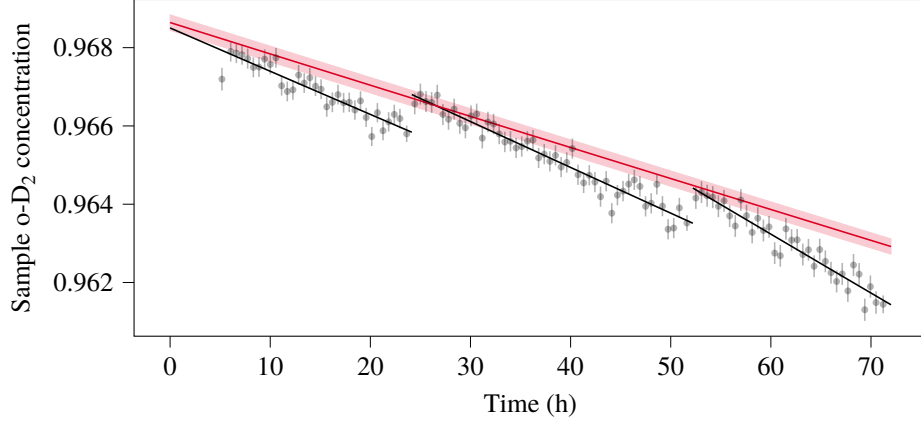
Figure 2.7: Determining the ortho-to-para conversion time constant in the deuterium sample bottle. The Raman cell was filled three times. Each time, we obtained a series of exposures and fitted the initial ortho-deuterium fraction and the conversion time constant in the Raman cell (black). To find the conversion rate in the sample bottle, we fitted another exponential function to the initial ortho-deuterium fraction at the time of the filling.

ortho fractions in the condenser, taking into account the vapour pressure ratio at the pressures at which the samples were taken, were respectively 87.64 % and 92.13 %.

During the entire period between the collection of both samples, the temperature in the condenser was maintained at a constant (6.1 ± 0.1) K (measured by temperature sensors 2RNK10 CT001 and CT002). At this temperature, the equilibrium ortho-deuterium fraction, $C_{e,p} = 1 - 1.4 \times 10^{-6}$. Assuming that the ortho fraction in the condenser evolves as

$$C_o(t) = C_{e,o} + \left(C_o(0) - C_{e,o}\right) e^{-t/t_0}. \tag{2.59}$$

The decay time constant t_0 can be computed as

$$t_0 = -\frac{\Delta t}{\log R} \quad \text{with} \quad R = \frac{C_{e,o} - C_{o1}}{C_{e,o} - C_{o2}}. \tag{2.60}$$

where C_{o1} and C_{o2} are the measured ortho fractions. We obtained a decay time constant of (100.3 ± 2.5) d.

The dominant uncertainties stem from the estimation of the ortho fraction. The uncertainties on Δt (on the order of 20 minutes) and $C_{e,o}$ (well below 1 ppm) are negligible. The uncertainty on the estimate of the decay constant can then be computed as follows:

$$\left(\frac{\sigma_R}{R}\right)^2 = \left(\frac{\sigma_{C_{o1}}}{C_{e,o} - C_{o1}}\right)^2 + \left(\frac{\sigma_{C_{o2}}}{C_{e,o} - C_{o2}}\right)^2, \tag{2.61}$$

where

$$\sigma_{t_o} = \sigma_R \frac{\Delta t}{R \left(\log R\right)^2}. \tag{2.62}$$

The natural ortho-to-para-deuterium conversion in the condenser vessel at $T = (6.1 \pm 0.1)$ K was measured to have a time constant of (100.3 ± 2.5) d. The uncertainty is purely based on the uncertainty of the Raman measurements and does not account for systematic effects.

This value falls between two time constants given for the natural spin conversion of deuterium: 79 d for solid deuterium at 4.2 K and 129 d for liquid deuterium at 20.4 K (Table 24.1 in [56]).

Perhaps counterintuitively, the timescales for $J = 1 \rightarrow 0$ conversion in solid hydrogens increase with temperature [56], which can explain the discrepancy with the reported value for 4.2 K.

(a)

(b)
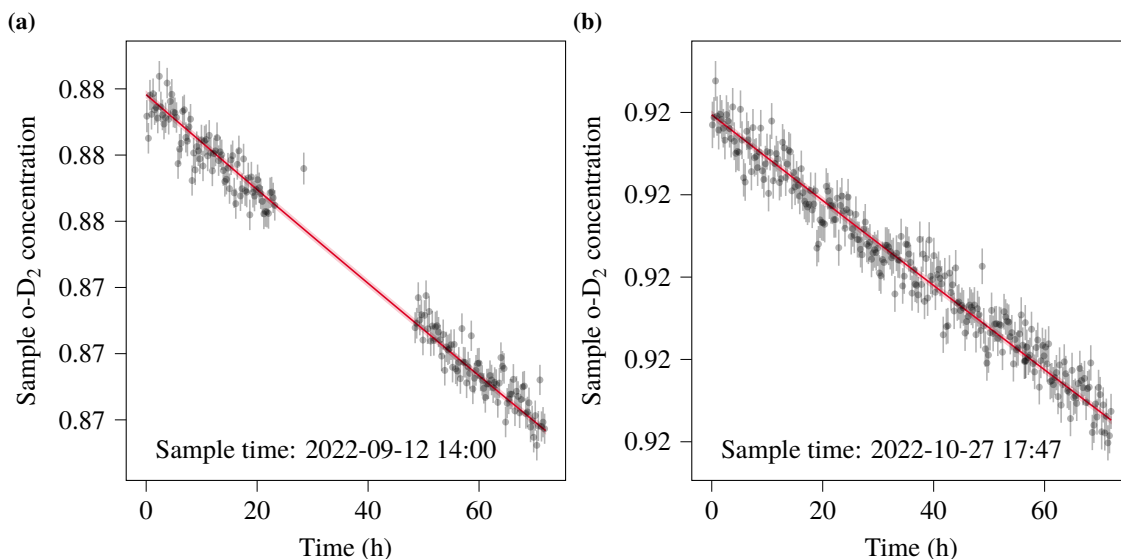
Figure 2.8: The ortho-deuterium fraction obtained from deuterium samples taken from the condenser in (a) September and (b) October 2022. Every grey point represents the ortho-deuterium fraction in the sample cell, calculated from a spectrum obtained by sigma clipping and averaging twenty 100-second exposures, with the corresponding 1σ error bar. The red lines represent the exponential decay fits, which asymptotically approach $\frac{2}{3}$. The sample ortho fraction at $t = 0$ is estimated from the fit and used to estimate the ortho-deuterium fraction in the condenser. The almost invisible pink-shaded area represents the 95 % confidence interval of the fit curve.

2.9   The HD concentration in the UCN source deuterium

Hydrogen's neutron capture cross-section is 640.8 times greater than that of deuterium [85]. The presence of hydrogen in the UCN source moderator leads to the capture of UCN before they are extracted from the deuterium.

Three molecular species can be made from atomic hydrogen and deuterium: H_2, HD and D_2.[4] In equilibrium, the molecular concentrations can be calculated from the atomic concentrations as follows:

$$[H_2] = [H]^2, \quad [HD] = 2[H][D], \quad [D_2] = [D]^2. \tag{2.63}$$

For small hydrogen concentrations, $[D] \approx 1$ and therefore $[HD] \approx 2[H]$. No H_2 lines are visible on the Raman spectra. We assume that these three hydrogens are in equilibrium in the UCN source and that $[H_2]$ is negligibly small (smaller than the uncertainty on $[HD]$).

The HD concentration is determined similarly to the D_2 concentration. Since HD is stable in the sample, we do not require a fit to account for conversion. Due to its low molecular weight, the vapour pressure for HD is higher than for D_2 (Figure 2.3).

The method for determining the hydrogen concentration is similar to the one for determining the ortho concentration. The Raman line strengths obtained from deuterium samples from the UCN source are compared with those obtained from a reference compound, containing known concentrations of deuterium and hydrogen deuteride. The difference from the determination of the ortho concentration is that the HD and D_2 concentrations are stable in the Raman cell. We need not fit a slope to account for conversion.

Since the hydrogen concentration in the UCN source deuterium is very small, the determination of the line strengths is more difficult. In addition to summing over the pixel values, we also used a binned Gaussian

---

[4]We neglect tritium since its concentration is so low that it cannot be detected using our Raman setup.
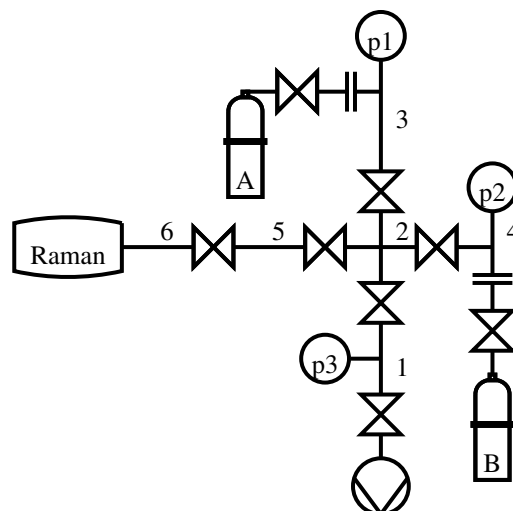
Figure 2.9: The setup for the production of the HD calibration mixture. The 50 mL gas bottles are labelled A and B. All other volumina, separated by valves, are numbered 1 through 6. The pressure gauges p1 and p2 have a range of respectively 0-10 and 0-1 bar, with a precision of 1 mbar. The gauge p3 is a full-range gauge.

fit to determine the line strengths, as described before. Unlike the ortho-deuterium fraction, which is known to be 2/3 at room temperature, the HD concentration has to be determined experimentally in the calibration sample.

The hydrogen concentration was determined earlier by N. Hild [84], who estimated $[HD] = 0.18\%$ and $[H] = 0.09\%$. This measurement was based on a calibration mixture of HD and D_2 supplied by a company, with a relatively large uncertainty on the HD concentration. We recently discovered a company which sells HD with a very high purity. This allows us to make our own calibration mixture, with a more precisely known HD concentration.

2.9.1   The HD calibration mixture

We created the calibration mixture by mixing D_2 and HD from two bottles of high purity. This gives us a smaller uncertainty on the HD concentration in our calibration mixture than in the calibration mixture used in [84].

Deuterium was used from the same bottle used for the calibration of the para-to-ortho ratio in the UCN source deuterium. This was ordered from Messer Schweiz AG.[5] In this bottle, the atomic concentration is 99.8% [98], such that $[D_2] \geq 99.6\%$.

The HD gas was chemically produced by Sigma-Aldrich GmbH[6]. The supplier determined the molar concentrations [99] using mass spectrometry and gravimetric analysis: $[HD] = 97.1\%$, $[D_2] = 1.2\%$, $[H_2] = 1.6\%$.

We filled two 50 mL bottles (labelled A and B) with HD and D_2, and mixed the contents directly in the Raman cell, using the setup shown in Figure 2.9. The HD and D_2 concentrations were calculated from the pressures.

---

[5]Messer Schweiz AG, Seonerstrasse 75, 5600 Lenzburg, Switzerland, https://www.messer.ch/
[6]Sigma-Aldrich, Kappelweg 1, D-91625 Schnelldorf, Germany, https://www.sigmaaldrich.com/

Table 2.5: The recorded pressures in the expansion stages of the measurement to determine the volume sizes of the HD mixing setup.

Starting from bottle A		Starting from bottle B	
Compartment expanded into	Pressure from p1 (mbar)[a]	Compartment expanded into	Pressure from p2 (mbar)[a]
A (atm.)[b]	982	B (atm.)[b]	981
3	782	4	702
2	670	2	612
5	630	5	578
6	532	6	495
4	435	3	440
1	75	1	79

[a] The uncertainty on the pressure is 1.5 mbar.
[b] The bottles were filled to atmospheric pressure.

Measurement of the volumes in the setup

To determine the HD and D_2 concentrations in the calibration sample, we need to know the pressures in bottles A and B, and the volumina of the Raman setup.

We determined the size of the different volumina in the setup (Figure 2.9). We filled bottles A and B with air at atmospheric pressure, and attached them to the setup. We kept the valves on the bottles closed and evacuated the rest of the setup using a Pfeiffer HiCube turbopump unit.[7] Then, the air from one bottle was expanded into the volumina of the setup, while the pressure in pressure gauge p1 or p2 was recorded. The pressure sensors are of the type Cerabar T PMC131, fabricated by Endress + Hauser.[8] The sensor p1 has a range from 0 to 1 bar, and sensor p2 a range from 0 to 10 bar.

This was repeated for the other bottle. The indicated pressures in both gauges were noted down at atmospheric pressure and in vacuum. The results are shown in Table 2.5.

The uncertainty on all pressure measurements is taken as 1.5 mbar.

The volumes were calculated using

$$\tilde{V}_i = V_A \frac{p_i}{p_0} \quad , \quad \left(\frac{\sigma_{\tilde{V}_i}}{\tilde{V}_i} \right)^2 = \left(\frac{\sigma_{p_i}}{p_i} \right)^2 + \left(\frac{\sigma_{p_0}}{p_0} \right)^2 , \tag{2.64}$$

where $\tilde{V}_i$ is the total volume of all the compartments up to and including compartment i, p_i is the pressure recorded after the air has expanded into compartment i, and p_0 is atmospheric pressure. V_A is the volume of bottle A. The bottles have the same volume. All volumes can be expressed in units of V_A. The uncertainty on V_A is irrelevant.

The volume of compartment i is then

$$V_i = \tilde{V}_i - \tilde{V}_{i-1} \quad \sigma_{V_i}^2 = \sigma_{\tilde{V}_i}^2 + \sigma_{\tilde{V}_{i-1}}^2 . \tag{2.65}$$

We obtained two estimates for the volume of each compartment, with unequal uncertainties. A chi-square test was done to ensure the uncertainties were not underestimated. Afterwards, the two separate measurements of every volume were combined by taking the uncertainty-weighted mean. The resulting volumes are presented in Table 2.6.

---

[7]Pfeiffer Vacuum GmbH, Berliner Strasse 43, 35614 Asslar, Germany, https://www.pfeiffer-vacuum.com/
[8]Endress + Hauser, Kägenstrasse 2, 4153 Reinach BL, Switzerland, https://www.endress.com/

Table 2.6: The estimated volumes in the HD calibration mixture setup.

Compartment	Volume from A (V_A)	Volume from B (V_A)	Volume combined (V_A)
1	10.84 ± 0.27	10.19 ± 0.24	10.48 ± 0.18
2	0.210 ± 0.007	0.206 ± 0.006	0.207 ± 0.005
3	0.256 ± 0.005	0.248 ± 0.011	0.255 ± 0.004
4	0.412 ± 0.013	0.397 ± 0.004	0.399 ± 0.004
5	0.093 ± 0.008	0.094 ± 0.007	0.093 ± 0.005
6	0.287 ± 0.010	0.285 ± 0.008	0.286 ± 0.006

The HD concentration in the calibration mixture

Bottle A was filled with deuterium from Messer AG and bottle B with hydrogen deuteride from Sigma-Aldrich. They were connected to the setup. The entire setup was evacuated (except the bottles). After that, the valves were closed. The first valves connected to the bottles were opened to allow the gases to expand into compartments 3 and 4, respectively. The pressures were recorded:

$$V_{D_2} = V_A + V_3, \qquad\qquad p_{D_2} = 2238\,\text{mbar},$$
$$V_{HD} = V_A + V_4, \qquad\qquad p_{HD} = 542\,\text{mbar}.$$

We opened the valves to let hydrogen deuteride expand into the Raman cell. The Raman measurement was started. Lastly, the valve connecting compartments 2 and 3 was opened, allowing the deuterium to enter the Raman cell. The deuterium had to reach the Raman cell by diffusion through the thin tubes. It took several hours until the concentrations of HD and D_2 in the Raman cell reached equilibrium. The ratio of HD line strengths and D_2 line strengths was evaluated over time (Figure 2.10). The equilibrium concentration had definitely been reached after 11 hours. All exposures after 11h were averaged together. This is the reference spectrum (Figure 2.11).

The final molar HD concentration in the calibration mixture is

$$[\text{HD}]_{\text{calib}} = \frac{p_{D_2} V_{D_2} [\text{HD}]_A + p_{HD} V_{HD} [\text{HD}]_B}{p_{D_2} V_{D_2} \left([\text{HD}]_A + [\text{D}_2]_A\right) + p_{HD} V_{HD} \left([\text{HD}]_B + [\text{D}_2]_B\right)} = (20.9 \pm 0.2)\,\%. \qquad (2.66)$$

The concentrations on the right-hand side are total concentrations, reported by the suppliers. The concentration on the left-hand side, $[\text{HD}]_{\text{calib}}$ is defined as $[\text{HD}]/([\text{HD}] + [\text{D}_2])$, ignoring H_2 and trace contributions. The uncertainty includes the uncertainty on the HD and D_2 concentrations in both bottles, taken as 0.1 %, as well as the uncertainty from the volume measurements.

Verification of the reference gas purities

We compared five spectra obtained from bottle measurements, obtained from 2020 to 2023, with the calibration mixture spectrum. We determined the HD concentration in the D_2 bottle to be $(0.269 \pm 0.016)\,\%$ (mean and standard deviation). The standard deviation of 0.016 % can be taken as the minimum systematic uncertainty on any Raman measurement.

Conversely, we cannot accurately confirm the D_2 concentration in the HD bottle, due to background features around the o-D_2 peaks that contribute a large relative error at low D_2 partial pressures. Nevertheless, we determined the D_2 concentration in the HD bottle is non-zero, and around 1 %, which agrees with the certification provided by Sigma-Aldrich.
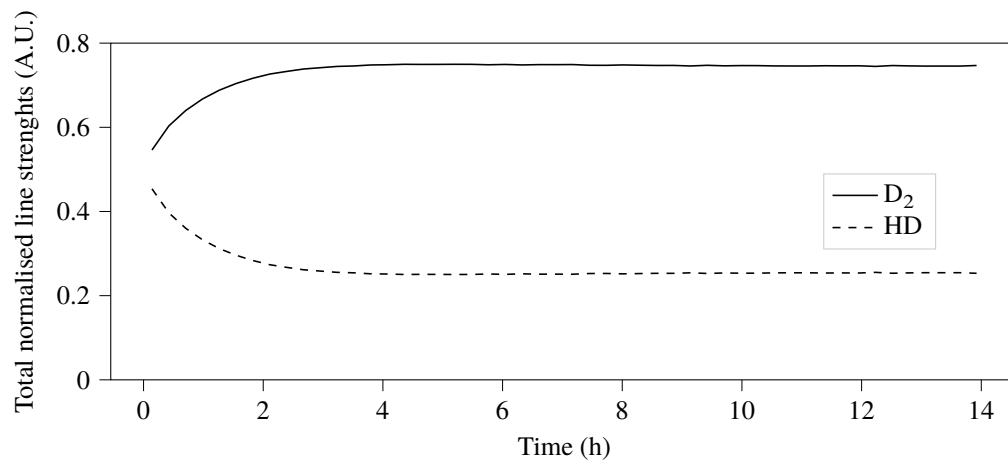
Figure 2.10: The sum of the D_2 and HD line strengths measured in the Raman cell, normalised, over the first 15 h. In the first hours, the HD concentration in the Raman cell decreases. The HD concentration slightly increases in the few hours after that. From $T = 11$ h, the concentration is stable. Note that the line ratio of line strengths alone cannot be interpreted as a ratio in partial pressures of D_2 and HD.
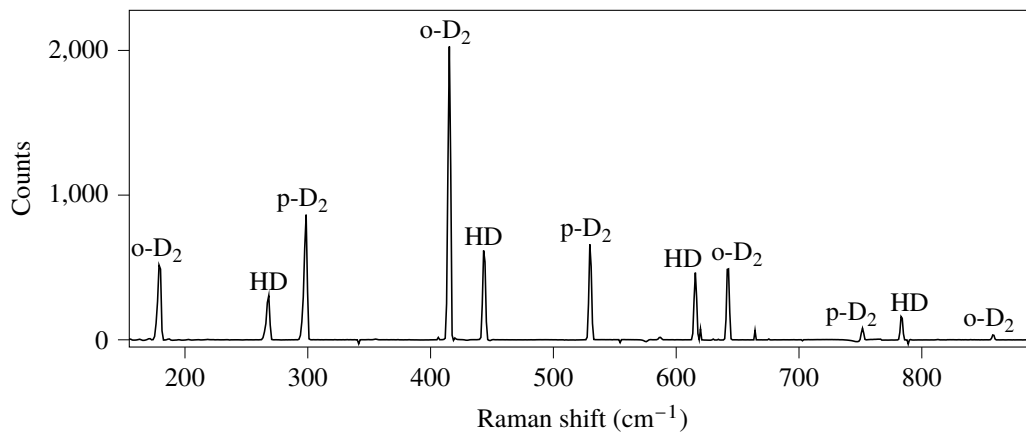
Figure 2.11: The calibration spectrum containing D_2 and HD. The shown spectrum is the average of all the exposures after $T = 11$ h. The background has been subtracted.
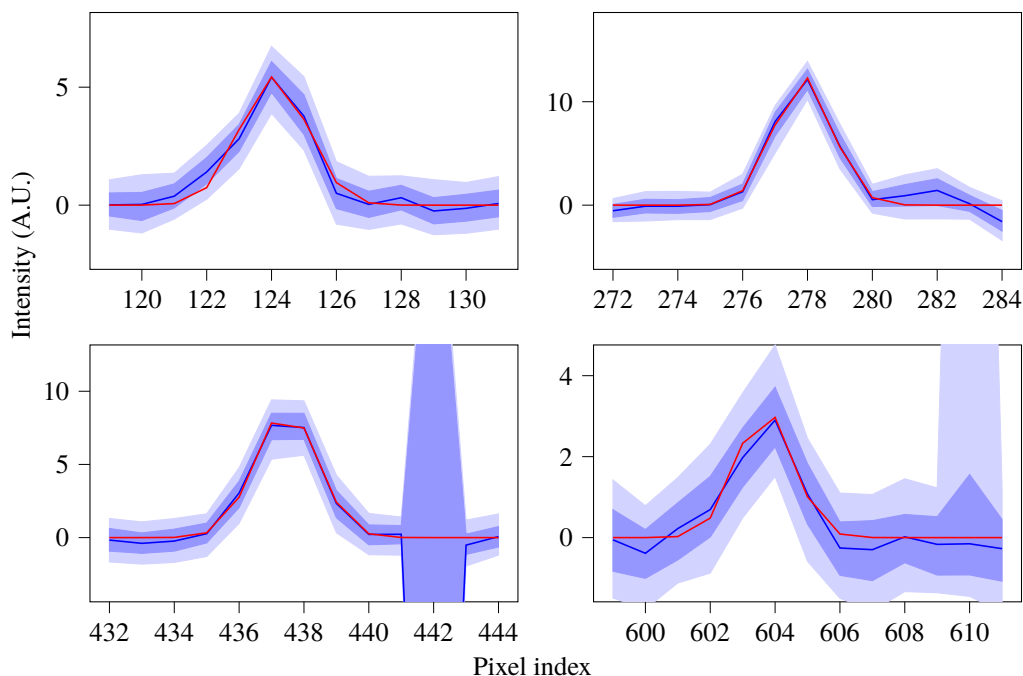
Figure 2.12: The first four HD peaks in the Raman spectra of UCN source deuterium, from sample 8, taken at 2022-09-12 (see Table 2.4). Sigma clipping was done with a group size of 50. The shades of blue represent the regions containing 95.5 % of the values, 68.3 % of the values, and the median. The red curves represent the averages of the binned Gaussian fits. Pixels 442 and 610 showed a high variability and were excluded in the analysis.

2.9.2 HD concentration in UCN source samples

Inspection of the HD lines

The HD line strengths are weaker in the UCN source deuterium Raman spectra than in the spectra from the bottled deuterium. This is because the UCN source has a lower HD concentration, and we cannot fill the Raman cell to the same pressures when we fill it from the sample bottle. For each measurement, we visually inspected the HD peaks to make sure the background subtraction was successful, and the peaks were clearly visible. An example is shown in Figure 2.12.

Data inclusion

Eight different measurements were included in the analysis. Sigma clipping was done with a group size of 50. The initial ortho concentrations in the samples were estimated. This was used to determine the ortho-to-para-deuterium ratio in the sample and the partial pressures. These were then used to determine the HD concentration in liquid deuterium in the UCN source from the measured HD concentrations for each set of exposures.

The results are shown in Figure 2.13. The plots represent the distributions of the values of the HD concentration in the liquid obtained from each group of 50 exposures. The blue plots were calculated by summing over the HD peaks, and the orange plots are based on the fit.

For most measurements, the summation method and the fit method give similar results. For the moderator spectrum obtained on 2021-05-04 (sample 5), the summation method yielded a lower estimate due to background over-subtraction next to one of the HD peaks. The fit method was not sensitive to this. We did

not use the estimate from the summed method in the determination of the UCN source HD concentration.

The measurement from deuterium from the tanks on 2023-03-06 was taken when all the deuterium was in the gaseous phase. The sample HD concentration was lower than in the other measurements because it was not enhanced due to the difference in vapour pressures. This resulted in a lower signal-to-noise ratio, and thus a larger spread of the obtained values. The difference with the other measurements is not statistically significant.

The measurements taken during the determination of the conversion rate in the condenser vessel, on 2022-09-12 and 2022-10-27 (samples 8 and 9), give a significantly lower estimate for the HD concentration. We did not expect this. A priori, one would expect the HD concentration in the condenser vessel to be larger or equal to the concentration in the moderator, as HD has a lower triple point (16.6 K) than D_2 (18.7 K, [56]).

When we correct for vapour pressure, we implicitly assume that the gas volume is small enough that the evaporation of D_2 and HD from the liquid does not affect the HD concentration in the liquid. Since the quantity of deuterium in the converter in September and October 2022 was only 400 g, this assumption may be violated in these measurements. Most of the deuterium was frozen inside the moderator. The quantity in the converter may have been so small that the evaporation significantly changed the concentrations. These two measurements were not used in calculating the final value.

2.9.3   Estimated HD concentration

For every deuterium sample of the UCN source i, an estimate of the HD concentration was obtained from every individual spectrum. The mean concentration, c_i, and standard error of the mean, σ_i were calculated. From the different measurements, these spectra were combined into a total mean by using a weighted mean:

$$c = \frac{\sum_i \frac{c_i}{\sigma_i^2}}{\sum_i \frac{1}{\sigma_i^2}}. \tag{2.67}$$

The weighted standard deviation of the sample means, σ was computed as

$$\sigma^2 = \frac{N}{(N-1)^2} \frac{\sum_i \frac{(c_i - c)^2}{\sigma_i^2}}{\sum_i \frac{1}{\sigma_i^2}}. \tag{2.68}$$

To be conservative, the standard deviation was taken as the uncertainty, and not divided by $\sqrt{N}$. The resulting estimates for the HD concentration in the source, are (0.214 ± 0.010) %, using the summing method, and (0.205 ± 0.012) % using the Gaussian fit method. The estimates of the HD concentration in the bottled deuterium, which we use as the minimum uncertainty on any Raman measurement, had a larger standard deviation of 0.016 %.

The systematic uncertainty from the uncertainty on the HD concentration in the calibration mixture equals 2×10^{-3} %.

We estimate the molar concentration of hydrogen deuteride in the UCN source as

$$[HD]_{\text{molar}} = (0.210 \pm 0.002_{\text{sys}} \pm 0.016_{\text{stat}}) \%. \tag{2.69}$$

This gives an atomic hydrogen concentration of

$$[H]_{\text{atomic}} = (0.105 \pm 0.001_{\text{sys}} \pm 0.008_{\text{stat}}) \%. \tag{2.70}$$

This result is higher than the value obtained by N. Hild (0.18 % molecular HD, 0.09 % atomic H), with a difference of less than 2σ.
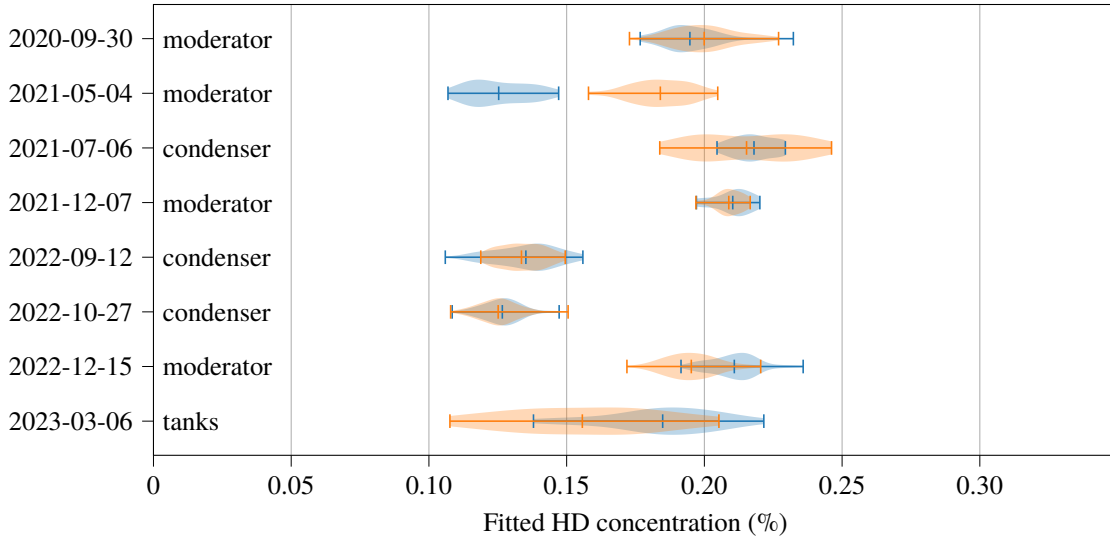
Figure 2.13: The distribution of HD concentrations estimated from different deuterium samples. The spectra we used in the calculation were sigma-clipped with a group size of 50. The estimated HD concentration depends on whether the HD transition line strengths were estimated using the summing method (blue) or the Gaussian fit (orange). The calculated concentrations tend to cluster around 2.1×10^{-3}. There were some issues with background over-subtraction around several HD peaks for the sample taken at 2021-05-04. This led to an underestimation of the HD concentration when the summing method was used. The estimated HD concentration in the two samples taken from the condenser at 2022-09-12 and 2022-10-27 is lower than in the other samples.

An HD concentration of 0.210% contributes $4.5\,\mathrm{s}^{-1}$ to the UCN loss rate in the moderator. Assuming an effective moderator temperature of 6.33 K and a p-D_2 concentration of 0.10%, the total loss rate in the deuterium is $30.8\,\mathrm{s}^{-1}$. The losses in HD are then 14% of the total UCN losses in the moderator. We decided that this does not warrant the investment of efforts and money to replace the deuterium in the UCN source by deuterium with a lower H concentration.

2.10   Discussion

We found that the para-deuterium fraction in the UCN source moderator varies from about 2.5% after startup of the UCN source to 0.5% at the end of the year. For para-deuterium concentrations below 1.0%, the para-to-ortho conversion upscattering of UCN due to para-deuterium happens at a lower rate than the inevitable losses due to thermal upscattering and neutron capture in deuterium.

The timescale for the ortho-to-para conversion in the sample bottle at room temperature has been determined to be $(157 \pm 6)\,\mathrm{d}$. The uncertainty is from the fit only. The conversion rate in the bottle depends on how well the bottle has been pumped.

The timescale for the para-to-ortho conversion in the condenser, at 6 K has been determined to be $(100.3 \pm 2.5)\,\mathrm{d}$. This is between the values reported in the literature for solid deuterium at 4.2 K, 79 d, and liquid deuterium at 20.4 K, 129 d.

The molar HD concentration in the UCN source has been measured to be $(0.210 \pm 0.002_{\mathrm{sys}} \pm 0.016_{\mathrm{stat}})\%$. conforming to the previous value of 0.18% [84] within 2σ. The UCN capture rate in hydrogen accounts for around 14% of the total UCN losses inside the deuterium. This estimate assumes the absence of high Z contamination in the source and does not take into account UCN interactions with the frost layer on the surface of the solid deuterium. Based on these results, we conclude that it is unnecessary to replace the deuterium

in the UCN source with higher-purity deuterium, which would be expensive and technically difficult.

CHAPTER 3

CONDITIONING OF THE UCN SOURCE SOLID DEUTERIUM MODERATOR

3.1 Overview

Every proton beam pulse deposits about 10 MJ of heat into the spallation target. Part of this energy is transferred to the solid deuterium moderator by gamma rays and neutrons (on the order of $0.2\,\mathrm{J\,cm^{-3}}$ [84]), increasing the deuterium temperature and causing sublimation of the deuterium surface. This deuterium vapour rapidly desublimates back onto the moderator surface after the proton beam pulse has passed, due to the high cooling power of the UCN source. These cycles of sublimation and desublimation create a low-density, rough layer on the sD_2 surface (sometimes referred to as 'snow' or 'frost'). This rough structure backscatters some upwards-travelling UCN, reducing the UCN extraction efficiency of the moderator [48], [100].

The UCN output can be recovered by repairing the sD_2 surface in a process called conditioning [31], [101]. This process consists of a simultaneous application of heating and a reduction of the cooling power applied to the deuterium. The rough frost layers sublimate. When cooling power is gradually restored, the deuterium vapour desublimates and is deposited as a smooth layer.

Before 2021, this was done by stopping the proton beam pulses during conditioning and using the outer and inner lid heaters (30 W in total) to warm the D_2 surface. This procedure requires an operator to be present in the UCN control room, which makes it impractical to do this form of conditioning outside normal working hours. The procedure takes about 2 to 3 hours [100], during which any UCN experiment is interrupted. The evolution of pressure in the moderator is shown in Figure 3.1, during this conditioning, the UCN source downtime was 135 minutes.

We developed a new method: conditioning with pulses. This method lasts only one hour: the cooling power is ramped down over twenty minutes, kept at a low state for twenty minutes, and ramped up again over twenty minutes. Proton beam pulses provide the required heating, so the electric heaters are not used. Conditioning with pulses takes less time and does not interrupt UCN production, although it temporarily reduces the UCN yield due to increased thermal upscattering. Some time series of the set value of the cooling valve, the vapour pressure in the UCN moderator vessel, and the UCN production are shown in Figure 3.2. Conditioning with pulses was first used in autumn 2021 while taking data for the mirror neutron experiment [96]. At this time, it was started by the UCN source operator, on request.

During conditioning with pulses, it is not required for an operator to be present in the UCN control room, it can be fully automated. This gives us the flexibility to do conditioning as often as we need to provide the maximum UCN output over time, even at night, or on weekends. In autumn 2022, we tested whether automated conditioning is reliable enough to be used as a mode of operation for n2EDM.
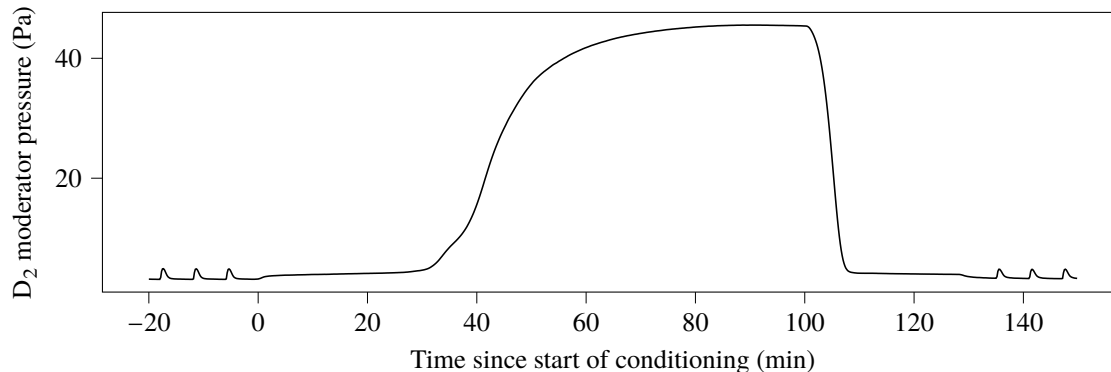
Figure 3.1: Evolution of D_2 moderator pressure during conditioning with heaters. The small peaks in the regions $T < 0$ and $T > 135\,\text{min}$ are the result of heating from UCN pulses. The UCN source downtime during this conditioning cycle was approximately 135 minutes.

We investigated the improvement of mean UCN yield over time when we used conditioning with proton beam pulses, either automatic or manual, compared to manual conditioning with heaters. We found that automated conditioning with pulses can yield an improvement up to 41 % over manual conditioning with heaters. The actual improvement depends on the amount of beam charge that the UCN source has received since the deuterium moderator was frozen. The optimal conditioning frequency ranges from once per day to four times per day. We investigated the reliability of automated conditioning with pulses and concluded that this is a feasible *modus operandi* for the UCN source during the n2EDM experiment. Further, we looked at the reproducibility of the conditioning and how much this depends on the timing of the start of conditioning relative to the timing of the proton beam pulses or on the pressure in the moderator vessel reached during conditioning. Within the available data, we did not find any dependence of the UCN yield after conditioning on these factors. Finally, we investigated whether the improvement in the decay rate of the UCN yield over weeks or months could be due to improved thermal conductivity of the deuterium. While the sensor data of the UCN source cooling system are not trivial to interpret, the evidence is compatible with this hypothesis.

3.2   Test of conditioning with pulses

In November 2022, we conducted a test to assess the reliability of automated conditioning. That is, the ability of automated conditioning to consistently recover UCN yield after some time of pulsing. For several days, we maintained constant UCN production, while scheduling automated conditioning with pulses every eight hours. To assess the benefit of automated conditioning with pulses, we will compare this series of data with data we took in June 2022, where we measured the decay of UCN output over a time without conditioning.

In both series of measurements, small Cascade detector B was mounted on beamport South, using the NiMo coated beam guide adapter (an aluminium tube, with a conical segment, reducing the inner diameter from 130 mm to 100 mm, the diameter of the flange on the small Cascade detectors). The superconducting magnet on beamline South was ramped up to 5 Tesla between the first and the second series of measurements. This led to a reduction in the UCN rate measured on beamport South. The experimental parameters are listed in Table 3.1.

3.2.1   Analysis of UCN data

The analysis was conducted on the UCN measured at beamport South. The ratio of UCN counts measured at beamport West-1 and South has been previously shown to be stable with a standard deviation of 0.14 % over several days of operation (see chapter 5).
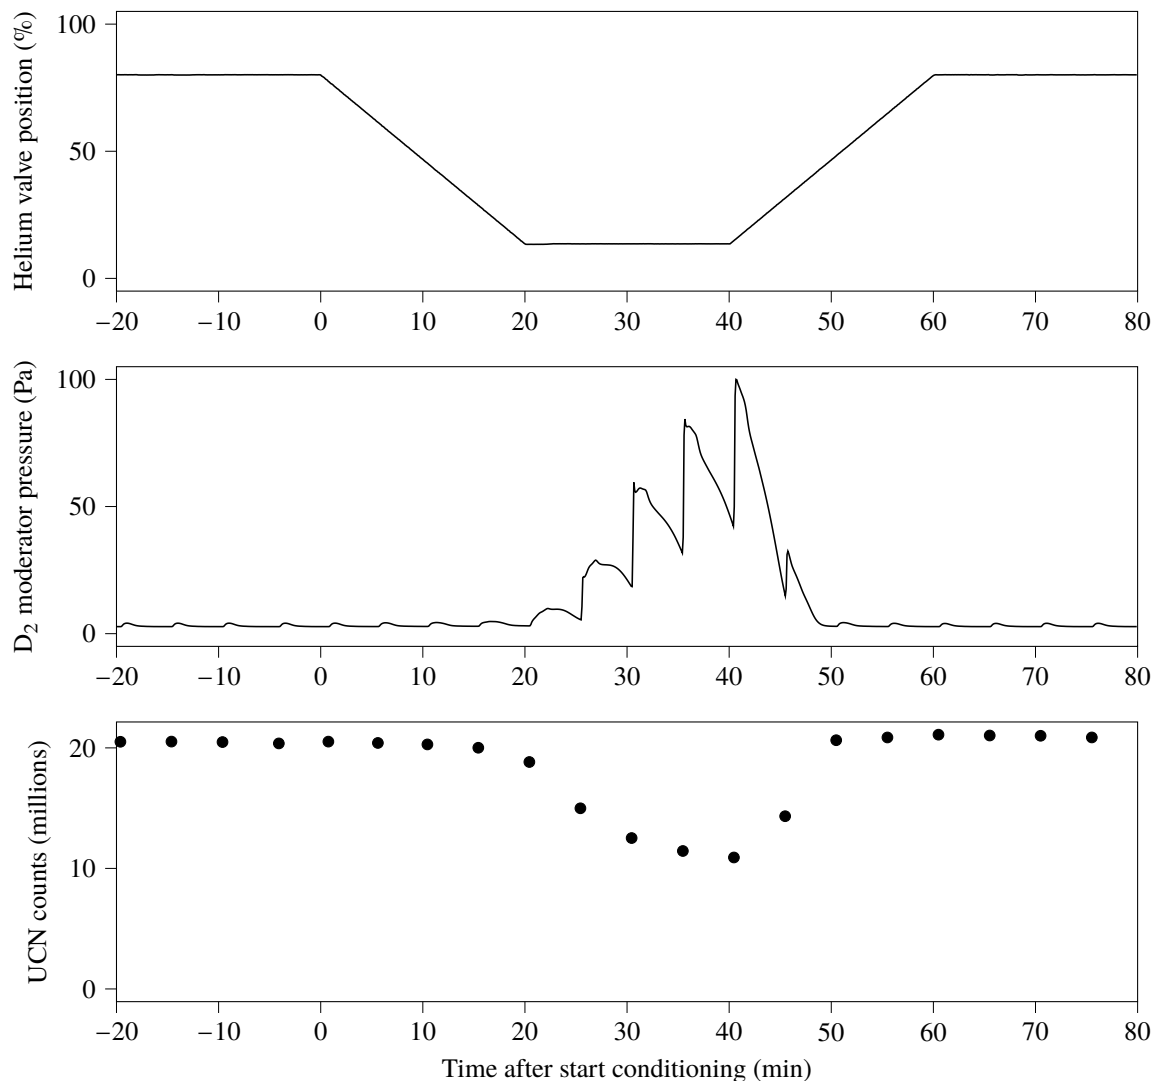
Figure 3.2: Evolution of helium cooling power, moderator pressure and UCN yield during conditioning with pulses. Top: The position of valve 2KJH10 AA003, controlling the helium coolant flow, is shown (100 % means the valve is fully open). Middle: The pressure in the moderator, measured by 2RNK10 CP006, is shown. The sharp pressure increases are the result of proton beam pulses. Bottom: The UCN counts in detector West-1, summed from 12 to 200 seconds after the pulse, their position along the horizontal axis indicates the start time of the corresponding pulse. Initially, the valve is 80 % open. The conditioning starts at $t = 0$, the valve is slowly closed to 13.6 % at $t = 20$ min. The valve stays partially closed until $t = 40$ min. During this time, the deuterium accumulates heat from proton beam pulses, and the pressure rises more with each pulse, as the deuterium sublimates. The UCN output is reduced by 40 % to 50 %, which can be explained by increased thermal upscattering, at an increase in deuterium temperature of 4 K to 5 K [84]. From $t = 40$ min to $t = 60$ min, the valve is opened again, to cool the deuterium. The pressure drops, and the UCN counts recover, indicating that the deuterium vapour has desublimated onto the moderator surface.

Table 3.1: Experimental parameters for the conditioning experiments in 2022.

Run	Without conditioning	Automated conditioning with pulses
Start time	2022-06-24 18:00	2022-11-02 15:05
End time	2022-06-27 09:30	2022-11-06 23:05
Duration	63.5 h	104.0 h
Conditioning	Once before start	Every 8 h
Detector West-1	Big	Big
Detector South	Small B	Small B
Mounting tube South	NiMo coated Al adapter	NiMo coated Al adapter
Superconducting magnet	0 T	5 T

The UCN source produces a lot of neutrons with energies above the Fermi potential of NiMo (220 neV). They can only be reflected under small grazing angles. This means they have a non-zero chance of passing through the UCN guide system and reaching the detector, but they have a small lifetime (seconds) in the UCN source and guide system [102]. They are thus unstorable on the timescale of typical UCN experiments (minutes). During the proton beam pulse, these unstorable neutrons contribute to the count rate in the detector. The count rate of unstorable neutrons diminishes quickly after the end of the pulse and becomes negligible approximately 4 seconds after the end of the pulse. Since we care about storable UCN, we conducted the analysis on the UCN counted between 12 and 200 seconds after the start of every pulse.[1]

Normalisation by proton beam current

The UCN yield is proportional, among other things, to the product of the proton beam current on the spallation target, and factors related to the state of the deuterium moderator. Since the latter is what we want to probe, we divide out the integrated beam current delivered onto the spallation target, which is not constant.

For every pulse, the beam-normalised UCN yield, U_b, is defined as:

$$u_{b,i} \equiv \frac{\int_{12\,\mathrm{s}}^{200\,\mathrm{s}} dt_i\, r_i(t_i)}{\int_{0\,\mathrm{s}}^{8\,\mathrm{s}} dt_i\, I(t_i)}, \tag{3.1}$$

where r is the UCN count rate (in counts per second) in the Cascade detector, I is the proton beam current (measured by beam monitor MHC1). t is the time after the start of the pulse, and i is the pulse index.

Since the conditions for the two measurement runs were not the same (Table 3.1), to compare the loss of UCN output caused by the degradation of the D_2 surface, we should consider the UCN yield in that run as a relative to the maximum UCN yield produced in that run. To be less sensitive to a single outlier, we divide by the average of the ten largest beam normalised UCN yields in that run. For each run, the sequence of relative UCN yields, u_i, is then defined as

$$u_i \equiv \frac{u_{b,i}}{\frac{1}{10} \sum_{j=0}^{9} \hat{u}_{b,j}}, \tag{3.2}$$

where $\hat{u}_j$ are elements of the sequence of beam normalised UCN yields for that run, reindexed in monotonically decreasing order.

The relative UCN yields are displayed in Figure 3.3. We lose about 50 % of the usable UCN output if we do not apply conditioning for the duration of one weekend (63.5 h) (the orange, smooth line). When applying automated conditioning with pulses approximately every 8 h, we observe a small decrease in UCN yield between conditioning, followed by a short but sharp decrease in yield during conditioning (the blue line). This is described in detail below.

---

[1]This interval is defined as $[12\,\mathrm{s}, 200\,\mathrm{s})$
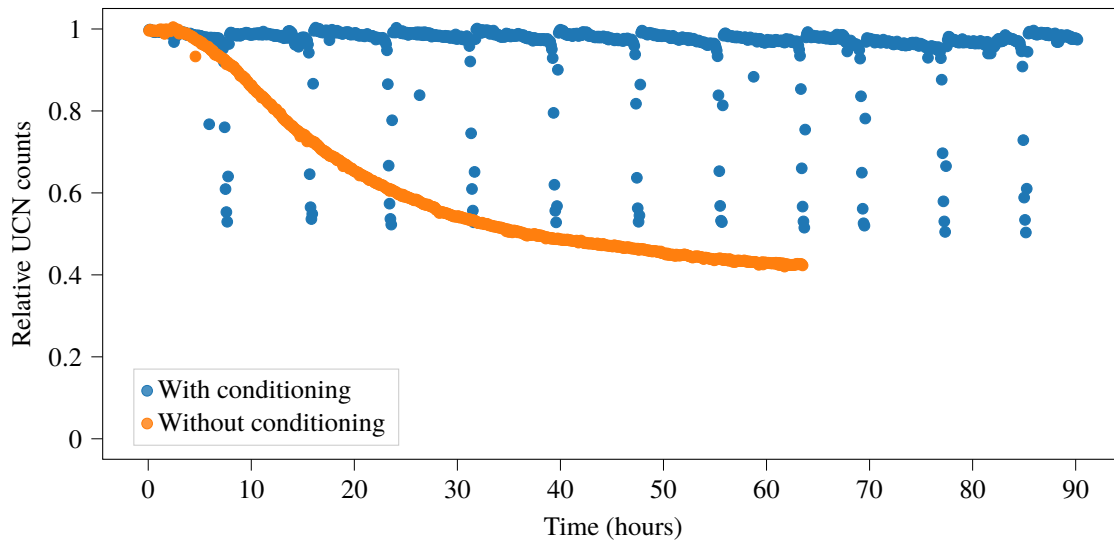
Figure 3.3: Evolution of the relative UCN yield with and without conditioning. We receive one UCN pulse every 5 minutes, 288 pulses per day. If we do not apply conditioning, we lose about half the UCN output over a weekend. The output is maintained when we apply conditioning every eight hours.

Estimation of expected UCN yields in selected scenarios

We used the data obtained in the two experiments, to estimate the UCN yield, averaged over a week, in different scenarios.

UCN yield for manual conditioning with heaters The scenario used for the calculation is that conditioning is done once per workday, always at the same time. We assumed conditioning lasts for 3 h. Four out of five times per week, conditioning is followed by a 21 h production period, and on Friday, it is followed by a 69 h production period (since no conditioning takes place during the weekend). From the data, it follows that the average UCN yield during the first 21 h after conditioning is 0.840, and for 69 h, this is 0.593 (using linear extrapolation, since the data go until 63.5 h). It follows that the weekly average yield is 0.664, see Table 3.2.

UCN yield for automated conditioning with pulses Over the entire analysis period (104 h), the mean relative UCN yield during the run with automated conditioning was 0.958. (This is compatible with the mean relative UCN yield in the first 7 h of the run without conditioning: 0.979.)

The mean relative UCN yield while not conditioning was 0.979. During conditioning, the relative UCN yield dropped to 0.522. Averaged over the duration of conditioning, the relative UCN yield were 0.809. Together, the average over 8 hours is 0.957. This result is shown in Table 3.2.

UCN yield for manual conditioning with pulses The calculation of the mean relative UCN yield in this scenario is similar to the scenario with manual conditioning with heaters. The difference is that conditioning with pulses only lasts 1 h. This means we have more production time, and a higher mean relative UCN yield.

The expected mean relative UCN yield during conditioning is a function of the yield when starting conditioning. This becomes important when the production time is long. To estimate this, I interpolated between the yield during conditioning measured in the test of automated conditioning with pulses in November 2022, and data from a conditioning we did at 2023-12-05 10:00, with an initial relative yield of 0.38, and a mean relative yield of 0.49.

Optimising the conditioning frequency   The optimal frequency for conditioning depends on whether the UCN produced during conditioning are useful for the experiment. Two numbers are given: the optimal conditioning frequency including the conditioning pulses, and the optimal frequency excluding these pulses.

We fitted a cubic spline to the UCN yield and calculated the optimal conditioning frequency from the fit. This ensures that the optimisation procedure is not sensitive to short-term variation in the UCN yield between pulses, but only to the long-term drifts. It also means that the procedure is not sensitive to missing data, for example, due to crashes in the detector software.

The mean yield during production is

$$\langle u \rangle_p(n) = \frac{1}{n} \sum_{i=0}^{n} u_i, \tag{3.3}$$

where $n = T/(5\,\text{min})$ is the number of pulses in the production period. The mean yield during conditioning, $\langle u \rangle_c(n)$ is given by the interpolation function described above.

The mean yield of both the production period and the conditioning period is then given by

$$\langle u \rangle(n) = \frac{n}{n+12} \langle u \rangle_p(n) + \frac{12}{n+12} \langle u \rangle_c(n). \tag{3.4}$$

If the UCN produced during conditioning do not contribute to experimental statistics, the mean yield is instead

$$\langle u \rangle(n) = \frac{n}{n+12} \langle u \rangle_p(n). \tag{3.5}$$

The results are shown in Figure 3.4.

If the UCN produced during conditioning are used to generate experimental data, the optimal production time between conditioning is 5 h, with a mean yield of 0.960 of the maximum. If the UCN produced during conditioning are not used, the optimal duration between conditioning is 10 h, with a mean yield of 0.866. See a detailed description in Table 3.2.

The curves of the mean UCN yield as a function of the duration between conditioning (Figure 3.4) are flat around the maximum, with a steeper slope on the left. This suggests that if the optimal conditioning interval is not known accurately, it is safer to overestimate it than underestimate it.

Discussion

To estimate the improvement in average UCN output gained by implementing automated conditioning, we take the ratio of the average normalised UCN output calculated in the previous two sections. With automated conditioning, averaged over a week, we would get 1.440 ± 0.005 times as many UCN as in the case of manual conditioning. The summary of the analysis is shown in Table 3.3.

The caveat is that we are comparing data taken in November with data taken in June. Immediately after freezing the deuterium, the UCN source output deteriorates more quickly, than after weeks or months of operation. The optimal conditioning frequency thus varies over a year.

3.3   Optimal conditioning frequency and integrated beam current

The calculation of the optimal UCN yield in the preceding section used the UCN yield degradation from a single series of data, taken shortly after the startup of the UCN source, and the freezing of the deuterium. Based on our experience of operating the UCN source, the degradation rate of the UCN yield is high after the deuterium has just been frozen, and it decreases over weeks or months of operation. Since the optimal

Table 3.2: Relative UCN yields, calculated for different types of conditioning.

Conditioning type	When	Situation	Duration (h)	Relative UCN yield incl. cond.	Relative UCN yield excl. cond.
Manual conditioning with heater	Weekday	Conditioning	3	0	0
		Production	21	0.841	0.841
		Subtotal	24	0.736	0.736
	Weekend	Conditioning	3	0	0
		Production	69	0.593	0.593
		Subtotal	72	0.568	0.568
	Week	Total	168	0.664	0.664
Manual conditioning with pulses	Weekday	Conditioning	1	0.617	0
		Production	23	0.822	0.822
		Subtotal	24	0.803	0.788
	Weekend	Conditioning	1	0.509	0
		Production	71	0.588	0.588
		Subtotal	72	0.583	0.580
	Week	Total	168	0.708	0.699
Automated conditioning with pulses	Any day	Conditioning	1	0.809	0
		Production	7	0.979	0.979
		Total	8	0.957	0.856
Automated conditioning optimised for UCN yield including pulses	Any day	Conditioning	1	0.809	0
		Production	5	0.990	0.990
		Total	6	0.960	0.825
Automated conditioning optimised for UCN yield excluding pulses	Any day	Conditioning	1	0.750	0
		Production	10	0.953	0.953
		Total	11	0.935	0.866

Table 3.3: Improvement of UCN yield compared to manual conditioning with heater.

Conditioning type	Conditioning frequency	UCN yield improvement incl. cond.	UCN yield improvement excl. cond.
Manual, with heater	Once per workday	0%	0%
Manual, with pulses	Once per workday	7%	5%
Automated, with pulses	Once per 8 h	44%	29%
Automated, with pulses	Once per 7 h	45%	24%
Automated, with pulses	Once per 11 h	41%	31%
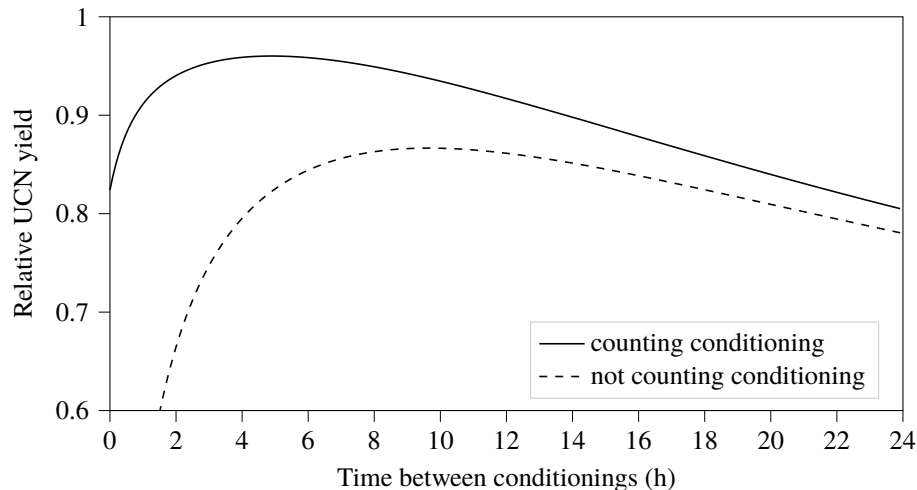
Figure 3.4: The expected relative UCN yield varies as a function of the time between conditioning. The optimal time between conditioning is longer if the UCN produced during conditioning are not used. The optimum time between conditioning, rounded to the hour, is 5 h if the UCN produced during conditioning are used, and 10 h if they are not used.

conditioning interval depends on the degradation rate of the UCN yield, the optimal conditioning becomes longer throughout the year.

To quantify how the degradation of the UCN yield and the optimal conditioning frequency depend on the integrated beam current, I applied the methods described in the previous section to the mirror neutron dataset.

3.3.1   The mirror neutron dataset from 2021

In 2021, the UCN group at PSI conducted an experiment to detect the anomalous loss of neutrons as a function of the magnetic field, in the search for neutron to mirror-neutron oscillations [48], [96]. The UCN data from this experiment can be used to investigate the change in UCN yield of the UCN source over a large time span.

Mirror matter has been proposed as a dark matter candidate, and an explanation of the baryon asymmetry of the universe [103], [104]. According to this hypothesis, there exists a copy of the standard model that does not interact with ordinary matter through the electroweak and strong interactions. The theory allows for oscillations between neutrons and mirror neutrons, which should have a resonance at the magnetic field strength for which the sum of the mass and potential energy of a neutron would equal that of a mirror neutron.

An overview of the experiment is shown in Figure 3.5a, taken from [96]. A stainless steel storage vessel was filled with UCN, during the 50-second filling phase, after which the beamline shutter was closed. A shutter, at the bottom of the storage vessel, was then opened for a 10-second monitoring phase, to count the number of UCN in the chamber. For 180 seconds, all shutters were closed, storing the UCN in the vessel. Afterwards, in the counting phase, the shutter at the bottom was opened, to count the remaining UCN in the vessel. The storage time of UCN in the vessel was computed from the ratio of UCN count rates in the counting and monitoring phases. Helmholtz coils, surrounding the vessel, were used to vary the magnetic field.

(a)

(b)
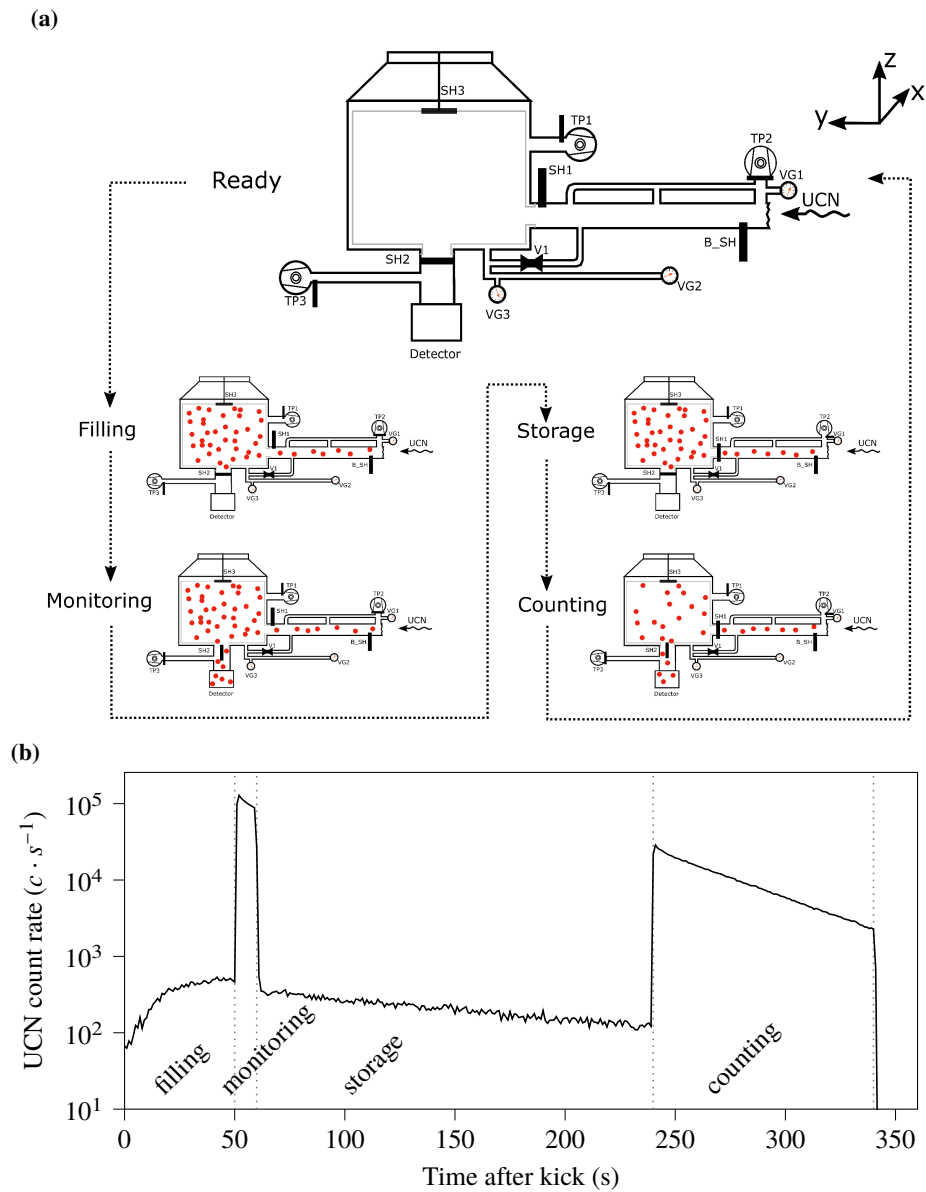
Figure 3.5: Operation of the mirror-neutron experiment. (a) A schematic of the experiment (published in [96]), showing the phases of one cycle. During filling (0 to 50 s), SH1 is closed and SH2 is opened, so UCN enter the storage vessel but do not fall into the detector. During storage (60 to 240 s), both shutters are closed. During monitoring (50 to 60 s) and counting (from 240 s), SH2 is closed to prevent backflow to the UCN source, and SH1 is open to count the UCN. The storage time is calculated from the ratio of count rates between the monitoring and counting phases. (b) The UCN count rate during the different phases. The count rate is not zero during filling and storage, SH2 does not stop all UCN.

3.3.2 Data reduction

Identification of conditioning and freezing times We determine the integrated beam current on target during every pulse from the proton beam current measured by HIPA beamline monitor MHC1. We sum up the integrated beam current from a specific point in time until every pulse to compute the cumulative beam current. We developed a method to detect the moments when the deuterium moderator was frozen or conditioned so that the cumulative beam currents can be computed from these moments forward.

When the deuterium is not frozen, the pressure in the moderator is above the triple point of deuterium, 171 mbar. By finding the moment when gauge 2RNK10 CP009 indicates a reading lower than 171 mbar, we can automatically determine the moment the deuterium is solidified. The cumulative beam current for the first proton beam pulse after freezing is thus 0.

During conditioning of the UCN source deuterium, the pressure inside the moderator according to 2RNK10 CP006 rises above 23 Pa. The finishing point of conditioning is defined by the time when the helium valve is opened to the normal operating position (80 % open). During conditioning with pulses, this occurs approximately 15 minutes after the pressure falls below 23 Pa again (see Figures 3.1 and 3.2). For the automatic detection of conditioning, the time when conditioning is finished was defined as 15 minutes after the moderator pressure (measured by 2RNK10 CP006) drops below 23 Pa. These timestamps were then used to calculate the cumulative beam current since conditioning, for every subsequent pulse. In the case of conditioning without pulses, the only requirement is that this timestamp is between the stop and restart of the production pulses.

The times of freezing and conditioning obtained using this method were verified with the written logbook. For every pulse, the cumulative beam current up to that moment was then calculated starting from the moment the deuterium was frozen, and from the end of the last conditioning. The UCN data were partitioned into series based on between which two conditionings they were collected.

Beam normalised UCN yield The beam normalised UCN yield (u_b) was calculated like in Equation 3.1, starting the integration at 12 seconds, but integrating up to 300 seconds (instead of 200 s). Since the mirror neutron experiment stored the neutrons until 240 s after the pulse, a significant fraction of UCN were counted after 240 s.

Data cleaning All the measurements where the beam current was below 95 % of the average beam current were excluded. The comparison of the UCN yield from different pulses is only meaningful if the measurements were conducted in the same way. Measurements following the standard mirror neutron measurement procedure (illustrated in Figure 3.5) were analysed. Measurements following non-standard measurement sequences, such as direct shots into the Cascade detector without closing the shutters, were discarded. For a measurement to be included, the UCN yield from 12 to 200 seconds must be between 6.00×10^5 and 1.285×10^6 counts. The ratio of counts from the periods [0 s, 12 s], [0 s, 300 s], had to be between 0.006 and 0.012. If fewer than 100 pulses took place between conditionings, this whole series of pulses was excluded, since the dataset was too small to fit a meaningful curve.

For each series separately, the normalised UCN yield was calculated from the beam normalised UCN yield, using Equation 3.2. One expects the highest normalised UCN yield just after conditioning. A few series started with so-called norm pulses, which have a duration of two seconds, and are done with the flapper valve in the UCN source opened. These were excluded.

Lastly, to detect remaining outliers, the relative difference between the normalised yield obtained during a pulse and its neighbours was defined as

$$D_i = \frac{u_i - \frac{1}{2}(u_{i-1} + u_{i+1})}{\frac{1}{3}(u_{i-1} + u_i + u_{i+1})}. \tag{3.6}$$

All counts for which $|D_i|$ exceeds 0.12 were excluded.

Table 3.4: Inclusion criteria for the mirror neutron dataset

Inclusion criterion	Pulses excluded by criterion	Remaining pulses
UCN data from 2021-07-14 to 2021-10-29		18 816
Sufficient beam current (> 95 % of average)	73	18 743
Counts from 12 to 200 s within range	2562	16 181
Expected ratio counts from 0 to 12 s to 0 to 300 s	47	16 134
Enough pulses in series	130	16 004
Normalised yield > 0.9, in first 100 pulses	31	15 937
Relative difference with neighbours < 0.12	391	15 582

The full list of inclusion criteria and the number of excluded and included pulses is shown in Table 3.4. Importantly, the calculation of cumulative beam current was done before the data cuts and therefore not affected by the cuts.

3.3.3 Spline fitting

To each series, a model was fitted to predict the normalised UCN yield for each pulse, as a function of the integrated beam current since conditioning:

$$u = f(C, \vec{\beta}) + \epsilon, \tag{3.7}$$

where u is the normalised UCN yield, C is the integrated beam current (i.e. charge) on target since the last conditioning, $\vec{\beta}$ is the set of fit parameters, and ϵ is the fit error.

This smooths out the statistical noise in the data and allows us to substitute interpolated values for missing data (excluded during data cleaning, or due to a crash of the detector software).

To preserve generality, a spline model was used. While lacking intuitive explanatory power, this model could be more successfully fitted to the series with relatively few parameters. A set of points is defined, (C_s, u_s). The C_s values are pre-defined metaparameters, and the u_s values are fitted such that cubic interpolation between these points minimises $\Sigma_i \epsilon_i^2$.

The Bayesian Information Criterion [105] was calculated for several fits, to optimise the number of points required for the spline to have a good fit to the data, while avoiding overfitting. The optimal spacing between the points $\Delta C = 100 \cdot 1.8$ mA $\cdot$ 8 s (i.e. approximately one point every 100 pulses). The number of fit parameters thus depends on the number of pulses in the series.

One of the series, together with the accompanying spline fit, is shown in Figure 3.6a. In Figure 3.6b, three fitted curves are shown, for representing series that were started at different times after the deuterium was frozen in the moderator. There is a large difference in the decay rate of the UCN source output.

3.3.4 Change of the UCN yield deterioration rate over many conditionings

Since the spline fits do not necessarily have the same shape, it is less straightforward to compare the parameters, to make a statement about the change in decay rate of the UCN source over time. We compared the change in UCN output of the source, at a specified time after the end of the last conditioning, for the different series.

For every series of pulses, we looked at the expected normalised UCN yield at 12, 24 or 48 beam hours after conditioning. In the rest of this chapter, one beam hour is defined to contain 12 pulses of 8 seconds, with an average beam current of 1.8 mA. It is thus a measure of cumulative beam current, not chronological time.

(a)

(b)
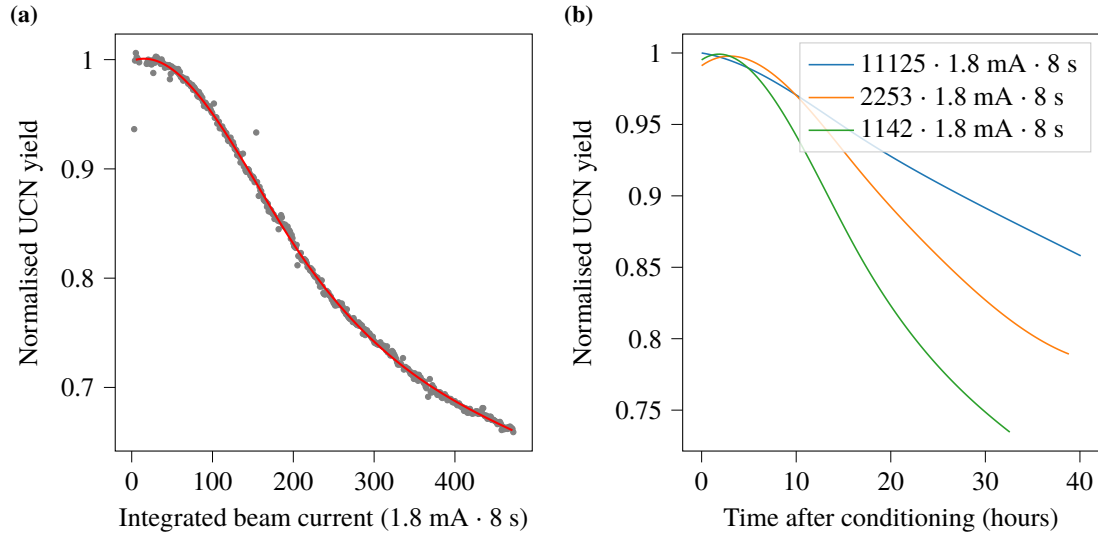
Figure 3.6: (a) The degradation of the normalised UCN yield, as a function of the integrated beam current since the last conditioning. The grey dots represent the data points, and the red curve the spline fit. Two outliers are visible, these were eliminated before fitting the curve. (b) Three fits, from series that were started at different times after the deuterium was frozen in the moderator. The degradation of the UCN yield happens markedly slower with higher integrated beam currents.

The normalised UCN yield after 12, 24 and 48 h is displayed in Figure 3.7. When the UCN source is started up, right after the deuterium has been frozen, the degradation rate of the UCN yield is high, indicated by a small UCN yield after 24 h. After several weeks of continuous operation, for every series the degradation rate of the UCN yield is much lower, indicated by a higher yield after 24 h. By fitting exponential curves, one can make a statement about the asymptotic behaviour.

The fits indicate that the expected normalised UNC yields, 12 h, 24 h or 48 h after conditioning, increase over weeks, and flatten off asymptotically. The fitted asymptotes that are indicative of the best (smallest) degradation rate that can be reached are

$$
\begin{aligned}
0.963 &\pm 0.005, & &12\text{ h after conditioning,} \\
0.914 &\pm 0.008, & &24\text{ h after conditioning,} \\
0.843 &\pm 0.032, & &48\text{ h after conditioning.}
\end{aligned}
$$

3.3.5   Optimal conditioning frequency

Since the rate of change in the normalised UCN yield decreases over time, so does the optimal frequency of conditioning. The mean yield over time, including the conditioning period, can then be calculated using Equations 3.4 or 3.5, depending on whether the UCN produced during conditioning are respectively used or not used. For some series, the estimated mean normalised UCN yield is shown in Figure 3.8.

The optimal conditioning interval is the one that maximises the mean normalised UCN yield over time. This was calculated for every series (see Figure 3.9. The interval between subsequent conditionings should increase if the UCN yield deterioration flattens off. By fitting exponential functions, we can estimate the asymptotic behaviour. When the UCN source is operated for a very long time, the optimal time between conditioning approaches (24.0 $\pm$ 1.6) h if the UCN produced during conditioning are discarded, and (10.1 $\pm$ 0.4) h if they are counted.
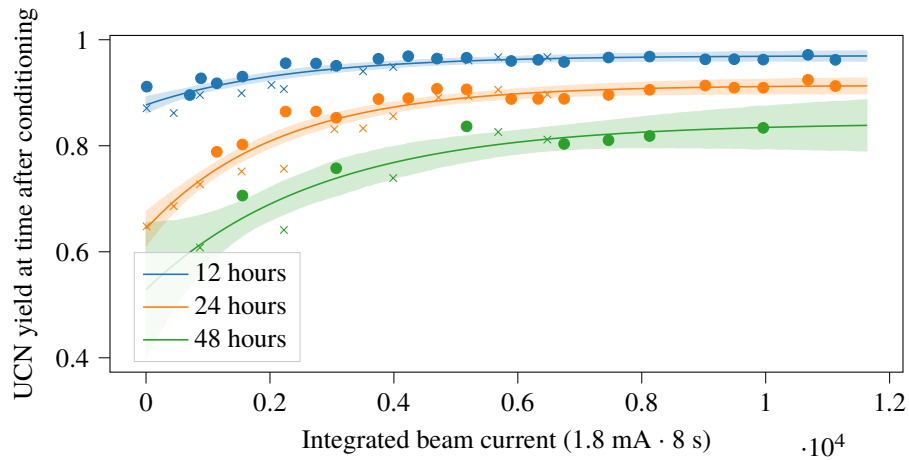
Figure 3.7: The normalised UCN yield at 12, 24 or 48 hours after conditioning. Every set of three points represents one series of pulses, between two conditionings. On the horizontal axis, the cumulative integrated beam current is shown, since the deuterium was frozen, at the start of the series. The deuterium was refrozen in the middle of the data taking. The data points marked with an × were taken after this. The curves are exponential fits, the corresponding shaded areas are the two-sigma error band of the fit. In the majority of series, conditioning was done before 48 beam hours had been reached, which is why there are fewer points for 48 h. This figure shows, how after weeks or months of operating the UCN source after freezing the deuterium, the degradation rate of the UCN yield is reduced.
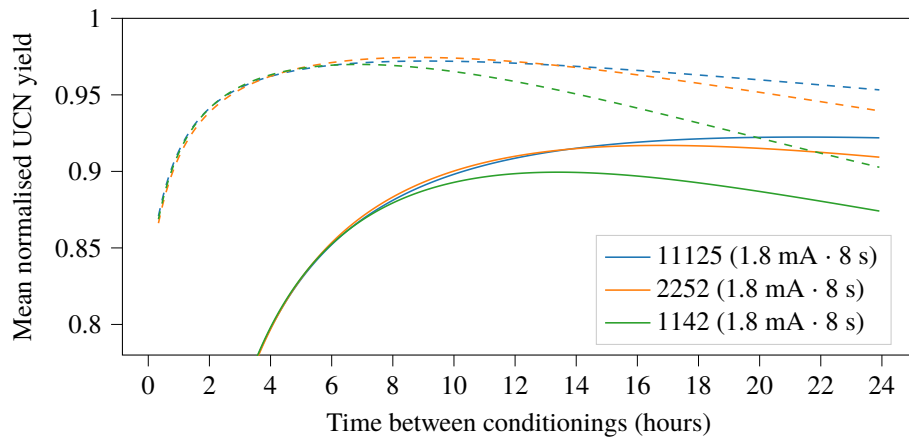
Figure 3.8: The mean UCN yield over time, as a function of the time between conditioning. The dashed lines represent the scenario where UCN pulses during conditioning are used, whereas the solid lines represent the scenario where they are not used. The series, for which this was calculated, are the same as in Figure 3.6b.

(a)

(b)
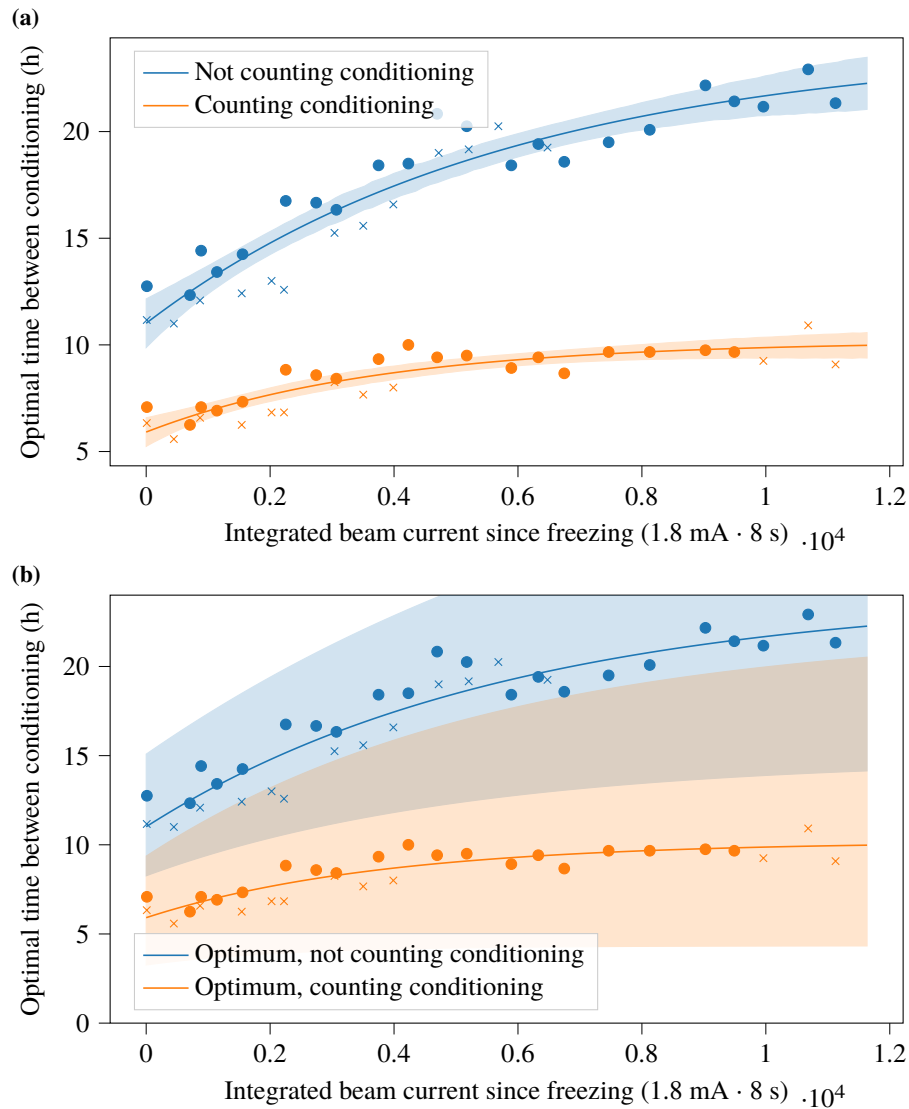
Figure 3.9: The optimal duration between two conditionings, as a function of integrated beam current since freezing. The curves represent exponential fits. In (a), the shaded areas represent the two-sigma error band of the fits. In (b), the shaded areas represent the flat area around the maximum, in which 99 % of the highest mean yield can be expected (outside this area, the mean yield is reduced rapidly). The deuterium was refrozen in the middle of the data taking. The data points marked with an × were taken after this.
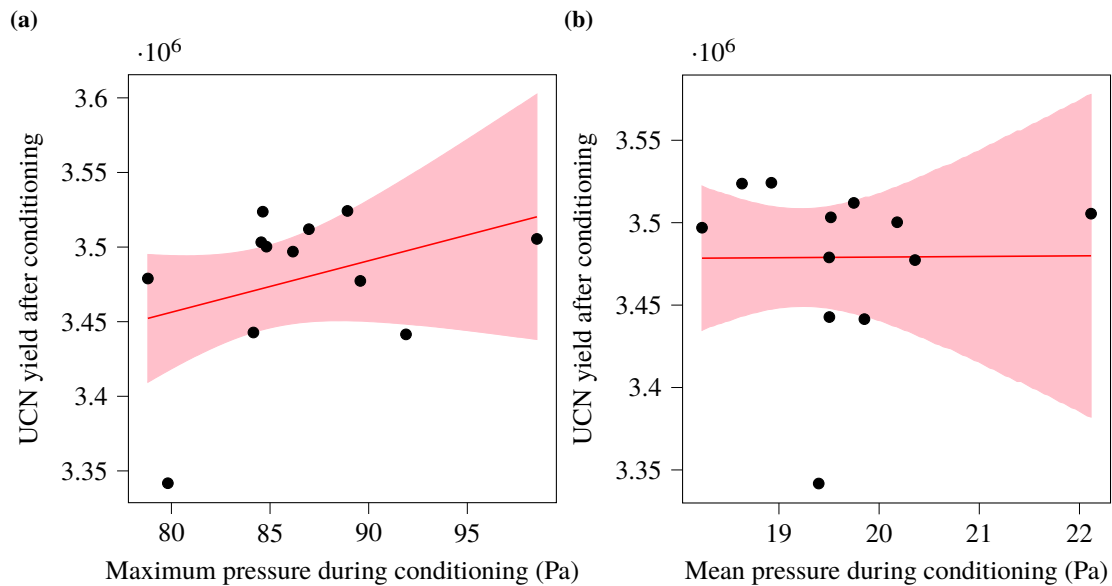
Figure 3.10: The relation between the mean of the twenty highest beam normalised UCN yields (12 s to 200 s) after conditioning, and (a) the maximum moderator pressure during conditioning, and (b) the mean moderator pressure during conditioning (1 h) on the other hand. Every black dot represents a mean of the twenty highest yields after one conditioning periods. The lines represent linear fits, and the shaded areas represent the 95 % confidence intervals for the fits. The correlations are not significant. The slope in (a) disappears completely if the outlier in the bottom left corner is removed. The data are from the automated conditioning test, done in November 2022.

3.4 Timing effects, pressure and UCN output after conditioning

When conditioning is done with pulses, in principle the UCN yield after conditioning could be a function of

- the start time of conditioning, relative to the timing of proton beam pulses,

- the pressure reached in the moderator vessel, during conditioning, and

- the number of missed pulses.

If the UCN output depends on any of these parameters, one needs to be aware of it to guarantee the effectiveness of conditioning.

The computer systems controlling and monitoring the UCN source are decoupled from the systems controlling the HIPA beamlines. In the past, the timing of the start of conditioning relative to the proton beam pulses has been random. To investigate whether it is worth the effort to synchronise the start of conditioning with the proton beam pulses, we investigated the data obtained in 2022 and 2023. From the available data, there is no evidence to suggest that mean UCN yield can be improved by synchronising the conditioning to the proton beam pulses.

In the data obtained during the automated conditioning measurement in November 2021, there was no significant correlation between the mean pressure in the moderator during conditioning, measured by 2RNK10 CP006, and the average of the 20 most productive UCN pulses in the three hours after conditioning (Fig. 3.10).

In the data from 2021 and 2022, there were no conditionings where proton beam pulses were missed. We lack the data to tell how much this would reduce the effectiveness of conditioning with pulses. This means that one need not worry about missed pulses from a conditioning standpoint.

3.5 Changes of the UCN yield degradation over operation time

As demonstrated in this chapter, the UCN yield decay rate between two conditionings decreases, when the UCN source is operated for a long time. This is a result of changes in the physical properties of the deuterium moderator. When the cooling of the UCN source moderator is interrupted, and the deuterium has to be refrozen into the moderator, these improvements disappear.

After the deuterium has been frozen inside the moderator, the ortho-to-para ratio slowly approaches its equilibrium (see Chapter 2). Natural self-catalysed equilibration is a slow process, with a rate around 6×10^{-4} h^{-1} [56], [84], [106]. During UCN source operation, neutrons convert para-deuterium to ortho-deuterium. Gamma rays break up deuterium molecules. Before the resulting free deuterium atoms recombine, they catalyse para-to-ortho conversion in other molecules [56]. Hence, the ortho concentration increases faster if more proton beam pulses are supplied to the UCN source. The radiation-induced conversion rate in the UCN source, for a beam current of 1.74 mA, was measured to be $(6.9 \pm 0.8) \times 10^{-7}\,\mathrm{s}^{-1}$, by N. Hild (Section 8.4.3 of [84]).

We hypothesise that the improvement in heat conduction (see Figure 2.2 in Chapter 2) caused by the increase in ortho concentration, causes heat to be removed from the deuterium more rapidly, leading to less sublimation of deuterium. If this is the case, the cooling power of the UCN source (the product of helium coolant flow and the temperature difference between in-flowing and out-flowing helium) should have a higher peak, and drop off more quickly after conditioning. In addition, the increase in pressure should be smaller.

3.5.1 Analysis

Three data series from the UCN source were used in the analysis, separated by several weeks. The first series was from 2021-09-22 14:00 to 20:00, the second from 2021-10-10 14:00 to 20:00, and the third from 2021-10-31 14:00 to 20:00. The sensors used were the helium flow sensor: 2KJH10 FF003, the thermometers measuring the in and out-flowing helium: 2KJH10 CT002 and CT003, and the pressure gauge measuring the moderator pressure: 2RNK10 CP006. The setup is described in detail in [48].

In every series, we averaged the sensor information over all the pulses. We estimated the helium flow rate through the moderator using the method described in [48]. Together with the temperature difference, we used this to estimate the cooling power. The results are plotted in Figure 3.11.

Several issues complicate the analysis. In [48], the cooling power was estimated during the freezing of the deuterium. There were no pulses, and thus no rapid changes in temperature or pressure. Conversely, during UCN pulses, the helium thermally expands and contracts. The flow sensor is located at the side where helium flows in. The backflow of expanding helium temporarily reduces the measured helium flow and the estimated cooling power. This makes it difficult to obtain an accurate estimate. Nevertheless, from 100 s after the pulse, the helium flow is reduced in the last series, and the estimated cooling power is also lower. This is compatible with the hypothesis.

Over the three series, the moderator pressure decreases progressively. The measured pressures from 240 to 300 s after the kick were $(3.44 \pm 0.05) \times 10^{-2}$ Pa lower for Series 2 than for Series 1, and $(5.76 \pm 0.06) \times 10^{-2}$ Pa lower for Series 3 than for series one (the reported uncertainty is the standard error of the mean).

The absolute value of the pressure is not relevant, because there is a non-zero amount of helium in the deuterium system, which offsets the pressure. However, the amount of helium typically increases over time, as it does not diffuse against a pressure gradient, and can be only removed by pumping. The presence of helium cannot account for the decrease in pressure over time.

The mean beam currents were (1.834 ± 0.003) mA, (1.854 ± 0.005) mA, (1.870 ± 0.002) mA (the reported uncertainty is the standard error of the mean). The decrease in pressure happened despite this small, but significant increase in beam current.
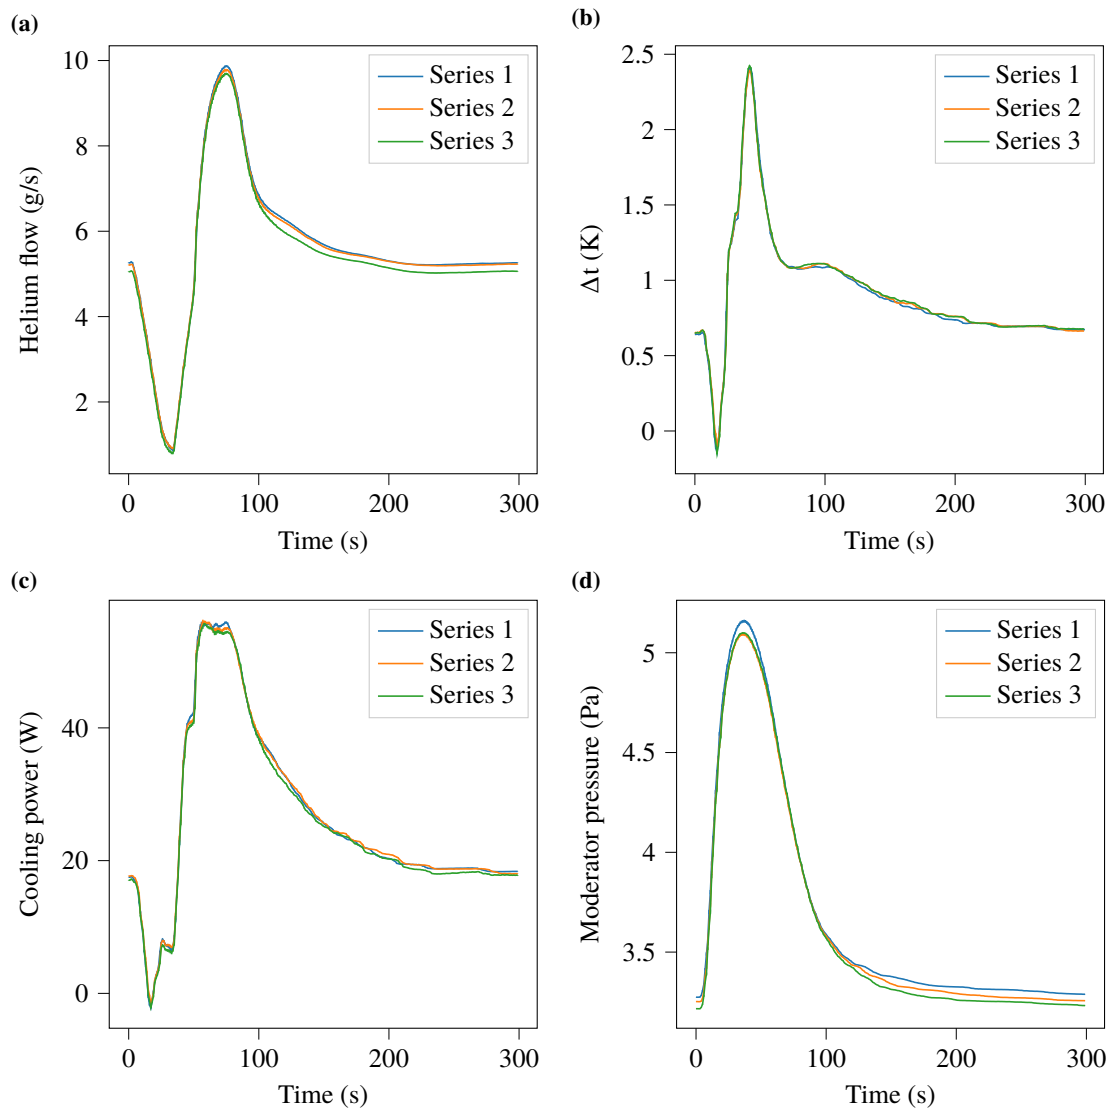
Figure 3.11: (a) The helium flow rate during the proton beam pulse. (b) The temperature difference between in and out-flowing helium during the pulses. (c) The estimated cooling power. (d) The measured pressure.
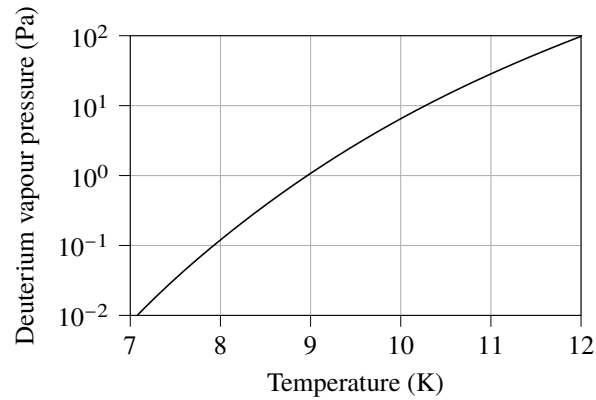
Figure 3.12: The vapour pressure of deuterium, in the 7 to 12 K range [56].

While 5.8 Pa seems like a small decrease in pressure, the vapour pressure of deuterium at 7.7 K is equal to 5.8 Pa (see Figure 3.12).

While it is difficult to extract the exact cooling power reliably, the pressure data add evidence in favour of the hypothesis that increased heat conductivity of the deuterium reduces the deterioration rate of the UCN yield of the UCN source over time.

3.6   Discussion

In this chapter, we have demonstrated that automated conditioning with pulses is reliable. Automated conditioning with pulses can be used in routine operation of the UCN source during experiments to maximise the average UCN yield over time. Automated conditioning with pulses significantly improves UCN output, compared to manual conditioning with pulses, or manual conditioning with heaters. Since it can be scheduled at any time of the day, the frequency can be optimised.

Assuming the UCN experiments are not taking data during conditioning, the optimal conditioning frequency for conditioning starts out from twice per day just after the deuterium has been refrozen in the moderator following shutdown of the UCN source. After at least a month of intensive UCN source operation, the optimal conditioning frequency would be reduced to about once per day. At which frequency conditioning will be done during n2EDM data-taking runs is still to be decided. This depends on whether useful experimental data can be obtained during conditioning.

In addition, we have shown that the timing of conditioning relative to the proton beam pulses and the possibility of having a missed pulse during conditioning are not of practical concern for the effectiveness of conditioning.

We set out to test whether the reduction in UCN yield decay rate over time over several conditionings is due to the increased thermal conductivity of the deuterium due to the improvement in ortho concentration. While it is difficult to get direct evidence about this from the cooling power, the pressure data support this hypothesis.

SOLID DEUTERIUM MASS IN THE MODERATOR VESSEL AND UCN INTENSITY

4.1 Overview

The UCN yield of the UCN source depends on the deuterium mass in the moderator vessel. Increasing the deuterium mass in the moderator vessel increases the total UCN production rate. The lifetime of UCN in the deuterium is finite (approximately 30 ms [107]) and depends strongly on the para-deuterium concentration and the temperature. The UCN extraction efficiency is lower for UCN produced in the lower strata of the moderator, as they need to traverse more deuterium before reaching the surface. Increasing the deuterium mass in the moderator thus leads to diminishing returns in UCN yield. In addition, the UCN production density scales negatively with height inside the moderator, since the thermal neutron flux approximately scales with the inverse squared distance to the centre of the UCN source spallation target. See for example Figure 4.16 of [48]. Above the optimal filling level of the moderator, the increased UCN production rate no longer outweighs the reduction in extraction efficiency.

In the years leading up to 2022, we had not yet reached the optimum filling level. In 2021, a record UCN yield was achieved by extracting deuterium adsorbed in the converter vessel and adding it to the moderator [48].

In this chapter, I describe how we filled the moderator vessel with 5.68 kg of deuterium, corresponding to an estimated filling level of 162 mm, in 2022. While this is the largest deuterium mass we have ever used for UCN production at PSI, it did not result in a record UCN yield. We observed that the pressure rise in the moderator vessel during proton beam pulses was larger for this deuterium mass, compared to smaller deuterium masses, even though the proton beam current was smaller. We suspect that the cooling effectiveness at the deuterium surface is reduced, resulting in more thermal upscattering.

4.2 Introduction

4.2.1 Norm pulses

We have been using norm pulses to compare the UCN source yield over years since the UCN source was taken into operation [45]–[48]. These are two-second pulses, during which the flapper valve of the UCN source stays opened, and NLK South is closed.[1] Because the pulses only last two seconds, it is not necessary

---

[1]Since 2019, the Böttstein flap stays closed, due to a leak in the pneumatic system.

to reduce the beam current before the proton beam is sent to SINQ again. By not operating the flapper valves, the resulting UCN yields are not affected by timing effects of mechanical components. The UCN yield is measured using the big Cascade detector, mounted on beamport West-1, with an adapter. The reported yields are counted from the start of the pulse (in contrast to, for example, the transmission measurements, where we analysed the UCN arriving in the interval from 12 to 200 seconds after the pulse). During the proton beam pulse, the proton beam current is measured using HIPA beam monitor MHC1. The normalised UCN yield is then computed using

$$\text{normalised UCN yield} = \text{UCN yield}\frac{2.2\,\text{mA}}{\text{beam current}}. \tag{4.1}$$

The beam current is normalised to 2.2 mA because this was the maximum beam current for standard UCN source operation when we started using norm pulses to compare the UCN source output over time.

In addition to norm pulses, we usually also compare production pulses, as these are the type of pulses used to provide UCN to experiments. Production pulses are 8 s pulses with normal operation of the flapper valve (closing after the pulse, opening 20 s before the next pulse). To get the highest UCN intensity in West-1, in this chapter, we compare production pulses where we kept NLK South closed.

4.2.2   Deuterium adsorption in the converter

The amount of deuterium in the UCN source is limited by the capacity and pressure rating of the storage tanks. Adding significant amounts of deuterium to the UCN source would require a redesign of the deuterium system.[2] In 2021, we developed a method to transfer a larger fraction of the deuterium already contained in the system to the moderator.

In the years until and including 2020, we managed to transfer up to 4.8 kg of deuterium from the converter to the moderator vessel, out of the total of 6.0 kg in the system. At least (1125 ± 22) g of deuterium remains adsorbed in the oxisorb® in the converter after the transfer of liquid deuterium from the converter to the moderator vessel at 20 K to 22 K. By warming up the converter, we can release more deuterium, at the cost of increasing the para-deuterium concentration. This was first tried in 2021 [48] and again in 2022.

4.3   The relationship between deuterium mass in the moderator and UCN yield

A simplified schematic of the compartments and valves of the deuterium system that are relevant in this chapter is shown in the introduction, in Figure 1.6. A detailed overview is shown in Appendix B.

4.3.1   Procedure for filling the moderator vessel with deuterium

In 2022, we filled the moderator vessel with the largest amount of deuterium so far using the following procedure.

- On 2022-04-20, we started freezing the deuterium from the storage tanks in the condenser. The condenser was cooled down to 5.9 K. On 2022-04-22, valve 2RNK10-AA001 was closed, interrupting the path to the storage tanks. The deuterium was warmed up to 19.6 K and melted. On 2022-04-25, the deuterium was transferred from the condenser to the converter to convert para to ortho-deuterium

---

[2]The deuterium system is equipped with rupture discs, that safely vent all the deuterium into the atmosphere, in case of overpressure. These discs rupture at 2.5 bar (2RNK10 BP901) and 2.8 bar (2RNK10 BP902, see Appendix B). Currently, when all deuterium in the UCN source is gaseous, the pressure reaches around 1.25 bar. We have to take a large safety margin, taking into account hot days, and uncertainty on the rupture pressure.

at 20.5 K. On 2022-04-27, the deuterium was transferred to the moderator. From 2022-05-5 to 2022-05-09, the converter was slowly warmed up to 32 K to transfer more deuterium to the moderator. We stopped the transfer when we had 5.28 kg of deuterium in the moderator.

- The UCN source had been operated for about a month and a half, increasing the ortho-deuterium concentration and improving the UCN yield. On 2022-08-22, we applied norm pulses and measured a record UCN yield. Normalised to 2.2 mA, the normalised yield was 6.17×10^6 UCN.

- A lot of deuterium was still adsorbed in the OXISORB in the converter. On 2022-08-24, we slowly heated the converter to 60 K, evaporating part of the adsorbed deuterium into the storage tanks. The mass of the released deuterium was determined by a pressure measurement.

- We decided to fill an additional 400 g of deuterium into the moderator vessel. On 2022-09-06, 400 g of deuterium was frozen from the storage tanks into the condenser. We made use of the natural para-to-ortho conversion in the condenser vessel, which was kept at 6.2 K until we were ready to transfer it to the moderator. We took samples from the deuterium in the condenser to determine the para-deuterium concentration using Raman spectroscopy on 2022-09-12 and 2022-10-27. We found that the ortho concentration had improved from (87.98 ± 0.15) % to (92.39 ± 0.10) % during this time (see Chapter 2).

- On 2022-11-02, we measured a record normalised UCN yield during norm pulses, of 6.49×10^6 UCN, still with 5.28 kg of deuterium in the moderator vessel.

- In November, we transferred the 400 g of deuterium from the condenser to the moderator vessel, adding up to a total of 5.68 kg. On 2022-11-07, we melted the deuterium in the moderator so that it would mix with the deuterium from the condenser and have a homogeneous ortho concentration. From 2022-11-08 to 2022-11-11, we transferred the deuterium from the condenser to the moderator, in gaseous form, via the valve 2RNK10 AA003, bypassing the condenser. From 2022-11-15, we started freezing the moderator again.

- To increase the ortho-deuterium concentration using radiation-induced conversion, we kept applying proton beam pulses from 2022-12-01 to 2022-12-14. We measured the UCN yield during norm pulses several times. The highest normalised yield we obtained was on 2022-12-14: 6.28×10^6, with 5.68 kg of deuterium. This is lower than the yield measured on 2022-11-02.

- On 2022-12-15, we melted the deuterium. We took a vapour sample to determine the ortho-deuterium concentration using Raman spectroscopy. Immediately afterwards, we evaporated the deuterium into the storage tanks to determine the deuterium mass in the moderator. Based on the pressure in the storage tanks of 1040 mbar, measured using gauge 2RNK10 CP001, it was determined that the total mass in the moderator was (5678 ± 10) g.

4.3.2 UCN yields

In Figure 4.1a, the normalised UCN yield obtained during norm pulses is shown for different dates. For comparison, two measurements were added, done at the end of 2021 with a moderator filling of 5.33 kg, which at the time resulted in a record UCN yield.

The next four measurements were done in 2022, with a deuterium mass of 5.28 kg. The first two were done at the beginning of the year and had relatively low UCN yields. The next two were done after operation of the UCN source for several months. Due to natural and radiation-induced para-to-ortho-deuterium conversion, and annealing of the crystal structure of the deuterium, the UCN yield is much higher. As of the writing of this text, the norm pulses done on 2022-11-02 produced the highest normalised UCN yield ever measured at the PSI UCN source.
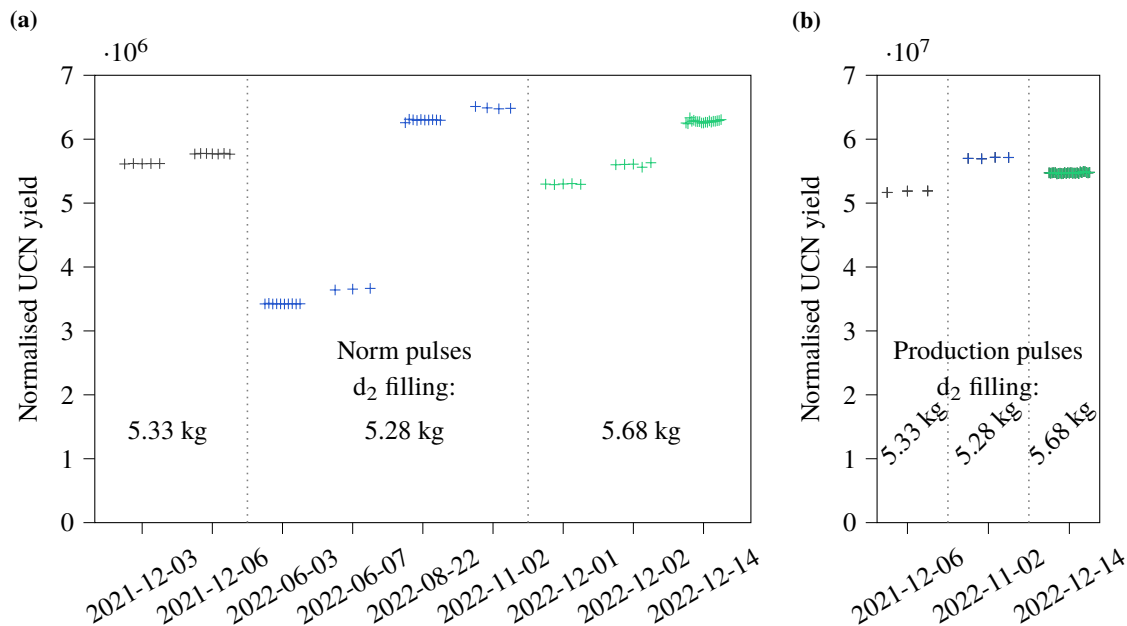
Figure 4.1: Normalised UCN counts for different deuterium masses in the UCN source moderator vessel using (a) norm pulses and (b) production pulses. The UCN yield was normalised to a beam current of 2.2 mA, with a duration of 2 s for norm pulses and 8 s for production pulses. The UCN were counted at beamport West-1 using the big Cascade detector. NLK South was closed during all measurements. The number of UCN counted during production pulses is about factor 10 higher than during norm pulses.

We then added the 400 g of deuterium stored in the condenser to the moderator vessel. The deuterium was melted and refrozen. When UCN source operation was resumed, the normalised UCN yield was significantly lower than on 2022-11-02. After continuous UCN source operation for approximately two weeks, the normalised UCN yield increased significantly but did not reach the level of 2022-11-02.

For the date with the highest UCN yield for a given deuterium mass (2021-12-06, 2022-11-02 and 2022-12-14) the average normalised yields during norm pulses and production pulses are plotted in Figure 4.1b, and displayed in Table 4.1. The yields on 2022-11-02 were the highest ever measured at the PSI UCN source, with a normalised yield of 6.49×10^6 during norm pulses and 5.71×10^7 during production pulses. The UCN yields were normalised to a proton beam current of 2.2 mA. The production yield is the most relevant number, as this is the yield supplied to the experiment, with the caveat that this depends on the filling time of the experiment. The UCN source at PSI is the most powerful UCN source on the planet currently in operation, and therefore the best facility to conduct the n2EDM experiment [108].

4.4 Para-deuterium concentration

Since para-deuterium causes upscattering of UCN, differences in para concentration contribute to the measured differences in the normalised UCN yield during the norm pulses. After the norm pulses supplied on 2021-12-06 and 2022-12-14, we took vapour samples from the deuterium in the moderator to determine the para-deuterium concentration using Raman spectroscopy (see Table 2.4). The concentrations were (0.45 ± 0.02) % and (0.57 ± 0.02) %, respectively (table 4.1).

We did not take a sample from the moderator after the measurement on 2022-11-02. To estimate the amount of para-deuterium added to the moderator from the condenser, one could extrapolate using the empirically found natural conversion rate in the condenser. This has to be subtracted. Radiation-induced conver-

Table 4.1: The maximum number of UCN measured during norm pulses and production pulses with different amounts of deuterium in the moderator vessel. The CN yields were scaled to a proton beam current of 2.2 mA. The pressure rise is displayed for norm pulses only.

Date	D_2 mass (kg)	Filling level[a] (mm)	Ortho concentration (%)	Norm pulse UCN yield[c] (UCN)	Pressure rise[b] (mA)	Production UCN yield[c] (UCN)
2021-12-06	5.330 ± 0.010	152	99.55 ± 0.02	5.77×10^6	0.44	5.18×10^7
2022-11-02	5.277 ± 0.014	151	99.00 ± 0.30	6.49×10^6	0.34	5.71×10^7
2022-12-14	5.677 ± 0.010	162	99.43 ± 0.02	6.28×10^6	0.42	5.47×10^7

[a] Uncertainty of 2 mm for all values.
[b] Uncertainty of 0.03 Pa for all values.
[c] Uncertainty smaller than the least significant digit.

sion (estimated rate for 1.74 mA of proton beam current for liquid deuterium at 20 K: $6.9 \times 10^{-7}\,\mathrm{s}^{-1}$ [84]) has decreased the para-deuterium concentration from 2022-11-02 to 2022-12-15. Taking both of these processes into account, the estimate of the ortho concentration on 2022-11-02 is (99.0 ± 0.3), with a rather large uncertainty. The para-deuterium concentration cannot explain why the normalised UCN yield on 2022-12-14 was lower than on 2022-11-02. It is possible that the volume of deuterium in the moderator on 2022-12-14 was so large, that the surface could not be cooled effectively, reducing UCN extraction due to thermal upscattering.

4.5 The deuterium filling level in the moderator vessel

Figure 4.2 displays a schematic cross-section of the moderator vessel. The shape of the moderator vessel is not perfectly cylindrical but is better approximated as a truncated cone (its cross-section is an isosceles trapezium). At the base, the inner radius equals $r_0 = 229$ mm. At a height of $h_1 = 160$ mm, the inner radius equals $r_1 = 243$ mm. In the central axis, there is a cylindrical insert that contains the helium in- and outlet, with a radius of $r_i = 22$ mm. When calculating the volume of deuterium in the moderator vessel as a function of height, this has to be subtracted. We neglect the rounded corners in the bottom of the vessel to simplify the estimation of the volume. The deuterium volume, as a function of filling height, is approximately

$$V(h) = \frac{\pi}{3} \frac{h_1}{r_1 - r_0} \left[\left(r_0 + \frac{r_1 - r_0}{h_1} h \right)^3 - r_0^3 \right] - \pi r_i^2 h. \tag{4.2}$$

At 5 K, the molar volume of deuterium equals $19.93\,\mathrm{cm}^3\,\mathrm{mol}^{-1}$ [56]. The molar mass of molecular deuterium equals $4.0282\,\mathrm{g}\,\mathrm{mol}^{-1}$, so the density at 5 K equals $0.2021\,\mathrm{g}\,\mathrm{cm}^{-3}$. Using the density of deuterium at 5 K, we estimated the mean height of the deuterium in the moderator (Table 4.1). The reported uncertainty on the mean filling level is systematic, it is mainly influenced by the uncertainty in density and the simplifications made in the geometry in the calculation of the volume. This means that the uncertainty on the differences of the filling levels is smaller than 2 mm. This calculation assumes that the deuterium surface is flat. The filling level should therefore be taken as the mean height of the deuterium above the bottom of the moderator vessel.

Note that when the moderator is filled with 5.68 kg of deuterium, the deuterium surface is higher than the coolant channels. This supports the hypothesis that a reduction in cooling effectiveness explains the lower normalised UCN yield with 5.68 kg of deuterium.
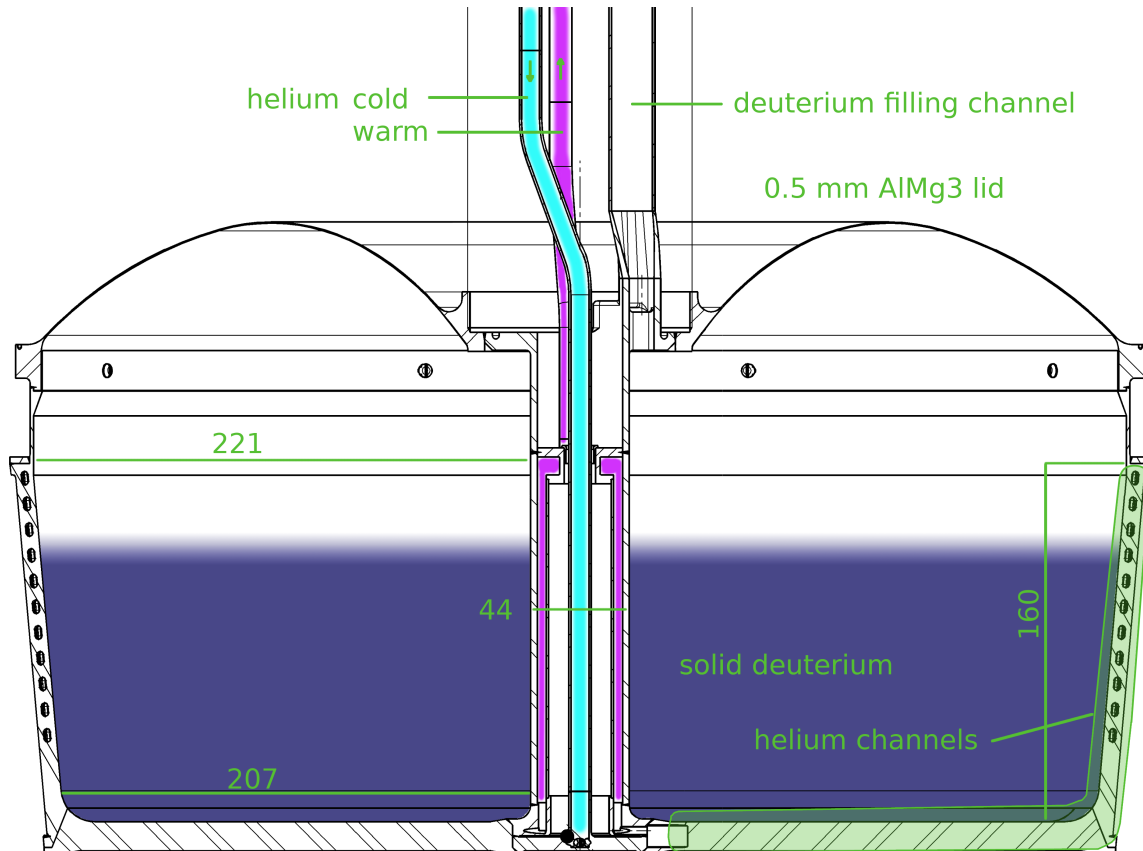
Figure 4.2: A schematic cross-section of the moderator vessel. The measurements in green were used for the estimation of the filling height, as a function of deuterium volume. The indigo shading represents the solid deuterium. The helium coolant channels reach a maximum height of 155 mm above the bottom surface of the vessel.

4.6 Inferring temperature from pressure in the moderator vessel

4.6.1 Pressure increase

During proton beam pulses, the pressure in the moderator vessel increases. This is due to the heating of the rest gas (helium) in the vessel, as well as the sublimation of deuterium from the surface of the moderator. The pressure change during the pulse has been used in previous work to determine the temperature of the deuterium surface [84].

For each of the three deuterium masses included in the analysis, we investigated the pressure rise in the moderator vessel during the most productive norm pulses (the last date for each mass in Figure 4.1). The average pressure rise during norm pulses is shown in Table 4.1.

The moderator pressure is monitored using the pressure gauge 2RNK10 CP006. This is a INFINICON SKY CDG045D capacitive gauge, with an accuracy of 0.15 % of the reading.[3] During these norm pulses, the residual moderator pressure due to helium in the deuterium system was around 3 to 4 Pa, giving the reported pressure readout accuracy of 0.15 %, corresponding to the vapour pressure of deuterium at 6.8 K [56]. Since the deuterium temperature is lower than 6.8 K when no proton beam pulses have been provided for more than 5 minutes [48], [84], we assume that the pressure at this point is entirely due to helium, and can be subtracted.

CP006 is at room temperature, separated from the moderator by a 10 m long, 22 mm wide tube. Due to the thermal transpiration effect, the pressure at the warm end of this tube is higher than at the cold end [56]. Therefore, CP006 overestimates the moderator pressure by factor 5 for pressures below 0.1 Pa, and by around 25 % at 3.5 Pa [109] (see also Figure 3.7 of [84]).

Taking into account the thermal transpiration effect, we estimated the maximum moderator temperature during norm pulses at 8.5 K on 2022-12-14 with 5.68 kg of deuterium, and 8.3 K on 2022-11-02 with 5.28 kg of deuterium (see Figure 3.12 in Chapter 3). A higher temperature was reached even though the beam current was 9 % smaller on 2022-12-14 (1.59 mA), compared to 2022-11-02 (1.75 mA), see Table 4.1.

The fact that the pressure increase was even higher on 2021-12-06 can be explained by the fact that the beam current was much higher on this date (1.94 mA).

We estimated loss rates using the equations described in section 2.2.1 of Chapter 2. The hydrogen deuteride and para concentrations were taken to be 0.21 % and 0.75 %, respectively. The estimated loss rate was averaged over the temperature ranges [5.3 K, 8.3 K] and [5.5 K, 8.5 K]. The estimated loss rates are $32.2\,\mathrm{s^{-1}}$ for $T_{max} = 8.3\,\mathrm{K}$ and $36.1\,\mathrm{s^{-1}}$ for $T_{max} = 8.3\,\mathrm{K}$. The estimated loss rate is about 5 % higher on 2022-12-14, compared to 2022-11-02. There is a large uncertainty on the difference in loss rate, due to the assumptions made in the estimation of the temperature, as well as the ortho concentration on 2022-11-02. Nevertheless, the difference in the estimated loss rate is of the same order of magnitude as the difference in normalised UCN yield, 3.3 %. This supports the hypothesis that a reduction in the effectiveness of the cooling power at the deuterium surface is responsible for the reduction in UCN yield on 2022-12-14, compared to 2022-11-02.

4.6.2 Base helium pressure

The helium pressure (the offset that was subtracted in Figure 4.3) was higher on 2022-12-14 (3.37 Pa) than on 2022-11-02 (2.71 Pa). If this was entirely caused by a temperature increase of helium at the same density, this temperature increase would have to be larger than 1 K. This is not the case. On days without proton beam pulses in October 2022 (for example 2022-10-20), the pressure reading in the moderator vessel was 2.65 Pa (measured using 2RNK10 CP006), and the pressure in the condenser was 3.27 Pa (measured using 2RNK10 CP008). Considering that the condenser was filled with about 2 L of deuterium and the moderator vessel with 25 L, the helium volume in the condenser was twice as large (42 L) as that in the moderator vessel.

---

[3]https://www.inficon.com/en/products/vacuum-gauge-and-controller/sky-cdg045d
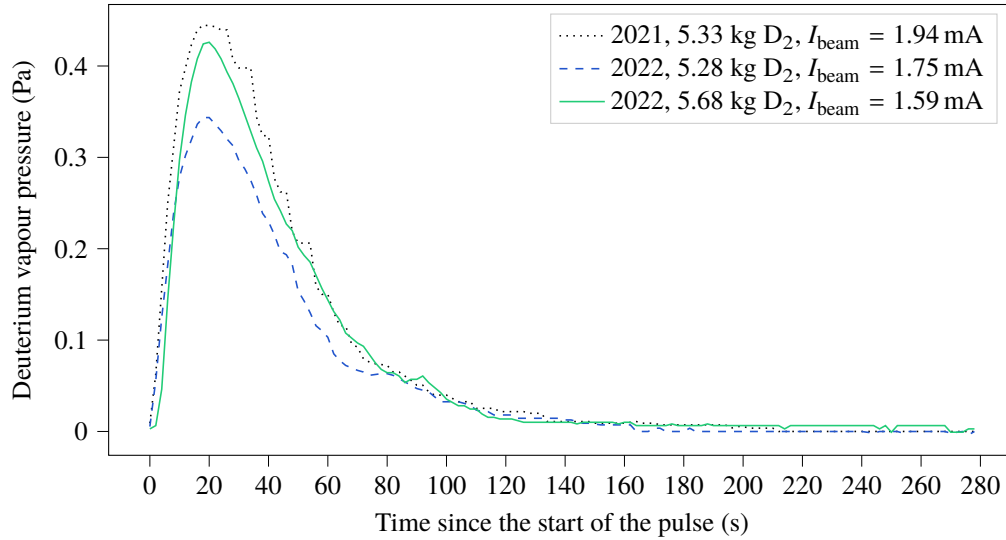
Figure 4.3: The pressure rise in the moderator vessel during a norm pulse. The pressure was measured using 2RNK10 CP006. One pulse was chosen from the day we measured the largest normalised UCN yield, for each of the three deuterium masses (see Figure 4.1). The offset pressure, caused by helium in the deuterium system, was different for every deuterium mass. This was subtracted to be able to compare the curves. The fact that the norm pulses in 2021 caused the largest pressure increase can be explained by the higher beam current in 2021.

After opening the valve, the pressure would equalise to 3.06 Pa. The estimated temperature increase is then

$$\frac{T_{5.68\,\text{kg}}}{T_{5.28\,\text{kg}}} = \frac{2.65\,\text{Pa}}{2.71\,\text{Pa}} \cdot \frac{3.37\,\text{Pa}}{3.07\,\text{Pa}} \approx 1.076. \tag{4.3}$$

4.7   Discussion

On 2022-11-02, we measured a record number of UCN during norm pulses of 6.49×10^6 UCN normalised to 2.2 mA of proton beam current. We also measured a record number of UCN during production pulses of 5.71×10^7 normalised to 2.2 mA of beam current. This was achieved with 5.277 kg of deuterium in the moderator, corresponding to a filling level of 151 mm, and an ortho-deuterium concentration around 99 %. This UCN yield was higher than the best yield obtained at the end of 2021, as well as, contrary to expectation, the yield obtained at the end of 2022, with a record amount of deuterium in the moderator (5.677 kg, filling level 162 mm). We investigated why the UCN yield was not obtained with the largest amount of deuterium in the moderator.

Para-deuterium concentrations cannot explain why the normalised UCN yield on 2022-12-14 was higher than on 2021-12-06, but 3.3 % lower than on 2022-11-02.

When the moderator is filled with 5.68 kg of deuterium, the deuterium extends above the maximum height reached by the helium coolant channels. Reduced cooling effectiveness at the deuterium surface could explain why the normalised UCN yield reached with 5.68 kg of deuterium was lower than the yield reached with 5.28 kg. This hypothesis is further substantiated by UCN loss rate estimated from the moderator pressure. This was 5 % higher during the pulses with 5.68 kg of deuterium, even though the beam current was 9 % smaller. The fact that the thermal conductivity of deuterium is lower at larger para-deuterium populations cannot explain the results, as the para-deuterium fraction was measured to be small on 2022-12-15 (0.57 %).

After 2022-12-15, HIPA operation was interrupted for the winter maintenance. Therefore, we do not

know if the UCN yield from 5.68 kg of deuterium would have kept increasing if we had kept applying proton beam pulses. The cumulative beam on target since the deuterium was frozen on 2022-12-14 was 3.26×10^3 standard pulses (1 standard pulse = $1.8\,\mathrm{mA} \cdot 8\,\mathrm{s}$). This is comparable with the value on 2022-08-22, (4.32×10^3), and much lower than the value on 2022-11-02 (5.51×10^3). It is well possible that further application of proton beam pulses would have reduced the number of lattice defects and decreased the para-deuterium concentration, leading to a higher output.

In 2027, the insert in the UCN source, containing the deuterium and helium pipes, the (damaged) flapper valve and the moderator vessel, will be replaced. We are investigating whether we can replace the aluminium lid, which is 0.5 mm AlMg3 thick, with a UCN transmission of approximately 50 % [53]. Based on the findings in this chapter, we recommend that the new moderator vessel be constructed with scooling channels which are several centimetres higher. This way, we can experiment with higher filling levels of the moderator vessel, while ensuring that all the deuterium is cooled effectively. We expect that we then can achieve even higher UCN yields.

We updated the procedure to prepare the UCN source moderator as a result of the study to initially solidify 5.3 kg of deuterium in the moderator vessel at the beginning of the UCN production period.

A COMPARISON OF THE UCN INTENSITIES
AT THE UCN SOURCE BEAMPORTS

5.1 Overview

To improve our understanding of the UCN source, we compared the ratio of the UCN intensities at all of its beamports, West-1, West-2 and South. The intensity ratio between beamports South and West-1 is of particular interest, since these are the beamports where the main experiments are mounted. The stability of the UCN intensities of the beamports over several days is relevant when we want to use a detector on one beamport to monitor a UCN experiment at another beamport, as was done in Chapter 6.

To be able to compare the UCN intensities at West-1 and South, we calibrated Cascade detectors Bob and Charlie, relative to each other. We found a large difference between the sensitivity of Cascade detectors Bob and Charlie and provide possible explanations for this difference.

We expected the same UCN intensities and time dependencies for beamports South and West-1. We found that the intensity of storable UCN at beamport South is $(81.6 \pm 0.2)\,\%$ of the intensity at beamport West-1. The fluctuation of this ratio over time is contained within the uncertainty.

The difference between West-1 and South is larger in the first seconds after the pulse. Using MCUCN [76] simulations, we determined that this is likely because one of the flaps in the flapper valve has to permanently stay closed during UCN source operation. This effect mostly affects neutrons with energies too high to be stored in the UCN source. These can only be counted if they reach the detector with few bounces, and are thus more sensitive to the geometry of the UCN source. Using MCUCN simulations, we found that the UCN intensity at beamport South for UCN storable in n2EDM is reduced by 28 % if one of the flaps remains closed, regardless of which flap.

5.2 UCN intensity measurements at the beamports

A CAD drawing of the UCN source is shown in Figure 1.2, in the introduction. We conducted several series of measurements of the UCN intensity at the beamports, summarised in Table 5.1.

We conducted a long series of measurements (series 1, 1297 proton beam pulses, of which 1029 led to valid measurements, due to some detector downtime) to measure the ratios of the UCN intensities at the beamports and their stability over time. Our three small Cascade UCN detectors, named Alice, Bob and Charlie, were used in the measurements. Alice was installed on beamport West-2. Bob was installed on beamport South using a NiMo-coated aluminium adapter. On one side, the diameter of this adapter matches

Table 5.1: Summary of the UCN intensity measurements at the UCN source beamports. They are numbered by the way they appear in the text.

Measurement series	Start date	Detectors			Number of pulses
		South	West 1	West 2	
1	2021-05-12	Bob + aluminium	Charlie + plexiglass	Alice	1297 (1029 valid)
2	2021-06-07	Bob + aluminium	Big Cascade	Alice	4
3	2021-06-08	Charlie + plexiglass	Big Cascade	Alice	8
4	2021-05-12	Bob + aluminium	Charlie + aluminium	Alice	4
5	2021-05-12	Bob + aluminium	Charlie + plexiglass	Alice	5

the flange installed on the BPV (outer: 140 mm, inner: 130 mm). On the other side, it matches the small Cascade detector flange (outer: 110 mm, inner: 100 mm). Charlie was installed on beamport West-1 using a flange on the BPV with a 110 mm wide opening and a NiMo-coated plexiglass UCN guide (outer: 110 mm, inner: 100 mm).

In series 2 and 3, we determined the UCN detection efficiency ratio between detectors Bob and Charlie. The big Cascade detector was attached directly to the BPV at beamport West-1, and one of the small Cascade detectors (Bob or Charlie) was mounted on beamport South. Bob was mounted with the aluminium adapter, and Charlie with the plexiglass guide. Only a small number (respectively four and eight) pulses were applied to ensure that there was no significant degradation in the UCN yield of the UCN source. This allows for a direct comparison of series 2 and 3.

In Series 4 and 5, the setup was the same as in series 1, except that Charlie was mounted on beamport West-1 with an aluminium adapter, like the one used to mount Bob on South. These two measurements were used to correct the efficiency ratio of the detectors for the adapters used to mount them on the beamports.

The Big Cascade detector always had a bias voltage of 1500 V, the small cascades always had a bias voltage of 1400 V. V_{ref} for all detectors was set to 110 V. The measurements are summarised in table 5.1.

5.3 UCN intensity ratios between the beamports

5.3.1 Overview

The UCN counts were divided into two intervals. The first interval includes UCN counted with $t_{\text{count}} \in [0\,\text{s}, 12\,\text{s})$, and the second includes UCN with $t_{\text{count}} \in [12\,\text{s}, 200\,\text{s})$, where t_{count} is the time difference between the start of the proton beam pulse, and the moment the neutron was counted in the detector.

The UCN intensities were normalised using the proton beam current, measured by HIPA beam monitor MHC1, to compensate for variations in the proton beam current. We defined a standard pulse as $1.8\,\text{mA} \cdot 8\,\text{s}$. This is natural for a unit of integrated beam current on target, since the HIPA proton beam current during the measurements was close to 1.8 mA. Recorded UCN intensities can then be expressed in UCN counts per standard pulse.

The UCN intensities, measured at all three beamports for each pulse, are shown in Figure 5.1. The degradation of the UCN yield of the UCN source over time, due to the frost build up on the solid deuterium in the moderator [100] is clearly visible in all beamports. This behaviour is described in more detail in Chapter 3.

The ratios between the UCN counted at beamports West-1 and South to West-2 are shown in Figure 5.2. There is a clear time dependence. The UCN with energies high enough to reach beamport West-2 are also more likely to pass through the frost layer on the solid deuterium, which is responsible for the degradation in UCN yield.

The ratios between the UCN intensities at beamports South and West-1 is shown in Figure 5.3. The ratio is very stable for the UCN detected between 12 and 200 seconds after the pulse, but for the UCN detected within the first 12 seconds, there is a clear slope.

5.3.2   Stability of the UCN intensity ratios between the beamports

While an experiment is performed at one of the beamports of the UCN source, the UCN source yield can be monitored using a UCN detector at another beamport. Since the UCN intensity ratios at the beamports change over many proton beam pulses, such monitoring is generally more accurate for experiments using fewer kicks. The accuracy of monitoring can be improved by correcting for drifts in the UCN intensity ratio between the experiment beamport and the monitoring beamport.

To quantify the stability of the UCN intensity ratios between the beamports over time we calculated relative standard deviations (RSD) of these ratios, defined as

$$\text{RSD} = \frac{\sigma_R}{\langle R_i \rangle}, \quad \text{with} \quad R_i = \frac{C_{1,i}}{C_{2,i}}, \tag{5.1}$$

where σ_R and $\langle R \rangle$ are the estimated standard deviation and mean of the set $\{R_i\}$, $C_{1,i}$ is the number of UCN counted at beamport 1, after the i^{th} pulse. If there is a constant difference in detector efficiency between beamports 1 and 2, this affects the magnitude of σ_R and $\langle R_i \rangle$ to equal degrees, and is divided out.

Changes in the UCN count ratios between pulses can be a result of slow drifts of the energy spectrum of the UCN extracted from the moderator and/or statistical noise. By calculating the RSDs of the beamport ratios for either small or large groups of pulses, we can determine the stability of the UCN count ratios in different scenarios.

Long term monitoring   A detector on one of the beamports is used to estimate the number of UCN entering an experiment, set up at one of the other beamports, over the course of days. The UCN count ratio is not recalibrated. The stability is represented by the relative standard deviation of UCN intensity ratio over all pulses in this series of measurements. In addition, we report the difference between the smallest and largest ratio as a percentage of the mean (see table 5.2).

Short term monitoring   An experiment is done with few pulses (up to 50), relying on the assumption that the UCN source does not significantly change in this time. To infer the stability in this case, the series of measurement is divided into groups of 50 consecutive pulses. For each group, the RSD of the UCN intensity ratio at the beamports is calculated. From this set of RSDs, the maximum value (worst case) is reported in Table 5.2.

Short term monitoring with slope correction   An experiment is conducted on one beamport while a UCN detector on another beamport monitors the UCN yield. The experiment is done with few pulses (up to 50). To be able to correct for a linear drift in the UCN intensity ratio at both beamports, the experiment is in the same configuration during the first and last few pulses.[1]

To estimate the stability of the UCN intensity ratio at the beamports in this scenario, the series of measurements is divided in groups of 50 consecutive pulses. In each group, a line is fitted through the ratios. In this case, the RSD is defined as

$$\text{RSD} = \frac{\sqrt{\langle \hat{\epsilon}_i^2 \rangle}}{\langle R_i \rangle}, \tag{5.2}$$

where $\{\hat{\epsilon}_i\}$ are the estimated residuals for the linear fit. For the set of relative standard deviations from all groups of 50 pulses, the maximum value (worst case) is reported in Table 5.2.

---

[1]This was the situation during the transmission measurements of the UCN guides for n2EDM, see Chapter 6.

(a)

(b)
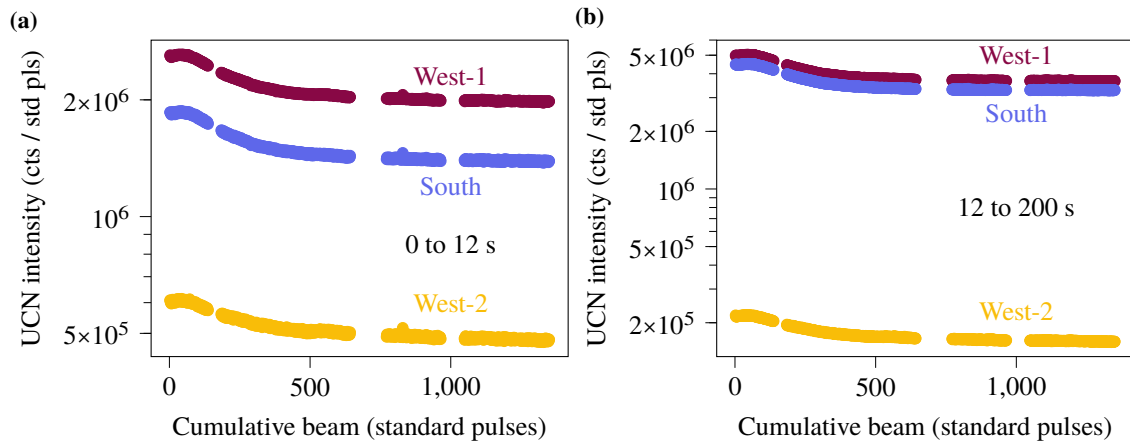
Figure 5.1: UCN count rates normalised to a standard proton beam pulse, integrated (a) between 0 and 12 seconds after the pulse, and (b) between 12 and 200 seconds after the pulse. Every marker represents the UCN counts in the specified time interval, in one detector, for a given pulse. Note the decrease in UCN yield after the pulse, which flattens off after some time.

(a)

(b)
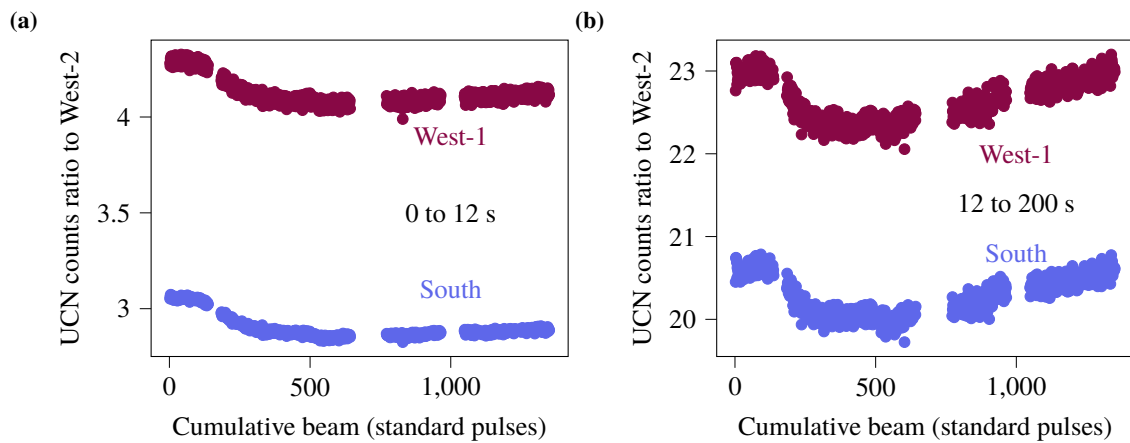
Figure 5.2: Ratio of UCN counted in beamports South and West-1, to West-2, detected (a) between 0 and 12 seconds after the pulse, and (b) between 12 and 200 seconds after the pulse.

(a)

(b)
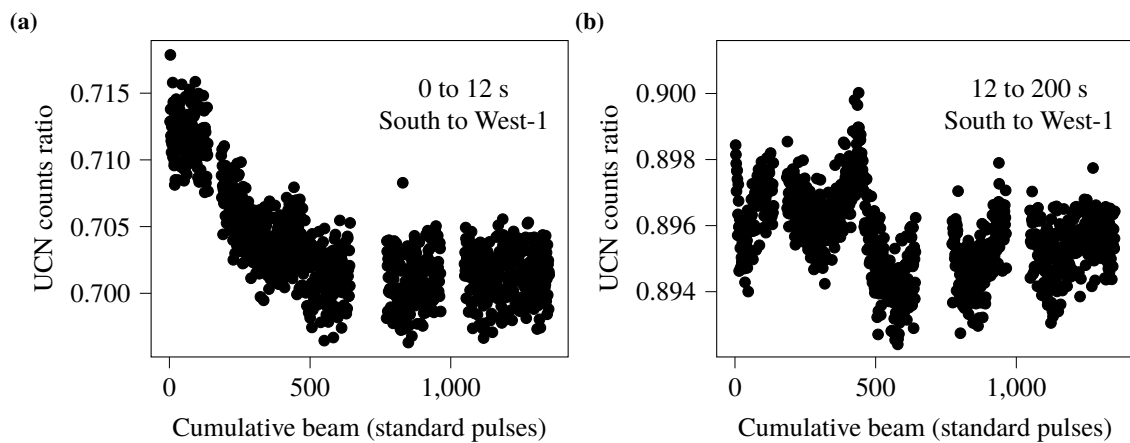
Figure 5.3: Ratio of UCN counted in beamport South to West-1, detected (a) between 0 and 12 seconds after the pulse, and (b) between 12 and 200 seconds after the pulse.

5.3.3 Results and interpretation

The UCN intensities over time are shown in Figure 5.1. The UCN count ratios between the different beamports are shown in Figure 5.2 and Figure 5.3. The calculated RSDs, for the scenarios described in the previous section, are displayed in Table 5.2.

The relative contribution of the higher energy neutrons is the highest in beamport West-2, because UCN with energies below 200 neV (at the level of beamports West-1 and South) can simply not reach it.

The variation of the ratio of UCN intensities between beamports is much larger in the interval 0 to 12 s than in the interval 12 to 200 s.

Neutrons with higher energies, and their distribution over South and West-1

In the interval 12 to 200 s, the UCN counts ratio for South and West-1 remains stable over time. However, significantly more neutrons are detected in West-1 than in South in the interval 0 to 12 s, and this difference grows over time (Figure 5.3). In the first 12 s after the start of the pulse, the energy spectrum is relatively wide, and neutrons up to $14\,\mathrm{m\,s^{-1}}$ (Figure 7 of [102]), 1×10^3 neV, reach the detector.[2] Neutrons with energies above 220 neV are unstorable, and do not significantly contribute to the UCN count rates in the interval 12 to 200 s. These neutrons are not UCN, but very cold neutrons (VCN). VCN are more likely than UCN to penetrate the frost layer that builds up in the moderator due to repeated pulsing [100]. In the interval 0 to 12 s, the number of UCN counted in the detectors decreases faster due to frost build up than the number of VCN. Because more neutrons are counted in West-1 than in South in the interval 0 to 12 s, and because this difference increases as frost builds up in the moderator, suggests that the VCN are more likely to reach the detector in West-1 than the one in South.

Several mechanisms could cause this energy screening. During the measurements, a broken NiMo-coated glass UCN guide in beamline South was replaced by a copper foil of over 50 cm in Length. Copper has a Fermi potential of 165 neV [30]. VCN would be disproportionally affected by this lower Fermi potential.

VCN have a higher loss-per-bounce rate than UCN. VCN are therefore more sensitive to the geometry of the UCN source: they can only reach the beamport if trajectories exist between the moderator and the detector with a relatively small number of collisions, at small glancing angles.

Inside the UCN source, a valve separates the UCN storage volume from the deuterium moderator. This valve consists of two flaps, named 'Böttstein' and 'Villigen', after the villages close to PSI in the direction (see Figure 5.4). Since 2019, the pneumatic system controlling the Böttstein valve leaks helium into the UCN source vacuum. To prevent pressure from building up, which could cause a temperature increase of the moderator due to the lack of insulation, the Böttstein flap is kept closed during UCN source operation [48]. This causes an asymmetry in the UCN source geometry. While the Villigen flap is open, it is positioned so that it could block some UCN headed towards the South beam guide, and reflect them towards the West-1 beam guide instead. This asymmetry disappears after both flaps are closed.

Monitoring uncertainty

The South:West-1 ratio of UCN counted from 12 to 200 s after the start of the pulse is stable (Figure 5.3b). The relative standard deviation over all the pulses in the first measurement series is 0.14 %, and the normalised maximum difference is 0.85 %. For short term experiments (up to 50 pulses), the maximum RSD equals

---

[2]Neutrons with energies larger than the Fermi potential of a surface are totally reflected if they graze the surface at an angle below the critical angle. Above the critical angle, the reflection probability scales with $|(p_1-p_2)/(p_1+p_2)|^2$, where k_1 is the wave number in vacuum $\sqrt{2mE}/\hbar$, and k_2 is the wave number in the material $\sqrt{2m(E-V_F)}/\hbar$. Even if a 1000 neV neutron hits NiMo perpendicularly to the surface, it has a 0.4 % probability to be reflected. This explains why some 1000 neV neutrons ($14\,\mathrm{m\,s^{-1}}$, Figure 7 of [102]) are detected, despite the very small probability that such neutrons are transmitted to the detector.

Table 5.2: Stability of the UCN intensity ratios at the beamports, expressed as relative standard deviations. The terms are explained in the text.

Beamports	Time interval	Long term		Short term	Slope corrected
		RSD (%)	Max diff (%)	Max RSD (%)	Max RSD (%)
W1 : W2	12 - 200 s	1.2	5.1	0.7	0.4
S : W2	12 - 200 s	1.2	5.2	0.8	0.4
S : W1	12 - 200 s	0.14	0.9	0.15	0.10
W1 : W2	0 - 12 s	1.7	8.2	1.0	0.6
S : W2	0 - 12 s	2.1	8.6	1.2	0.5
S : W1	0 - 12 s	0.57	3.1	0.4	0.3

0.15 %. These numbers can be interpreted as the relative uncertainty on a single UCN measurement at beamport West-1 using a monitor on beamport South (or vice versa).

The ratios of South:West-2 and West-1:West-2 are less stable than the ratio South:West-1. The maximum relative variation in these ratios over the whole series of measurements was 5.2 % (Table 5.2). If the UCN data are divided into groups of 50 pulses, the maximum RSD is 0.8 %. When correcting for the slope, the maximum RSD is halved (0.4 %). This number can be interpreted as the statistical error of a UCN measurement of a single pulse, monitored by the detector on beamport West-2, once the decrease of the UCN yield of the UCN source has been accounted for.

5.4 The relative efficiencies of small Cascade detectors Bob and Charlie

5.4.1 The relative efficiency including adapters

To make an absolute comparison of the UCN intensities at beamports South and West-1, we need to correct the ratio of UCN counts using the combined efficiencies of cascade detectors Bob and Charlie with their adapters.

To determine the relative UCN detection efficiencies, we mounted Bob and Charlie on beamport South in succession, respectively, in measurement series 2 and 3. Both detectors were mounted with the same adapters as in the long measurement (series 1). In both series of measurements, the big Cascade detector was mounted on beamport West-1 (see table 5.1). The comparison relies on the fact that the South:West-1 ratio is stable with a relative standard deviation of 0.15 % for UCN arriving from 12 to 200 s after the pulse.

The relative efficiency, $\tilde{\eta}_B / \tilde{\eta}_C$ is the ratio of the number of UCN detected at beamport South normalised by beamport West-1 between measurements 2 and 3:

$$\frac{\tilde{\eta}_B}{\tilde{\eta}_C} = \left\langle \frac{C_2^S}{C_2^{W1}} \right\rangle \left\langle \frac{C_3^S}{C_3^{W1}} \right\rangle^{-1}, \tag{5.3}$$

where $\tilde{\eta}_B$ stands for the efficiency of Bob together with the NiMo-coated aluminium adapter, $\tilde{\eta}_C$ is the efficiency of Charlie mounted with the NiMo-coated plexiglass adapter. C_2^{W1} represents the number of UCN counted in the detector on beamport West-1 in measurement 2 (from 12 to 200 s). The resulting efficiency ratio is

$$\frac{\tilde{\eta}_B}{\tilde{\eta}_C} = 1.097 \pm 0.002. \tag{5.4}$$

The uncertainty comes from the maximum RSD obtained for an experiment of 50 pulses (see Table 5.2). We used this efficiency ratio to correct the UCN intensities at the beamports. Since the efficiencies include the UCN transmission of the adapters, we need an extra correction to calculate the true detector efficiency.

5.4.2   Energy dependence of detector efficiency

The UCN time-of-arrival distributions in measurement series 2 and 3 were divided into 27 time intervals, chosen such that the amount of UCN detected in every interval was similar. For each time interval, the detector efficiency was calculated using Equation 5.3. The results are shown in Figure 5.5. The uncertainties were estimated by adding the relative Poisson errors from the UCN counts in quadrature.

In measurement series 1, 2 and 3, Charlie was mounted on the BPV with a plexiglass adapter, and Bob with an aluminium adapter. To calculate the true efficiency ratio of detectors Bob and Charlie, we need to correct for the UCN transmission of the adapters.

In Series 4 and 5 (Table 5.1), Charlie was mounted on BPV West-1 using either the aluminium or the plexiglass adapter. Using the same method to calculate the efficiency ratios of the detectors together with their adapters, we calculated the UCN transmission ratio of the aluminium and plexiglass adapters.

$$\frac{T_{\text{Pl}}}{T_{\text{Al}}} = \left\langle \frac{C_5^{\text{W1}}}{C_5^{\text{S}}} \right\rangle \left\langle \frac{C_4^{\text{W1}}}{C_4^{\text{S}}} \right\rangle^{-1} \tag{5.5}$$

Where $T_{\text{Pl,Al}}$ stands for the transmission of the plexiglass or aluminium adapter. The estimated transmission of the plexiglass adapter is (8.3 ± 0.2) % higher than that of the aluminium adapter for the interval 12 to 200 s and (7.2 ± 0.3) % higher for the interval 0 to 12 s.

The true efficiency ratio η_B/η_C is then

$$\frac{\eta_B}{\eta_C} = \frac{\tilde{\eta}_B}{\tilde{\eta}_C} \frac{T_{\text{Pl}}}{T_{\text{Al}}}. \tag{5.6}$$

The results are shown in Figure 5.5. For the interval 0 to 12 s, $\eta_B/\eta_C = 1.168 \pm 0.004$. For the interval 12 to 200 s, $\eta_B/\eta_C = 1.188 \pm 0.003$.

It is implausible that this time dependence of the efficiency ratio was caused by changes in the UCN yield or energy spectrum: measurement series 2 and 3 used few pulses, and we normalised the measurements using beamport West-1. This discrepancy needs further discussion.

5.4.3   Boron layer thickness and detector efficiency

The Cascade detectors are based on UCN capture in boron-10, which creates decay products which are counted in an ionisation chamber. A thin layer of boron-10 (approximately 200 nm [47], [110]) is coated on the back of the aluminium entrance foil of the detector (100 μm thick) using vapour deposition. The entrance foil closes off the ionisation chamber. The ionisation chamber is filled by a mixture of argon and carbon dioxide. Typically, a 1400 V bias voltage is applied. [111] Different thicknesses of entrance foils would cause different detection efficiencies. This is implausible, as all three small Cascade detectors were made with the same type of entrance foil. Differences in UCN detection efficiency could plausibly be caused by different thicknesses of the boron layer, as well as different charge collection efficiencies in the ionisation chamber.

The efficiency of UCN capture in Boron-10 in a cascade detector can be estimated using

$$\eta = (1 - A_{\text{Al}})A_{\text{B}}, \tag{5.7}$$

where A stands for the fraction of UCN that are absorbed in a layer. The absorption is calculated as

$$A = 1 - e^{-\Gamma_c \frac{d}{v}}, \tag{5.8}$$

where d/v is the duration the UCN with velocity v spend inside the layer with thickness d (v can be calculated using the UCN energy and the Fermi potential), and Γ_c is the capture rate of UCN in the material. The capture
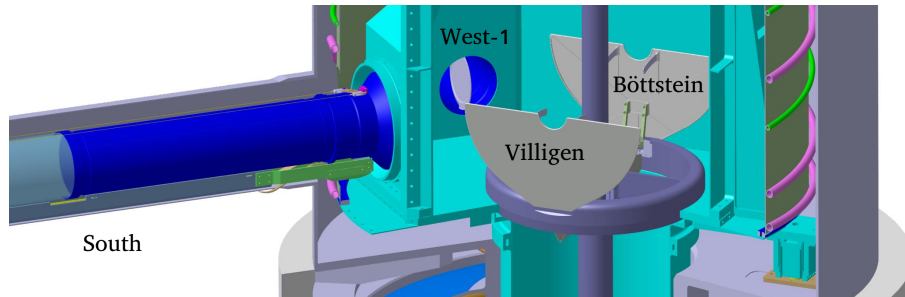
Figure 5.4: A CAD drawing of the position of the flapper valves and beam guide openings in the UCN source.
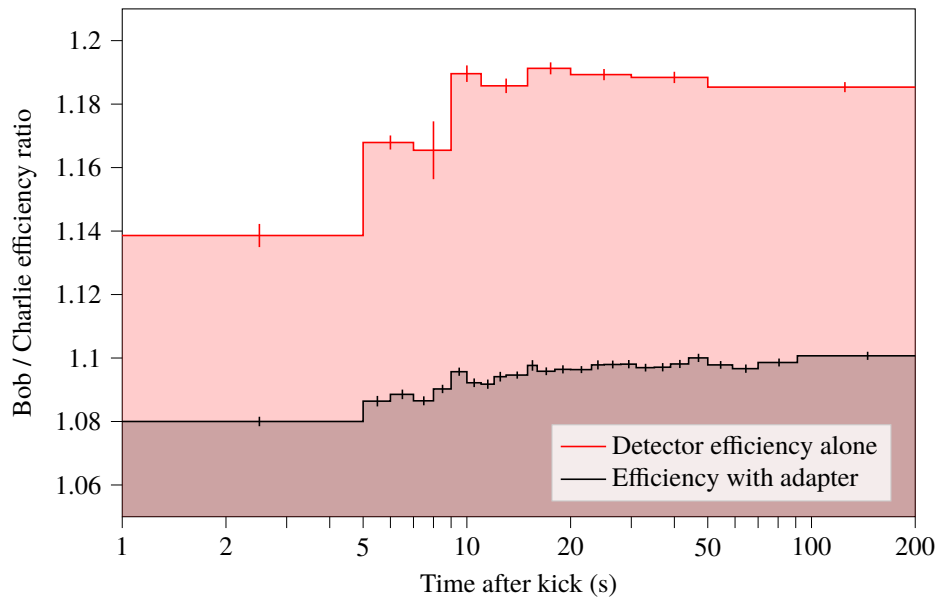
Figure 5.5: The efficiency ratio of Cascade detectors Bob and Charlie as a function of the time after the start of the pulse. The black curve is the efficiency ratio of both Cascade detectors, including the adapters used to attach them to the BPV during the long measurement (series 1). This quantity is relevant for calculating the intensity ratio at the beamports West-1 and South. The red curve is the true efficiency ratio of the Cascade detectors, corrected for the adapters. The efficiency ratio increases with time. The UCN spectrum softens as time passes after the kick. The plot suggests that the difference in efficiency for detectors B and C is larger for lower UCN energies. Possible causes are discussed in the text.
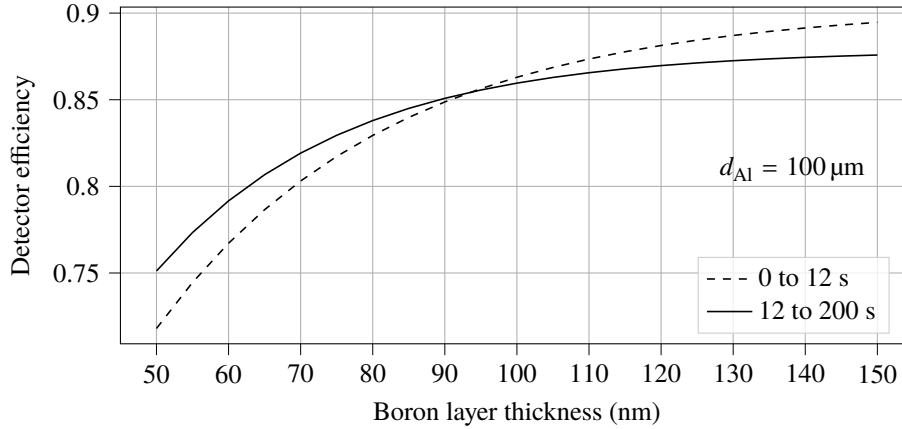
Figure 5.6: Estimated probability that a neutron impacting the Cascade detector entrance foil is absorbed in the boron-10 layer, as a function of the thickness of the boron-10 layer. The probability was integrated over the energy spectra of UCN detected in the intervals 0 to 12 s and 12 to 200 s. The energy spectra were estimated using MCUCN simulations. The entrance foil thickness is 100 μm.

rate can be computed from the capture cross-section as follows:

$$\Gamma_c = \sigma_c(v)nv, \tag{5.9}$$

where $\sigma_c(v)$ is the absorption cross-section at velocity v [85], and n is the number density of the material.

This efficiency can be averaged over the velocity distribution of UCN arriving at the beamport in MCUCN [76] simulations, in the interval 12 to 200 s.

The difference in detection efficiency in the interval 12 to 200 s could be explained if the boron layer is 30 nm thinner (see Figure 5.6).

5.4.4   Detector bias voltage and count rate

The UCN detection efficiency of the Cascade detectors depends on the charge collection efficiency of the ionisation chambers, and thus on the bias voltage.

On 25 October 2022, we mounted detector Charlie on beamline South and the Big Cascade detector on beamport West-1. We used norm pulses (2 s pulses, without operation of the flapper valve), and registered the UCN counts in both detectors. The bias voltage of one detector was varied. The bias voltage of the other detector was kept at the default value (1400 V for the small Cascade detectors, 1500 V for the big detector) to normalise the UCN counts. The UCN counts were normalised to 1 at the default bias voltage for each detector. The results are shown in Figure 5.7.

Even at the highest voltages, the background count rate was negligible ($< 0.1\,\mathrm{s}^{-1}$), indicating that no spontaneous breakdown occurred.

Based on Figure 5.7, the bias voltage on detector Charlie would have to be increased by hundreds of volts to make up for the 18.8 % difference in UCN detection efficiency between Bob and Charlie. It is possible that a difference in charge collection efficiency contributes to the difference in UCN detection efficiency.
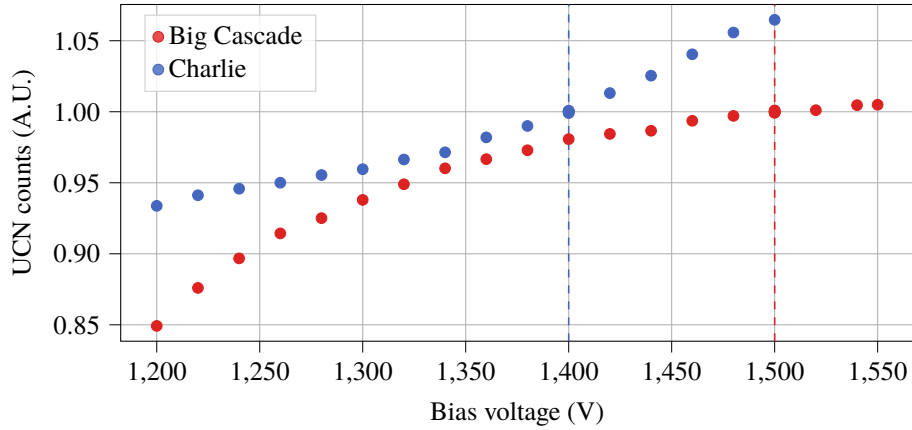
Figure 5.7: Normalised UCN counts in the Cascade detectors as a function of bias voltage. The dashed lines indicate the default voltage settings.

5.5 The UCN intensity ratio at beamports South and West

5.5.1 The intensity ratio for UCN arriving from 12 to 200 seconds after the pulse

Now that we have the efficiency ratio of the detectors and the adapters, we can calculate the ratio in UCN intensity at beamports South and West-1,

$$
\frac{I^S}{I^{W1}} = \left\langle \frac{C^S}{C^{W1}} \right\rangle \frac{\tilde{\eta}_C}{\tilde{\eta}_B} = 0.816 \pm 0.002, \tag{5.10}
$$

where I represents the neutron intensity, and C represents UCN counts. The mean is taken over all time-of-arrival distributions obtained in the first measurement series. The uncertainty comes from the RSD for a long-term experiment (Table 5.2) and the uncertainty on the efficiency ratio.

5.5.2 Time dependence of the UCN intensity ratio

We calculated the UCN intensity ratio for the same intervals as we used for the estimation of the detector efficiency ratios. For every time interval, we took the UCN intensity ratio between the beamports, averaged over all pulses in measurement series 1, and corrected this using the detector efficiency ratio corresponding to the time interval. The uncertainties were estimated as in the previous section. The result is shown in Figure 5.8.

In the first seconds after the start of the pulse, the UCN intensity is significantly higher at beamport West-1 than at beamport South. Several factors could contribute to this. Beamline South is 1 m longer than beamline West-1, which causes UCN to arrive later. During the pulse, the Böttstein flap remains closed while the Villigen flap is open (see Figure 5.4). This could reflect more neutrons towards West-1 than towards South.

Around 8 s after the start of the pulse (i.e. at the end of the pulse), the Villigen flap closes. This propels UCN towards West-1 [48], [102], causing a dip in the intensity ratio South/West-1. There is a small variation in the flapper valve timing.[3] This variation is likely responsible for the larger error bars around 8 s.

After the flapper valve has been closed, the ratio flattens off, peaks between 20 and 50 s, and then slightly decreases. The small decrease in the ratio after 50 s could also be caused by the fact that beamline South

---

[3]This is why the flapper valve is not moved during norm kicks.
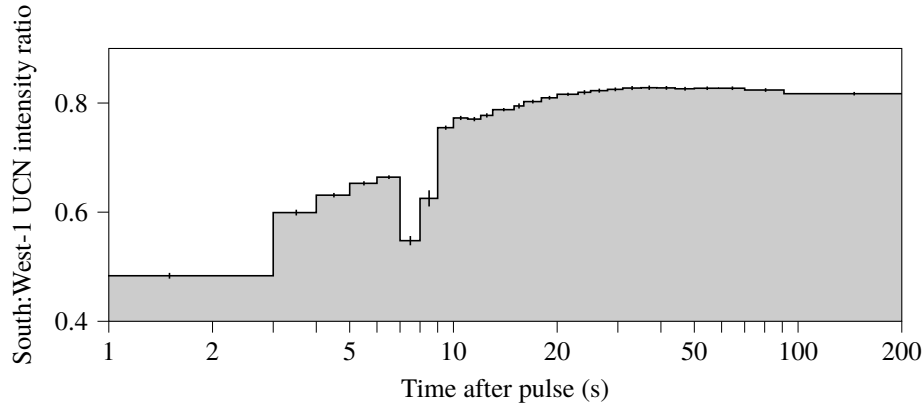
Figure 5.8: The UCN intensity ratio at beamports South and West-1 as a function of time after the start of the pulse. The ratio was corrected for detector efficiencies, unlike the ratio shown in Figure 5.3. Asymmetry of the flapper valve is the likely cause for the large discrepancy in the first seconds. The ratio changes momentarily around 8 s when the flapper valve of the UCN source closes. Differences in timing led to the large error bars around 8 s.
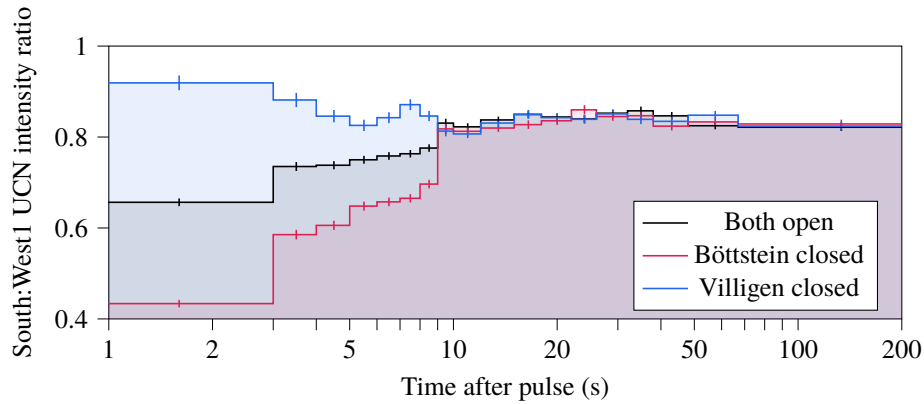
Figure 5.9: The UCN intensity ratio at beamports South and West-1, simulated using MCUCN [76]. Black: both flaps are open during the pulse, and close afterwards. Blue and Red: one flap opens and closes normally, while the other is permanently closed. During the pulse, the UCN intensity ratio between beamports South and West-1 heavily depends on which flap is opened. When we keep the Böttstein flap closed, the ratio closely corresponds with the measured ratio (Figure 5.8), except for the absence of the dip in the ratio around 8 s. The error bars represent Poisson errors.
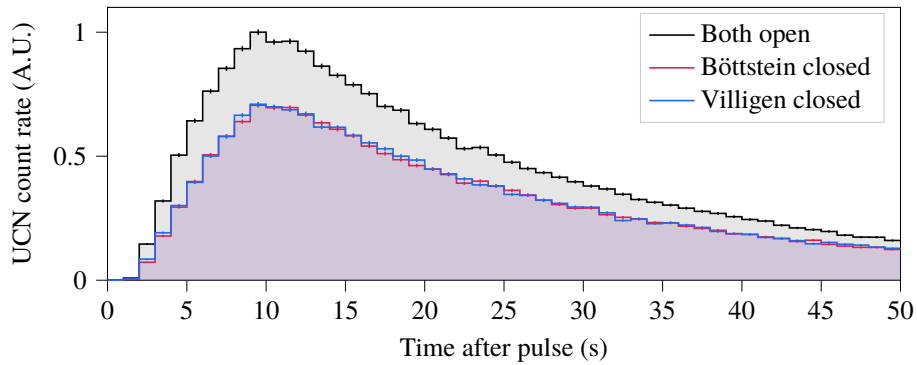
Figure 5.10: The simulated UCN count rate at beamport South, including only the UCN storable in n2EDM (205 to 389 neV, measured at the level of the solid deuterium in the moderator with 13 cm moderator filling).

is longer. The detectors at the beamports drain UCN from the system. Through a longer beamline, it takes longer for UCN from the source to replace the lost UCN.

5.5.3 Flapper valve operation in simulation

As mentioned before, the remarkable change in the UCN intensity ratio between West-1 and South, from 0 to 12 s, could be caused by the geometry of the UCN source. In particular, it could be caused by the fact that only the Villigen flap is operated. We tested this hypothesis using MCUCN simulations (see the introduction).

Simulated scenarios

The UCN source was simulated with detectors on all beamports. The Cascade detector entrance foils were simulated at the positions where normally the BPVs of beamports South and West-1 are.

The following scenarios were simulated:

- both flaps are open during the pulse, and close afterwards,

- the Villigen flap is open during the pulse, and closes afterwards, while the Böttstein flap remains closed (the current mode of operation),

- the Böttstein flap is open during the pulse, and closes afterwards, while the Villigen flap remains closed.

For all scenarios, 4×10^8 trajectories were simulated. The UCN intensity ratio South/West-1 was calculated as a function of time after the proton beam pulse. The results are shown in Figure 5.9.

Comparison with experiment

In the first 9 s after the start of the pulse, the UCN intensity ratio between beamports South and West-1 depends strongly on the state of the flaps. In the simulation in which the Böttstein flap remains closed, the UCN intensity ratio closely corresponds to the measured ratio.

The dip in the ratio around 8 s is absent. In MCUCN, the motion of the closing flapper valve is not simulated. Instead, the flaps jump between the open (87.5°) and closed (0.0°) positions via two intermediate positions (50.0° and 10°). Therefore, they cannot transfer momentum to the UCN. After the flaps are closed, there is no significant difference between the ratios in the scenarios. In all scenarios, the ratio peaks between 20 and 50 s, and slightly decreases afterwards, like in the measured results.

Even in the scenario where both flaps are open, the UCN intensity ratio South/West-1 is smaller in the first 8 s than afterwards. This can at least partially be explained by the fact that beamline South is 1 m longer than West-1, such that UCN have to survive more wall collisions to reach beamport South.

Relevance for n2EDM

We investigated whether the state of the flapper valve affects the UCN with energies storable in the n2EDM experiment, or only VCN. UCN need enough energy to reach the lower surface of the bottom precession chamber, 195 cm above the solid deuterium surface in the moderator (assuming 13 cm of deuterium filling). At the lowest point in the upper chamber (213 cm above the solid deuterium), the kinetic energy should not exceed the Fermi potential of the insulator ring coating (deuterated polystyrene, 165 neV [74]). UCN with energies between 205 neV and 389 neV are storable n2EDM. [4]

---

[4] The gravitational force acting on UCN at Earth's surface equals 105 neV m^{-1}.

Table 5.3: Reduction in the count rate of UCN storable in n2EDM when one of the flaps remains closed, in MCUCN simulations

Filling time (s)	Reduction in storable UCN intensity (%)	
	Villigen flap closed	Böttstein flap closed
5	40.4 ± 1.5	43.9 ± 1.6
10	33.6 ± 0.7	34.7 ± 0.8
30	28.3 ± 0.4	28.9 ± 0.4
50	27.5 ± 0.4	28.1 ± 0.4
70	27.3 ± 0.3	27.7 ± 0.4

The simulated count rate of UCN storable in n2EDM at beamport South, is shown in Figure 5.10. The n2EDM experiment is filled with UCN during a predefined filling time. In 2023, this was up to 50 s after the start of the pulse. Table 5.3 shows the reduction in the storable UCN intensity at beamport South, integrated over the filling time, if one of the flaps remains closed during the pulse. For a filling time up to 5 s, the reduction is slightly larger if the Böttstein flap remains closed than if the Villigen flap remains closed. For realistic filling times, the reduction is approximately 28 %, regardless of which flap stays closed.

We conclude that in the first seconds after the start of the pulse, slightly fewer storable UCN reach beamport South if the Böttstein flap is closed than if the Villigen flap is closed. This difference is much smaller than for VCN, and is irrelevant for realistic filling times for experiments like n2EDM.

5.6 Discussion

We have established that the ratio of UCN intensities at beamports South and West-1 between 12 and 200 s after the start of the pulse is very stable (relative standard deviation of 0.0015).

The height difference between beamline West-2 and the other beamlines acts as an energy filter. As the UCN source spectrum changes due to frost build-up on the deuterium moderator, there is significant drift in the UCN count ratios West-1:West-2, and South:West-2. However, if we average over up to 50 pulses and apply slope correction, we can use beamline West-2 as a monitor for South and West-1 with an error per pulse of about 0.4 %.

Beamport South has a lower UCN intensity than beamport West-1. Taking into account the UCN arriving between 12 and 200 s after the start of the pulse, the ratio of the UCN intensity at beamport South to beamport West-1 is 0.816 ± 0.002. This could be caused by the fact that beamline South is about 1 meter longer than beamline West-1. In the interval 0 to 12 s, the difference is significantly larger. Based on MCUCN simulations, this can in part be explained by the fact that the Böttstein flap of the UCN source does not open.

During the measurements, the UCN guide in the superconducting magnet was broken and temporarily replaced by a copper foil. Copper has a Fermi potential of 165 neV [56]. This could partially explain the lower UCN intensity at beamport South, compared to beamport West-1.

Small Cascade detector Bob is (18.8 ± 0.3) % more efficient than detector Charlie, for the UCN in the time interval 12 to 200 s. This number is corrected for the adapters used for Bob and Charlie. Charlie could have a thinner Boron layer than Bob, or a less efficient charge collection, or both.

PART II

THE N2EDM EXPERIMENT AT PSI

CHAPTER 6

UCN TRANSMISSION OF THE UCN GUIDES

6.1 Introduction

The statistical uncertainty of the n2EDM experiment is inversely proportional to the square root of the number of UCN that reach the detectors. The UCN guides from the UCN source to the precession chambers and from the precession chambers to the detectors need to be efficient in transporting UCN with low losses.

We measured the UCN transmission of the guides that we built into the experiment. The best UCN guides have a UCN transmission of over 95 %, normalised to a guide length of 1 m. However, there is a large variation in UCN transmission between individual guides. Some guides have a transmission significantly lower than what we would expect based on UCN transport simulations, using surface parameters that have been fitted to experimental values.

We used measurements and simulations to determine the cause of the variation in UCN transmission between the guides and concluded that the most likely explanation is that rough sections in the guides cause non-specular scattering of the UCN. In particular, this could be due to grooves at the end of some guides. These were caused during the glass-blowing process: the ends of the guides were rolled until the outside diameter matched the specification (140 mm) with a small tolerance. The recommended course of action is to replace at least the horizontal UCN guides for n2EDM to see if this improves the number of UCN in the chamber.

6.2 Methods

6.2.1 Measurement setup

We used two setups in the transmission measurements: the reference setup and the test setup (Figure 6.1). The reference setup consisted of a reference guide connected to beamport valve (BPV) West-1, followed by the big Cascade detector. In the test setup, the UCN guide under test is inserted between the reference guide and the Cascade detector.

We recorded the UCN count rate in the reference setup, for several proton beam pulses. Afterwards, we installed the test setup and again recorded the count rate for several pulses. The transmission is defined as the ratio of the UCN count rate in the test setup to the count rate in the reference setup. Both setups are shown in Figure 6.1.

The properties of the reference guides that we used were similar to those of the UCN guides for n2EDM: they are borosilicate glass guides, with an inner diameter of 130 mm, an outer diameter of 140 mm, and a

97
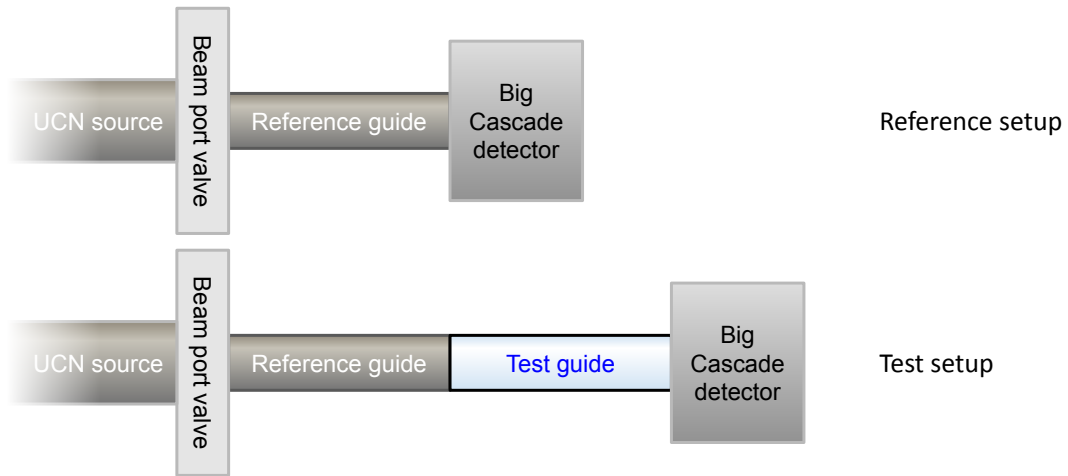
Figure 6.1: The two basic setups used in the transmission experiment. The setups were installed on beamline West-1.
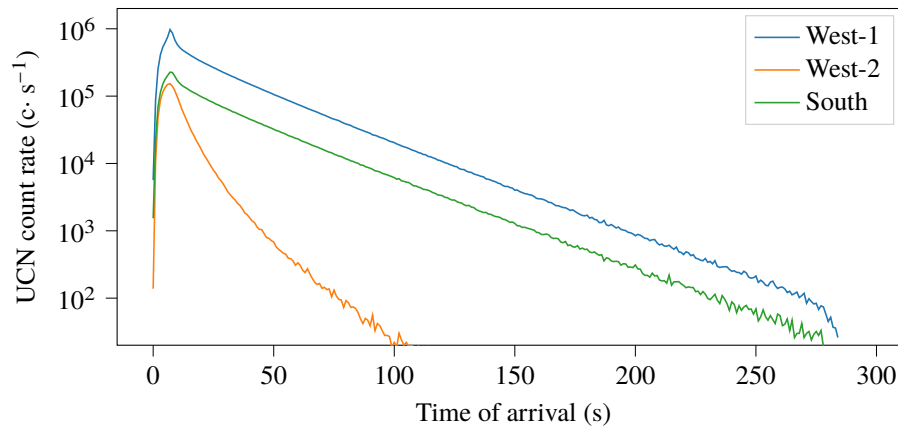
Figure 6.2: The measured UCN time-of-arrival distribution in the reference setup at all UCN source beamports. The time of arrival of the UCN is counted since the start of the proton beam pulse. The UCN fluxes in beamports South and West-1 have an almost equivalent time dependence up to a constant factor of 0.81 (see Chapter 5). The data visualised in this plot were collected using the big Cascade detector on beamport West-1 and a small Cascade detector on beamport South. The difference in detector surface area amplifies the difference in UCN count rates between beamports South and West-1.

wall thickness of 5 mm, coated on the inside with nickel-molybdenum (85:15 mass ratio). The connection between the test guide and the reference guide was made using a Wilson flange, so that the connection is glass on glass, without additional material in between that could affect the transmission. Throughout the experiments, two different reference guides have been used, with lengths of 300 and 1000 mm.

The reference guide serves two purposes. It eliminates the need for an adapter between the BPV and the cascade detector: the only difference between the two setups is the test guide. It also collimates the UCN from the 180 mm diameter of the UCN source to the diameter of the test guide. Collimation effects that occur in the reference guide therefore cannot affect the measured transmission of the test guide.

6.2.2 UCN rate monitoring

We can account for short-term fluctuations and long-term drifts in UCN production by using a beam monitor. During the transmission measurements in area West-1, small Cascade detector A was installed on beamport West-2, and small Cascade detector B on beamport South.

In Chapter 5, we found that over a thousand proton beam pulses, the West-1/South count ratio, for UCN arriving from 12 to 200 seconds after the start of the pulse, fluctuated by 0.14 % (standard deviation / mean), with a maximum spread of 0.8 % ((max − min)/μ). Beamport South is the preferred monitor for beamport West-1.

Cascade detector Bob on beamport South was down during some transmission measurements, requiring us to use detector Alice on West-2 as for monitoring instead. The West-2 beam guide is attached to the UCN Source storage vessel, two meters higher than the beam guides South and West-1[37]. This height difference acts as an energy filter. Only UCN with energies above 330 neV can reach the beam guide. These UCN have a short storage time in the UCN source. As a result, the time-of-arrival distribution of West-2 is characterised by a much faster decay of the UCN count rate (Figure 6.2). Over a thousand pulses, the ratio West-1/West-2 fluctuated with a standard deviation of 1.1 % (standard deviation / mean), with a maximum spread of 5.1 % ((max - min)/mean) (see Chapter 5).

6.2.3 Analysis

Data preparation

The integrated proton beam current on target during each pulse is determined from HIPA beam monitor MCH1, in order to correct for short-term UCN yield fluctuations. The cumulative integrated proton beam current up to each pulse is calculated by summing over the integrated proton beam currents. This quantity is required to fit linear models used for UCN yield drift correction. The cumulative integrated proton beam current is calculated for each pulse, regardless of whether UCN data were being recorded.

The UCN counts in the detectors (West-1, West-2 and South) are summed over intervals, separated by the following timestamps after the start of the pulse, in seconds: 0, 12, 30, 60, 100 and 200. These intervals were analysed separately. The neutrons arriving from 0 to 12 seconds after the start of the proton beam pulse are treated separately from the rest of the neutrons because we measure a lot of unstorable neutrons in this time period. These can be reflected under small grazing angles, and end up in the detector in a direct transmission experiment, but they are lost quickly after the end of the pulse. The other time bins were chosen arbitrarily, such that there is a reasonable number of UCN in every interval. Typically, we are interested in the interval containing all neutrons that arrived between 12 and 200 seconds after the proton beam pulse. The other subdivisions were mainly used for diagnostics.

Outliers with extremely high or low neutron counts can be caused by a multitude of factors. For example, due to missed proton beam pulses due to an interlock, improper configuration of the Cascade detectors, closed NLKs or BPVs, or a change in the operation of the flapper valve. After visual inspection of the data,

an inclusion range for the total number of UCN counted during and after one proton beam pulse was specified for every detector. Measurements outside these ranges were discarded. These ranges were $(1 \times 10^6, 3 \times 10^8)$ for West-1, $(1 \times 10^5, 5 \times 10^8)$ for West-2 and $(1 \times 10^5, 3 \times 10^8)$ for South.

Several measurements within these ranges seemed to be outliers. In these cases, the measurements were investigated individually, and they were discarded if they looked anomalous (e.g. if the measurement was interrupted, the measurement was not synchronised with the proton beam pulse, or the decay constant was very unusual). For most of these outliers, the cause could be identified.

Transmission fit

When South is used to monitor West-1, the transmission can be simply calculated as follows:

$$T = \frac{\langle C_{\text{test}} \rangle}{\langle C_{\text{ref}} \rangle} \cdot \frac{\langle M_{\text{ref}} \rangle}{\langle M_{\text{test}} \rangle}, \tag{6.1}$$

where T is the calculated transmission, C are the UCN counts in the test and reference setup, and M represents the counts in the monitors. A transmission value can be assigned to each time bin.

If the ratio of the monitor to the UCN counts in the reference setup is correlated to the state of the UCN source, and drifts over time, we can estimate the transmission by fitting a model that allows for linear drift:

$$\frac{C_i}{M_i} = f(Q_i; a, b, \{T_g\}) + \epsilon_i. \tag{6.2}$$

Here C_i, the UCN counts for proton beam pulse i, divided by the monitor, M_i, are predicted by a fit function f. The independent variables are integrated beam current, Q_i. The fit parameters are a scale factor a, a slope b, and the set of transmission values for the relevant guides, $\{T_g\}$.

The fit function has the form

$$f(Q_i; a, b, \{T_g\}) = \begin{cases} (aQ_i + b)T_g & \text{for test guide } g \\ (aQ_i + b) & \text{in the reference setup} \end{cases}. \tag{6.3}$$

When we apply the fit to sets of fewer than 50 measurements, nonlinear components of the drift in the UCN intensity ratio between West-1 and West-2 are not large enough to be relevant. An example of the fit is shown in Figure 6.3.

6.2.4   Comparison with the prestorage method

The UCN transport in a UCN guide has a complex dependence on its surface properties. There exist different definitions of UCN transmission, with corresponding measurement methods sensitive differently to the surface properties of UCN guides.

During the construction of the UCN source, the UCN transmission of the UCN guides in the beamlines was measured at the Institut Laue-Langevin (ILL) in Grenoble, France, using a prestorage method [112]. The setup was similar to the setup in Figure 6.1, except that instead of a reference guide, a storage guide was installed, with beamline shutters on both sides. This guide was filled with UCN, keeping the detector-facing shutter closed. Then the source-facing shutter was closed, and after a storage duration of five seconds, the detector-facing shutter was opened, draining the UCN through the test guide (if it was installed) into the detector. Since the storage volume was of a similar size as the test guide, UCN reflected back to the storage volume could return to the test guide after a short time.

In the direct transmission measurements described in this chapter, back-reflected UCN can travel all the way back to the UCN storage vessel of the UCN source, and may never reach the test guide again. This
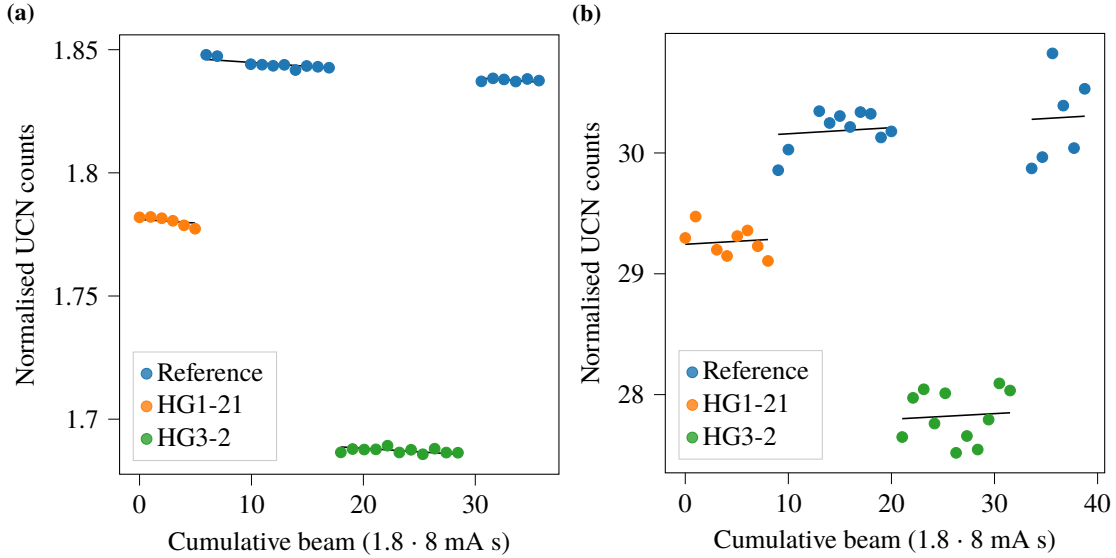
Figure 6.3: Two transmission fits, including drift correction, to UCN data obtained during the transmission measurements of UCN guides HG1-21 and HG3-2 (Table 6.1). Both plots show the normalised UCN counts arriving from 12 to 200 s after the proton beam pulse as a function of the cumulative beam on target. Every dot represents one pulse, the black lines represent the fit. In (a), the detector on beamport South (12 to 200 s after the pulse) was used as a monitor, whereas in (b), the detector on West-2 was used (12 to 200 s). South is more reliable as a monitor, but because of missing data, we sometimes had to use West-2.

method is thus more sensitive to the surface roughness of UCN guides than the prestorage method. For similar UCN guides, we expect lower transmission values than the ones found using the prestorage method (0.959 to 0.999, normalised to a guide length of 1 m [112]).

6.2.5   UCN transmission simulations using MCUCN

We used neutron transport simulations to better understand the measurement results. We simulated the entire UCN system, from the solid deuterium moderator in the UCN source to the boron-10 layer on the foil of the Cascade detector. The length (l), the Fermi potential (V_F), the diffusive scattering fraction (d) and the loss coefficient (η) of the test guide were the physical input parameters of the simulation, the rest of the simulation was unchanged. Transmission is defined in the same way as in the experiments: for each respective time interval, the transmission is the ratio of the count rate in the test setup over the count rate in the reference setup. The reference setup was simulated by setting l to 1 nm. For most parameter sets, we generated 100 to 150 million UCN trajectories starting from the solid deuterium surface.

Since the simulation of all setups is compared to the reference setup, we simulated more trajectories for the reference setup (1000 million) to reduce its contribution to the statistical uncertainty. In the reference setup, (709.6 ± 1.3) ppm of the UCN generated at the solid deuterium surface ended up in the detector on West-1. In the simulations, the state of the UCN source does not change, and the number of simulated UCN is exactly known. Therefore, the monitor was not needed for normalisation and not simulated. NLK South was kept closed in the simulations.
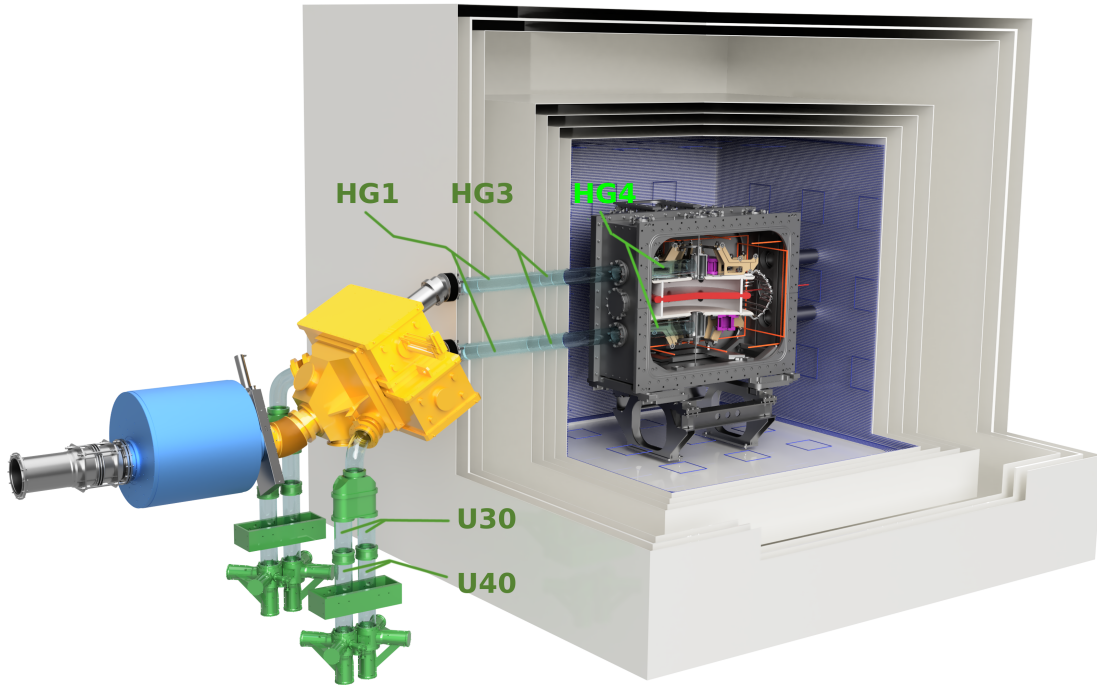
Figure 6.4: The position of some UCN guides in n2EDM.

6.3   Transmission measurements of UCN guides for n2EDM

Using the methods described above, we measured the transmission of the major UCN guides built into
n2EDM (see Figure 6.4 for an overview of their positions in n2EDM). An overview of the transmission
measurements of UCN guides is shown in Table 6.1. The measurements of several UCN guides were re-
peated to investigate the reproducibility of the results. At some point, we switched from using a 300 mm
NiMo-coated reference guide to a 1000 mm NiMo-coated reference guide. We measured the transmission
of several guides with both reference guides to investigate whether the choice of reference guide affected the
measured transmission value and concluded that this was not the case.

6.3.1   Comparison of the monitoring methods

The transmission values for the UCN guides can be estimated using different monitoring methods. The
statistic used to normalise the UCN counts in area West-1, were either the UCN counts in South (12 to 200
s), West-2 (12 to 200 s), West-2 (30 to 200 s) and the integrated proton beam, measured by the HIPA beam
monitor, MHC1. From 12 seconds after the proton beam pulse, the UCN flux ratio between beamports South
and West-1 is stable, with a relative standard deviation of 0.15 % (see Chapter 5). Beamport South will be
used to normalise the UCN counts in West-1 whenever possible. To determine the next best monitor, for
the measurements where South was not available, for each other the root mean squared difference of the
estimated transmission values to the transmission values estimated using South (12 to 200 s) was calculated.
These values were: 3.6 % for the proton beam monitor (MHC1), 0.63 % for West-2 from 12 to 200 s, 0.50 %
for West-2 from 30 to 200 s. It follows that the counts in West-2, from 30 to 200 s, are the best monitor for
the measurements when South failed.

Table 6.1: Summary of the transmission measurements of UCN guides.

Name	L (mm)	Description	Date	Reference guide
U29.8-1	298	USSA guide	2022-06-20	300 mm NiMo
			2022-06-23	300 mm NiMo
U30-1	300	USSA guide	2022-06-21	300 mm NiMo
U40-1	400	USSA guide	2022-06-20	300 mm NiMo
U40-2	400	USSA guide	2022-06-21	300 mm NiMo
SB45-1	450	Switch guide	2022-06-22	300 mm NiMo
HG4-1	595	Horizontal guide: switch → precession chamber	2022-06-22	300 mm NiMo
HG4-2	595	Horizontal guide: switch → precession chamber	2022-06-21	300 mm NiMo
HG4-3	595	Horizontal guide: switch → precession chamber	2022-10-10	1000 mm NiMo
			2022-10-11	1250 mm uncoated
HG1-21	625	Horizontal guide: switch → precession chamber	2022-06-10	300 mm NiMo
HG1-21	625	Horizontal guide: switch → precession chamber	2022-06-17	300 mm NiMo
HG1-3	625	Horizontal guide: switch → precession chamber	2022-06-24	300 mm NiMo
			2022-09-29	1000 mm NiMo
			2022-10-31	1000 mm NiMo
			2022-10-14	1250 mm uncoated
HG1Ti	625	Horizontal guide: switch → precession chamber, coated with titanium and NiMo	2022-09-27	300 mm NiMo
			2022-10-25	1250 mm uncoated
HG3	1250	Horizontal guide: switch → precession chamber	2022-06-08	300 mm NiMo
HG3-2	1250	Horizontal guide: switch → precession chamber	2022-06-17	300 mm NiMo
HG3-u	1250	Horizontal guide: switch → precession chamber, uncoated	2022-06-22	300 mm NiMo
			2022-10-12	1000 mm NiMo
			2022-10-24	1000 mm NiMo
HG4-1 U29.8-1	893	Combination of guides	2022-06-20	300 mm NiMo
HG1-3 HG1-Ti	1250	Combination of guides	2022-09-28	300 mm NiMo
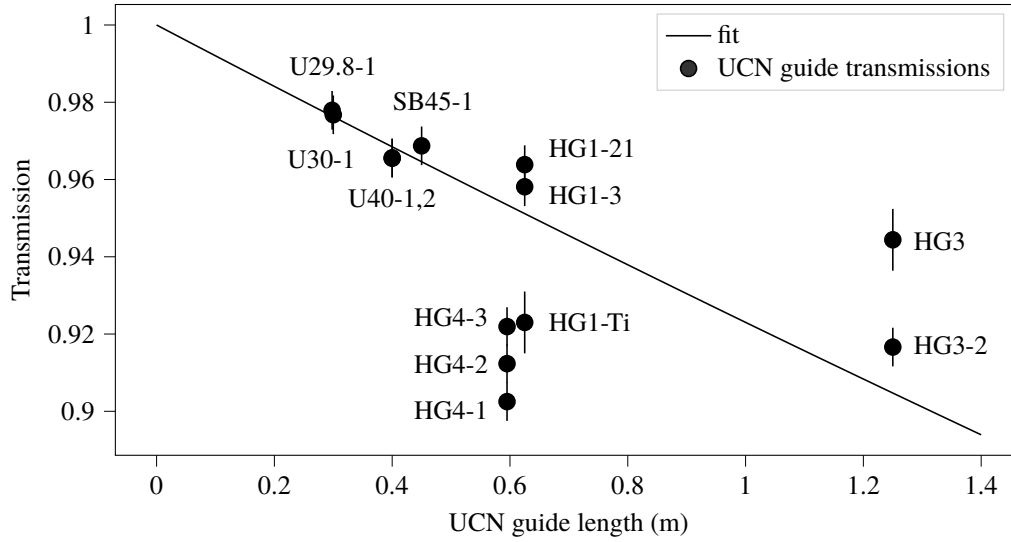
Figure 6.5: Measured transmission values as a function of UCN guide length. The fit is a simple exponential fit. The shown uncertainties are derived from statistics only. The estimated uncertainty including systematics is 0.007. A description of the location of the guides in the experiment is given in Table 6.1.

6.3.2   Transmission values

For the measurements where South was available as a monitor, South (12 to 200 s) was used to estimate the UCN guide transmissions. Otherwise, West-2 (30 to 200 s) was substituted as a monitor. The uncertainty on the transmission values was estimated from the fit. (The UCN Poisson errors were much smaller than the uncertainties from the fit.)

In the interval from 12 to 200 s, the largest absolute difference in the transmission values obtained for the same guide is 0.010, typically it is smaller. Note that South was not available as a monitor for some of these measurements. From this, we estimate the uncertainty on the transmission values as 0.005, or 0.008 when West-2 was used as a monitor.

The results are displayed in Table 6.2 and Figure 6.5.

For a UCN guide with a given diameter and surface properties positioned in the horizontal plane, the probability of losing a UCN with a given momentum in any segment of the guide should be equal. Then, the transmission should decay exponentially with the length of the guide:

$$T = T_0^{\frac{l}{l_0}}, \tag{6.4}$$

where T and l are the transmission and length of the UCN guide, and T_0 is the transmission at a standard length l_0. T_0 for $l_0 = 1$ m will henceforth be referred to as specific transmission. Specific transmission is a function of the UCN energy, the UCN guide diameter, and the surface properties of the guide.

The UCN guides were manufactured using the same process, therefore we assumed the specific transmissions to be comparable, which is not the case. Differences in measured transmission values could be either caused by irreproducibility of the UCN guide manufacturing process or by irreproducibility of the measurement setup. The latter could be caused by mounting errors during the installation of the measurement setup. These mounting errors could add a random offset to the transmission, every time a new guide was mounted. We ruled out the irreproducibility of the measurement setup as an explanation for the different specific transmissions of the UCN guides, as explained in the next section.

Table 6.2: Measured UCN guide transmissions (12 to 200 s)

Name	Length (mm)	Transmission[a]	Specific transmission (transmission normalised to 1 m)
HG1-21	625	0.964	0.943 ± 0.011
HG1-3	625	0.958	0.934 ± 0.011
HG1Ti	625	0.927	0.885 ± 0.011
HG3	1250	0.944	0.955 ± 0.006
HG3-2	1250	0.917	0.933 ± 0.006
HG4-1	595	0.903	0.842 ± 0.011
HG4-2	595	0.912	0.857 ± 0.011
HG4-3	595	0.922	0.872 ± 0.011
SB45-1	450	0.969	0.932 ± 0.015
U29.8-1	298	0.978	0.928 ± 0.022
U30-1	300	0.977	0.925 ± 0.022
U40-1	400	0.966	0.916 ± 0.017
U40-2	400	0.966	0.916 ± 0.017

[a] The uncertainty on the transmission values equals 0.007, determined in Section 6.4.3

6.4 The effect of UCN guide mounting errors on transmission values

We identified two possible mechanisms by which mounting errors could reduce the measured UCN transmission, represented in Figure 6.6.

If a test guide is attached to the measurement guide under an angle, the guides touch only at one point, leaving a gap between them. If there is an offset between the central axes of the guides, the overlap between the guides' interior cross-sections is imperfect. This would have a similar effect as a local reduction of the inner diameter of the guide, creating a bottleneck.

We conducted measurements to simulate the effect of gaps in the setup, or local reductions in diameter. We installed the transmission measurement test setup and measured the transmission. Metal rings were installed to simulate successively larger gaps or diameter reductions.

6.4.1 Gaps between the UCN guides

To investigate the effect of gaps between the test and reference guides on the transmission measurement setup, we modified the standard transmission setup (Figure 6.1) by inserting aluminium rings between the reference guide and the test guide (Figure 6.7a). We used the 300 mm NiMo-coated reference guide, and U298-1 (a 298 mm NiMo-coated guide for n2EDM) as a test guide.

The aluminium rings simulated gaps in the setup. Their outer diameter was 140 mm (like the UCN guides) and their inner diameter was 135 mm (5 mm larger than the inner diameter of the UCN guides). Aluminium has a low Fermi potential of 54 neV [30]. On collision with the aluminium ring, a significant fraction of the UCN will enter the aluminium and be captured. To be able to estimate the loss of transmission as a function of gap surface area, we used rings with various thicknesses: 2.0, 5.0 and 10.0 mm. We also measured with all rings combined, with a total thickness of 17.0 mm.

The UCN count rates were normalised using beamport West-2.

We found a linear relationship between the ring thickness and the transmission. The loss in transmission per unit thickness was $0.0062\,\mathrm{mm}^{-1}$, for neutrons counted in the interval 12 and 200 s. Neutron sinks, such as gaps or aluminium surfaces, remove UCN from the system over time. Therefore, the reduction in count
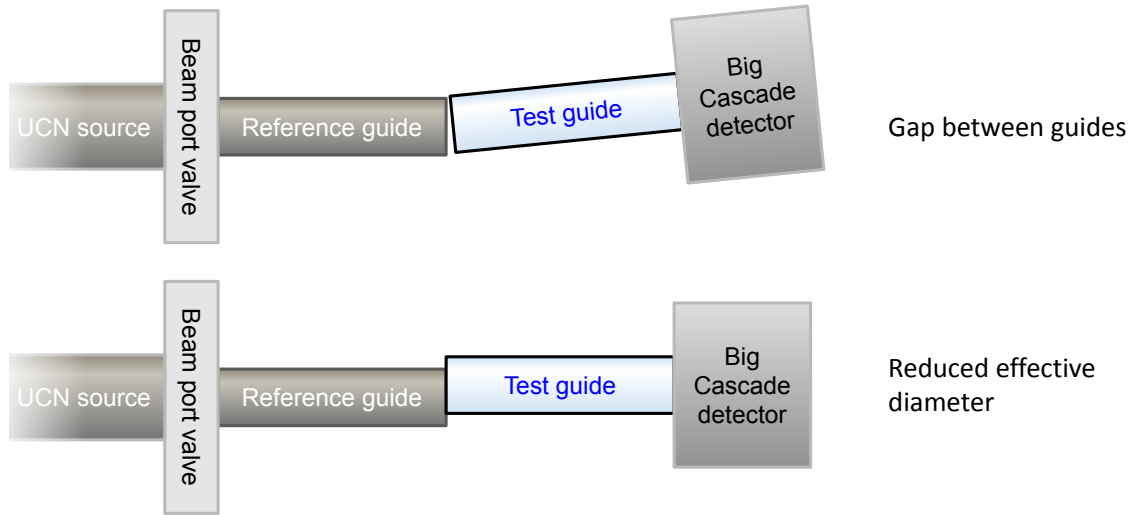
Figure 6.6: Two types of mounting errors that could reduce the measured transmission.

rate is larger as more time passes, as illustrated in Figure 6.8.

An angular misalignment of 1.2°, which is so large that it is easily visible to the naked eye and would be corrected, would create a gap of the same size as a linear separation between the guides of 1.5 mm. With a loss in transmission of 0.0062 mm^{-1}, this would reduce the transmission by 0.0076.

6.4.2   Effective local cross-section reduction

To investigate the effect of misalignment of the test guide with respect to the reference guide, we modified the standard transmission setup (Figure 6.1) by inserting steel rings, with reduced diameter, between the reference guide and the test guide (Figure 6.7b). We used the 1000 mm NiMo-coated reference guide, and HG1-3 (a 625 mm NiMo-coated guide for n2EDM) as a test guide.

The steel rings had a thickness of 1.0 mm, an outer diameter of 140 mm. The inner diameters were 110, 120, 125, 128 and 130 mm.

Steel was chosen because it has a Fermi potential of 188 neV [30], slightly lower than that of NiMo. UCN hitting the surface of the steel ring would most likely bounce back, as they would when hitting the exposed flat surface of a misaligned guide. Neutrons with an energy above 188 neV could potentially enter the steel. The rings had to be thick enough that such neutrons would be absorbed, yet thin enough not to cause a significant effect from reflections from the interior surface of the ring. The absorption cross-section of natural iron for 2200 m s^{-1} neutrons is 2.56 Bn [85], which translates into an attenuation length of 48 mm. According to MCUCN simulations and measurements [102], onwards from 12 seconds after the proton beam pulse, the most energetic UCN that still contribute to the measured signal have an energy of 300 neV. If these UCN enter the steel ring, these have a kinetic energy of 112 neV, after subtraction of the 188 neV Fermi potential of steel [30], and a velocity of 4.6 m s^{-1}. Since the absorption cross-section is inversely proportional to the neutron velocity, the attenuation length for these neutrons is 0.10 mm. The transmission probability for a 300 neV neutron is thus 4.5×10^{-5}. This probability is much smaller than the relative Poisson error on UCN counts in a single measurement. For UCN with energies below 300 neV, the transmission probability is even smaller.

The results are shown in Figure 6.9, where the transmission is plotted against the ring cross-section. The transmission depends on the UCN spectrum: higher energy UCN are more affected by the ring. More interestingly, the transmission is not proportional to the ring cross-section. Two effects could contribute to this. On the one hand, it is possible that UCN spend more time closer to the axis of the guide than to the

(a)

NiMo   Al   NiMo

130

135

d

(b)

NiMo   Steel   NiMo

130

D

1
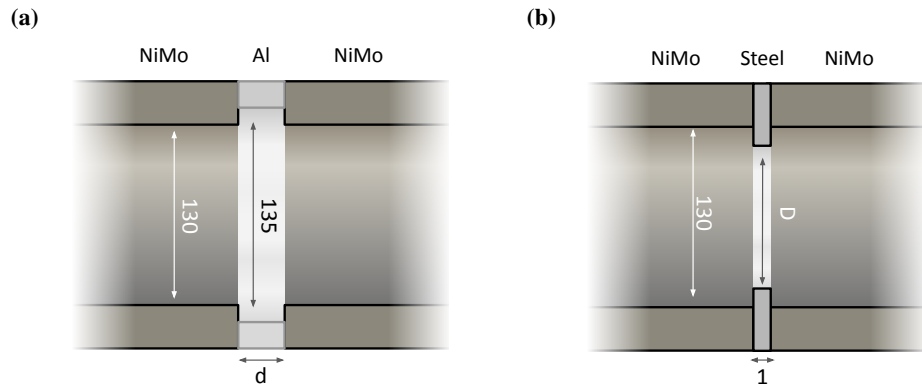
Figure 6.7: A schematic representation of the measurement setup we used to determine the effect mounting errors could have on the transmission values. Measurements are in mm. The two NiMo-coated guides are the reference guide and the test guide, as in Figure 6.1. (a) To investigate the effects of gaps in the setup, we inserted Aluminium rings with an inner diameter of 135 mm (5 mm larger than the inner diameter guide) with varying thicknesses, d. (b) To measure the effect of an effective diameter reduction, we used steel rings with a thickness of 1 mm and varying inner diameters, $D \leq$ 130 mm.

$\cdot 10^{-3}$

Loss per mm

6

4

2

0

0 to 12 s   12 to 30 s   30 to 60 s   60 to 100 s   100 to 200 s
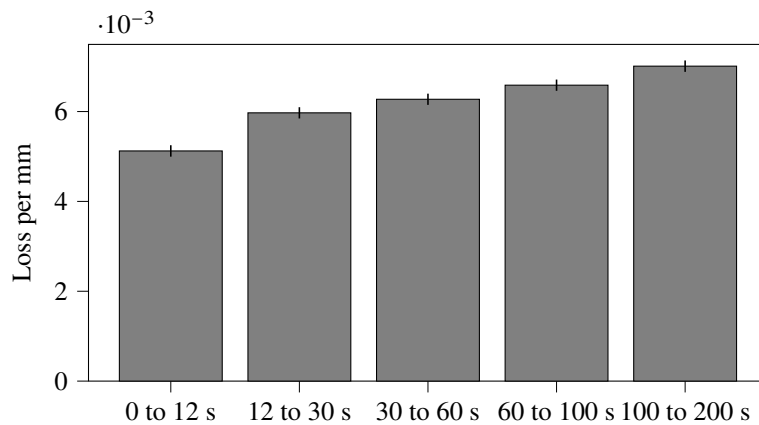
Figure 6.8: The reduction of the UCN transmission per millimetre of aluminium ring inserted between the UCN guides, calculated for different time bins. Beamport West-2 was used for normalisation.
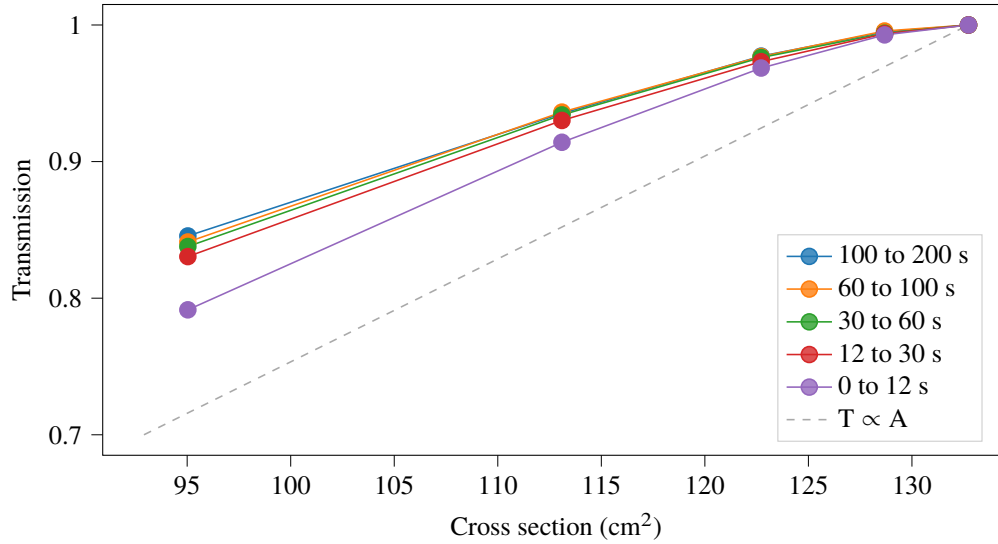
Figure 6.9: Transmission of UCN through a setup with 1 mm steel rings with varying diameters, placed between two UCN guides, as a function of the inner cross-section of the rings. Different curves represent different time bins. The transmission of the 130 mm ring (132.7 cm^2 cross-section) is one, by definition. The dashed line represents the scenario in which the transmission is proportional to the inner cross-section of the ring.

surface. On the other hand, since UCN are stored inside the guide, back-reflected UCN remain in the system (as opposed to UCN lost in the Aluminium rings in the previous section) and may pass through the ring at a later moment.

To calculate the axial displacement that would result in an equal reduction of cross-section as the insertion of a ring with diameter d, can take the formula for the overlap between two circles, with radius r, with their centres offset by distance d:

$$O = \int_0^{x_m} dx\sqrt{r^2 - x^2} - \frac{1}{2}d, \tag{6.5}$$

with

$$x_m = \sqrt{r^2 - \frac{1}{4}d^2}, \tag{6.6}$$

and solve it (numerically) for $O = r^2/r_{\text{guide}}^2$. To match the reduction in the cross-section of a 128 mm ring, the UCN guides would have to be off-axis by 3.1 mm. It is unlikely that the axial displacement is larger than this, because range of motion of the UCN guides is constrained by the Wilson flanges. If such large offsets were possible, they would have been visually detected by the experimenter and corrected. This corresponds to a loss in transmission of 0.0059. It is therefore implausible that a random axial offset between the UCN guides can explain the differences between UCN guide transmissions.

6.4.3   Reproducibility of measurements

The transmission of several guides has been measured multiple times.

On 2022-12-05, we measured the transmission of the 1 m NiMo-coated UCN guide three times, dismounting it and inverting its direction in between. The results of these repeated measurements are shown in Table 6.3.

From this series of measurements, one can estimate the statistical uncertainty of a single transmission measurement. The values obtained by using beamline South or West as a monitor are not independent. We

Table 6.3: Repeatability of UCN transmission measurements

Name	Date	Transmission South 12 to 200 s as monitor	Transmission West-2 30 to 200 s as monitor
HG1-3	2022-06-24	-	0.957
	2022-09-29	0.958	0.956
	2022-10-04	0.961	0.968
HG1-21	2022-06-10	-	0.964
	2022-06-17	0.963	0.971
HG1Ti	2022-06-09	-	0.917
	2022-06-10	-	0.913
	2022-09-27	0.927	0.928
1 m NiMo	2022-12-05	0.950	0.947
	2022-12-05	0.947	0.944
	2022-12-05	0.946	0.946

used the values using West-2 to estimate the reproducibility since fewer values are missing. For each set of measurements of a guide with at least 3 values, we take the standard deviation of transmission values obtained by using West-2 as a monitor, subtracting 1.5 degrees of freedom.[1] This gives standard deviations of 7.7 for HG1-3, 9.0 for HG1Ti, and 1.8 for the 1 m NiMo guide. Taking the root-mean-square of these, we obtain a standard deviation of 0.0069. This serves as the uncertainty on a single transmission measurement. Since the UCN count rate West-1 correlates much better with the count rate at South than the one at West-2 (see Chapter 5), we believe the reproducibility is better than 0.0069 if South is used as a monitor, but we do not have enough data to back this up. We therefore conservatively estimate the uncertainty on UCN transmission measurements, including the effects of imperfect mounting reproducibility, to be 0.007.

6.4.4 Interpretation

Mounting errors could introduce a random offset in the transmission measurements. We estimate that the maximum plausible loss in transmission due to angular misalignment is 0.0076, whereas the maximum plausible loss due to axial misalignment is 0.0059. This would add up to a maximum total loss in transmission of 0.0135.

One would expect the losses to follow a bell-shaped distribution, with a maximum between 0 and 0.0135. The standard deviation of this distribution is the uncertainty due to imperfect mounting reproducibility. The mean of the distribution can either be considered as a systematic bias to the transmission measurement, or part of the transmission value.

Since this distribution is unknown, we pessimistically assume that the losses due to mounting errors follow a uniform distribution:

$$\rho(x) = \begin{cases} \frac{1}{l} & \Leftarrow x - \mu < \frac{l}{2} \\ 0 & \text{otherwise} \end{cases}, \tag{6.7}$$

---

[1]The unbiased estimator of the population variance is obtained by multiplying the sample variance by $n/(n-1)$. This is called Bessel's correction. Even though this is an unbiased estimator of the variance, the square root of the unbiased estimator of the variance underestimates the standard deviation because in general, $\langle\sqrt{X}\rangle \neq \sqrt{\langle X\rangle}$. No general unbiased estimator exists for the population standard deviation, since this depends on the underlying distribution of the data [113]. A simple yet reasonable correction is to subtract 1.5 degrees of freedom instead of 1 when estimating the standard deviation [114].

with $l = 0.0135$ and $\mu = 0.0608$. The variance and standard deviation are

$$\sigma^2 = \int_{-\frac{1}{2}l}^{\frac{1}{2}l} dx \, \frac{x^2}{l} = \frac{l^2}{12}, \tag{6.8}$$

$$\sigma = \frac{l}{2\sqrt{3}} \approx 0.289 \, l. \tag{6.9}$$

This provides an estimate of 0.0039 for the statistical uncertainty due to imperfect mounting reproducibility.

In principle, this reasoning can be applied not only to the connection between the test guide and the reference guide but also to the connections between the Cascade detector and the test guide, as well as the reference guide and the BPV. The reference guide is not remounted between the installation of the test and reference setup, so mounting errors of the reference guide do not contribute to the variation of the measured transmission values. Allowing for variation of the connection with the reference guide increases the uncertainty by $\sqrt{2}$ to 0.0055.

This is comparable with the standard deviation of repeated measurements of the same guide, which is 0.0069.

In contrast, the differences in transmission between different guides are sometimes as large as 0.10 for guides with similar lengths. This effect is too large to be explained by mounting differences. We conclude that the differences are intrinsic to the guides.

6.5   UCN transmission of uncoated glass guides

To investigate the effect of incomplete coating coverage, we measured the transmission of uncoated glass guides. We used four different guides, with lengths 250, 500, 700, and 1250 mm (the last one is an uncoated HG3 guide for n2EDM). In addition, we measured the transmission of the 500 and 1250 mm guides mounted together, with a total length of 1750 mm. The reference guide was 300 mm, NiMo-coated. All guides had an inner diameter 130 mm. The 250 mm guide was slightly thicker than the other guides and did not fit in the Cascade detector. Therefore, when it was measured, the 250 mm uncoated guide was installed between the BPV and the reference guide. In this study, we assume that the order in which the guides are mounted does not matter within the precision of our measurement.

The results are shown in Figure 6.10. For each time interval, a simple exponential curve of the form

$$T(L; I, T_0) = I T_0^{l/l_0} \tag{6.10}$$

was fitted to the transmission values in this interval. Here I is the intercept (the transmission at $l = 0$), T_0 is the transmission at a reference length l_0 and l the UCN guide length. For $l_0 = 1$ m, we call T_0 the specific transmission.

Uncoated glass has a Fermi potential of approximately 88 neV [30], which is significantly lower than the Fermi potential of the NiMo-coated guides leading up to the glass guide. The part of the UCN spectrum with radial momentum components corresponding to kinetic energies above 88 neV is trimmed in the first collision(s). UCN with kinetic energies above 88 neV can still reach the detector if their radial momentum components correspond to energies below 88 neV. This trimming of the UCN spectrum is captured in the parameter I.

In reality, the UCN transmission should approach one as the UCN guide length approaches zero, as there are more and more neutrons that pass through the uncoated guide without wall collisions. We might have seen this effect if we had included UCN guides under 250 mm. In that case, we could fit the data with a double exponential. Equation 6.10 is an approximation that is valid for UCN guides that are sufficiently long that all UCN collide with the wall.
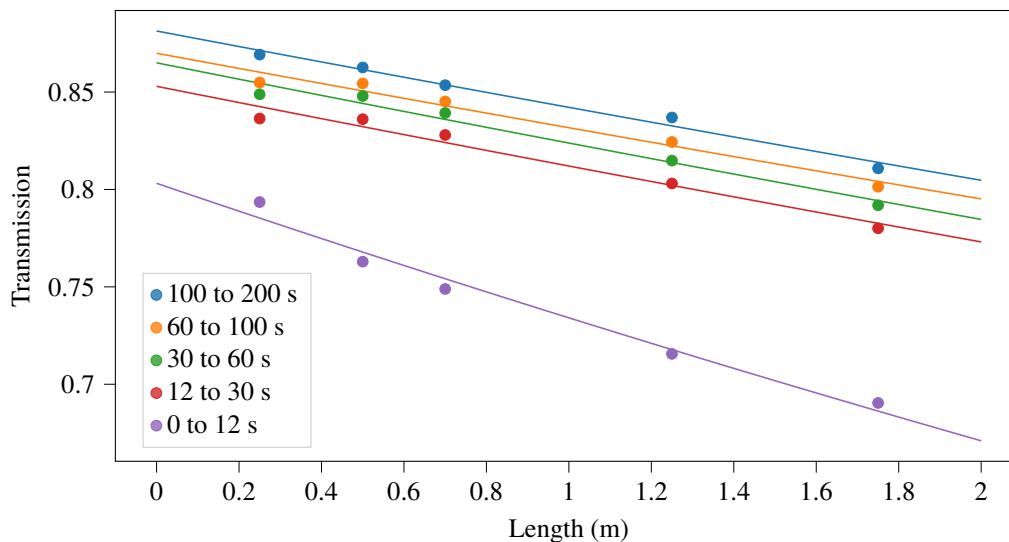
Figure 6.10: Transmission of UCN through uncoated glass UCN guides as a function of UCN guide length. The points represent different time intervals for each measurement. The lines represent exponential fits to the data. The measurement at $l = 1.75$ m was done by installing the 1250 mm and 500 mm guides in series. These points are therefore not independent of the rest. When we measured the 250 mm guide, the order of the test and reference guides was inverted, for technical reasons.

For neutrons arriving between 12 and 200 s after the proton beam pulse, the fitted specific transmission is 0.952 with an intercept of 0.861. This means that if the length is extrapolated to 0, the transmission would be 0.861, while every additional meter of guide would reduce the transmission with a factor 0.952.

The specific transmission of 0.952 is comparable to the best results from the transmission measurements of coated UCN guides for n2EDM (Table 6.2). If this is broken down into smaller time bins, the specific transmission does not change, but the intercept does, with later time bins having a higher intercept. This can be explained by the softening of the UCN spectrum over time. For neutrons arriving in the first 12 seconds after the proton beam pulse, the transmission, and the intercept are significantly lower, 0.803 and 0.914, respectively.

The calculated transmission values depend on the chosen time interval (Figure 6.10). This is true for coated and uncoated guides. However, for uncoated guides, this strongly depends on the way monitoring is done. When estimating the transmission in a given time interval, one can use the UCN counts in South in either the corresponding interval or the interval from 12 to 200 seconds, for normalisation. These are both valid methods that measure slightly different things. In Figure 6.11, the transmission of a coated and an uncoated guide are plotted for several time intervals using both methods of normalisation. For the coated guide, the transmission improves over time, which is most likely caused by the softening of the UCN spectrum. For the uncoated guide, this is only true if normalisation is done using the matching time interval in South; when the counts in South from 12 to 200 s are used, transmission decreases over time. The reason for this is that a system containing uncoated UCN guides leaks UCN. Even though this UCN loss happens in beamline West-1, the detector on beamport South is sensitive to this. When one uses the corresponding time interval for normalisation, one compensates for this loss.

6.5.1 Interpretation

The hypothesis that differences in transmission are adequately explained by incomplete coating coverage can now be tested. One guide with one of the worst transmission values was HG4-1 (595 mm): 0.903. We estimate its transmission would have been 0.836 if it had been uncoated, from the fitted parameters and

(a)

(b)
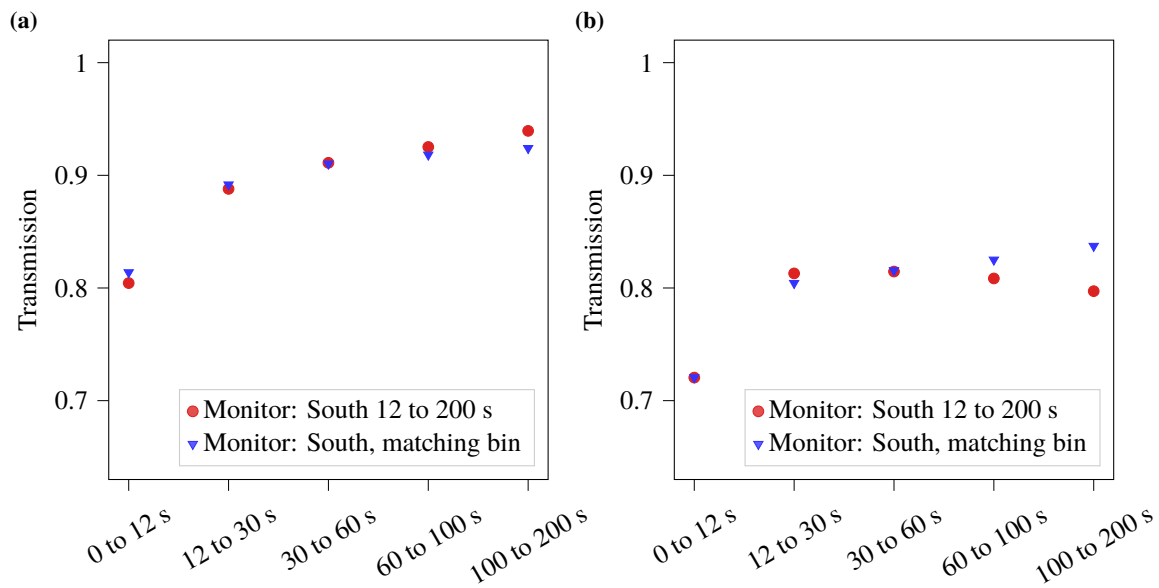
Figure 6.11: UCN transmission calculated for different time intervals, for (a) HG4-1, a 595 mm guide coated with NiMo, and (b) an uncoated 1250 mm guide. South is used as a monitor. The red marks represent the case where the counts in South integrated from 12 to 200 s are used to normalise the counts in all intervals in West-1. The blue marks represent the case where each of the counts in each interval is normalised using the corresponding interval in South. For the coated guide, the transmission gets better for later intervals, most likely due to the softening of the UCN spectrum, lowering the loss per bounce. For the uncoated guide, this depends on the way South is used to monitor West-1. If the counts from 12 to 200 s are used for normalisation, the transmission decreases for later time bins due to UCN losses. If the corresponding interval is used for normalisation, the UCN losses are partially compensated and the transmission improves for later intervals.

Equation 6.10. If the guide had the same specific transmission (0.952) but an intercept of 1, the transmission would have been 0.971. The measured transmission is thus about 0.07 lower than the ideal transmission, but still 0.07 higher than if it was completely uncoated. Assuming that the coated sections of the guide had a high specific transmission, the total transmission could be explained if half of the UCN entering the guide experienced at least one collision with an uncoated section.

Since the shortest uncoated guide already had a transmission of around 0.85, we lack the experimental data to estimate the required uncoated surface area that would reduce the transmission by 0.07. We can determine a lower bound for the minimum exposed surface area from the transmission measurements with aluminium rings. We found that each mm of aluminium ring thickness reduced the transmission (12 to 200 s) by 0.0062. An 11 mm thick aluminium ring would reduce the transmission by 0.07, with an inner surface area of $45 \, \text{cm}^2$.

Glass has a smoother surface and a higher Fermi potential than aluminium. For exposed glass to reduce the transmission by 0.07, the surface area should be more than $45 \, \text{cm}^2$. We determined that the uncoated area of the guides is well below $5 \, \text{cm}^2$ by visible inspection of the guides. Incomplete coating coverage thus cannot explain a difference in transmission larger than 0.008.

Unlike the coated guides, the four uncoated guides can be represented by a single specific transmission. This could mean that the variation in transmission is introduced by the coating. However, except HG3-u, the uncoated guides that we measured were not produced in the exact same way, and by the same manufacturer, as the UCN guides for n2EDM. It is possible that the glass blowing and forming process for the n2EDM guides introduced variation in transmission. We tested this hypothesis, as described in the next section.

6.6 UCN guides from tubes from two different manufacturers

The production process for the UCN guides for n2EDM caused ripples inside the guide. In this section, we investigated whether this could have caused the differences in transmission.

Part of the UCN guides for n2EDM have been ordered from GlasForm[2]. The guides were initially made with an outer diameter slightly larger than the specified 140 mm. To ensure the guides can be mounted in n2EDM, the glassblower rolled the ends of the guides until the outer diameter was 140 mm. This process causes ripples on the interior of the guide, as seen in Figure 6.12. The ends of the guide thus have a higher 'waviness' than the central section.

If these ripples in the first centimetres of the guide significantly reduce the UCN transmission, it could explain why the quality of the guides is so different. Under this hypothesis, it is also clear why the 1250 mm guides should have a good specific transmission because the ripples in these guides occupy a smaller fraction of their length.

To investigate this, we ordered two new 595 mm HG4 guides: the first one again from GlasForm, and the second one from Schmizo AG[3], a Swiss company which manufactures specialised scientific equipment from borosilicate glass to compare with the one from GlasForm. We coated both guides in the same process as the other guides, in a single production cycle, to ensure that variations in the coating production cannot contribute to the differences in surface quality.

6.6.1 Profilometry

We measured the surface roughness of the guides before coating them using a handheld profilometer, the HANDYSURF+[4]. The profilometer has a precision of 0.7 nm. The setup is shown in Figure 6.13. We measured

---

[2]GlasForm J. Städler u. Co, Güetli 15, 9100 Herisau, Switzerland, https://glasform.ch/kontakt
[3]SCHMIZO AG, Grünmattstrasse 4, 4665 Oftringen, Switzerland, https://schmizo.ch/
[4]https://www.zeiss.com/metrology/products/systems/form-and-surface/surface-measurement/handysurf.html

(a)   **(b)**
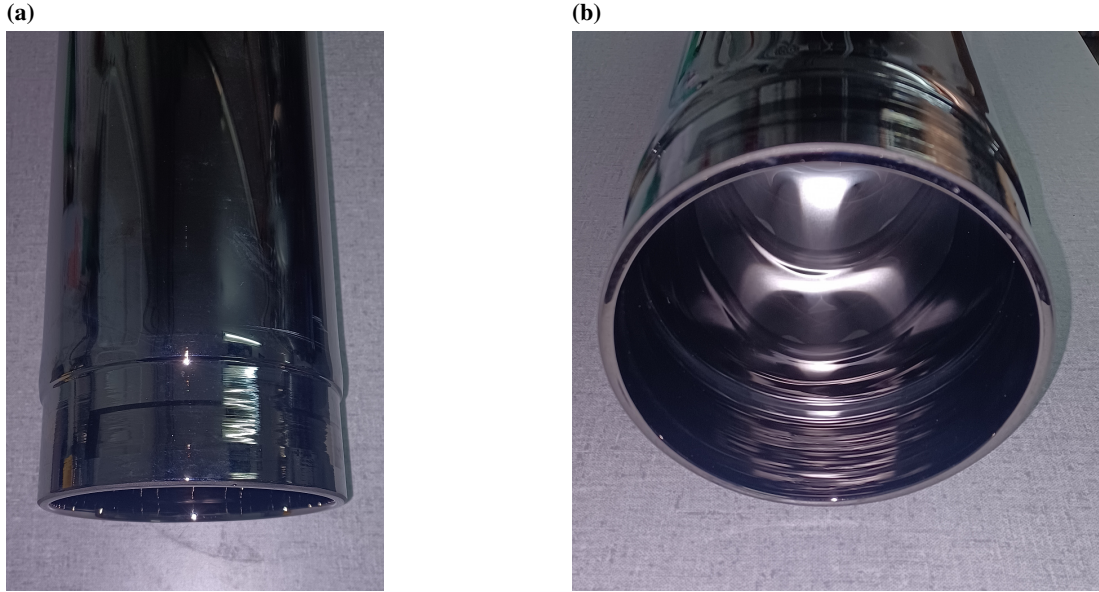
Figure 6.12: Photographs of one end of HG4-3: a NiMo-coated, 595 mm UCN guide for n2EDM, manufactured by GlasForm. (a) The last 45 mm of the guide were rolled to ensure the outer diameter is 140 mm, within the required tolerance. (b) The first centimetres of the interior surface of the guide are less smooth than the rest, as a result of the rolling. This is visible in the photograph, the reflections are clearer further inside the guide.

the surface roughness in several regions within 5 cm from the end of both guides and of sections in the interior (more than 5 cm from the end). Each measurement was done over a length of 4.8 mm, with a resolution of 4.2 μm. Some profiles are shown in Figure 6.14a.

The roughnesses were calculated from the surface profiles (Figure 6.14b). There was no significant difference between the roughnesses in the interior, or close to the ends of the Schmizo guide, the mean roughness was (6.7 ± 1.1) nm. The GlasForm guide had a mean roughness of (33 ± 5) nm within 5 cm of the ends, and (7.7 ± 2.3) nm in its interior.

The measurement confirms the visual observation that the interior of the GlasForm guide has a similar roughness to the Schmizo guide, but that the ends of the GlasForm guide are significantly rougher.

6.6.2   Transmission measurement setup

The transmission of the guides was measured in December 2023, at beamport West-2.[5] This required a modified setup (Figure 6.15), in comparison to the measurement at West-1.

Beamline West-2 is attached to the UCN source storage volume two meters higher than beamlines West-1 and South. In addition, the UCN in West-2 must pass through several bends before reaching the detector. Therefore, the energy spectrum and angular distribution of UCN at West-2 are different to West-1. We remeasured HG4-3 and the 50 cm uncoated guide in this new setup to understand the relation of the transmissions obtained in West-2 and West-1.

A stainless steel 90° bend was installed on the VAT shutter, followed by the 300 mm NiMo-coated reference guide. In the reference setup, the large Cascade detector was installed at the end of the reference guide. In the test setup, the test guide was connected to the reference guide, followed by the large Cascade detector.

---

[5]At this time, the neutron lifetime experiment, τSPECT [35], was taking data at beamport West-1, and n2EDM was taking data at beamport South.
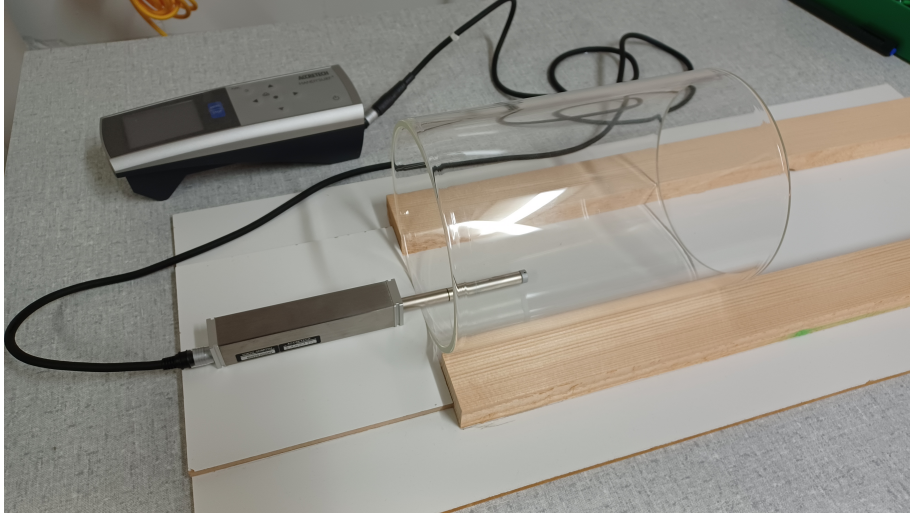
Figure 6.13: The setup for profilometry of UCN guides. Two wooden beams hold the UCN guide in place. The HANDYSURF+ profilometer is placed on a 5 mm plywood platform so that it is at the same height as the inner surface of the guide. A diamond tipped stylus is drawn over the guide surface to measure the profile.

Photographs of the setup are shown in Figure 6.15.

6.6.3   Analysis

As there was no UCN beamport available for monitoring, we had to rely on the HIPA proton beam monitor (MHC1) to normalise the measurements. This removes the effect of proton beam fluctuations, but it does not compensate for the drift in UCN yield caused by the degradation of the surface of the UCN source deuterium moderator. This effect is not linear. Therefore, where we normally fit a line (Eq. 6.3), we fitted a cubic spline (see Chapter 3). After every conditioning, a new spline was fitted.

During the operation of the τSPECT experiment [35], the NLK in beamline West-1 was opened and closed. This leads to a change in the UCN count rate. The status of NLK West-1 is logged by the UCN source DAQ system, we included it in the fit function. NLK South was open for the entire duration of the measurements, it was not included in the fit.

When n2EDM was taking data, usually BPV South was open at the start of the proton beam pulse, and was closed some time after, with variable timing. For some measurements, the BPV stayed open for the entire 300 s period between two consecutive pulses. Whether BPV South was initially open only had a small effect on the UCN count rate on beamport West-2. The fit could not be improved further by taking the exact timing of the closing of the BPV into account, which was determined using the Bayesian Information Criterion [105], [115].

The final fit function is:

$$f(Q, s_W, s_S, i; \{y_s\} r_W, r_S, \{T\}) = S(Q; \{Q_s\}, \{y_s\})(1 - s_W r_W)(1 - s_S r_S)T_i. \tag{6.11}$$

The independent variables are: Q, the cumulative beam current; s_W and s_S, the statuses of τSPECT and n2EDM, 1 if they are receiving UCN, otherwise 0; i, the UCN guide index. The fit parameters are $\{y_s\}$, the set of normalised UCN counts corresponding to the fixed set of cumulative beam currents, $\{Q_s\}$, in the cubic spline; r_W and r_S, the fractional reduction in UCN counts when τSPECT or n2EDM are operating; and $\{T\}$, the set of UCN guide transmissions, where the value of the reference transmission is defined as 1. S is the cubic spline.

(a)

(b)
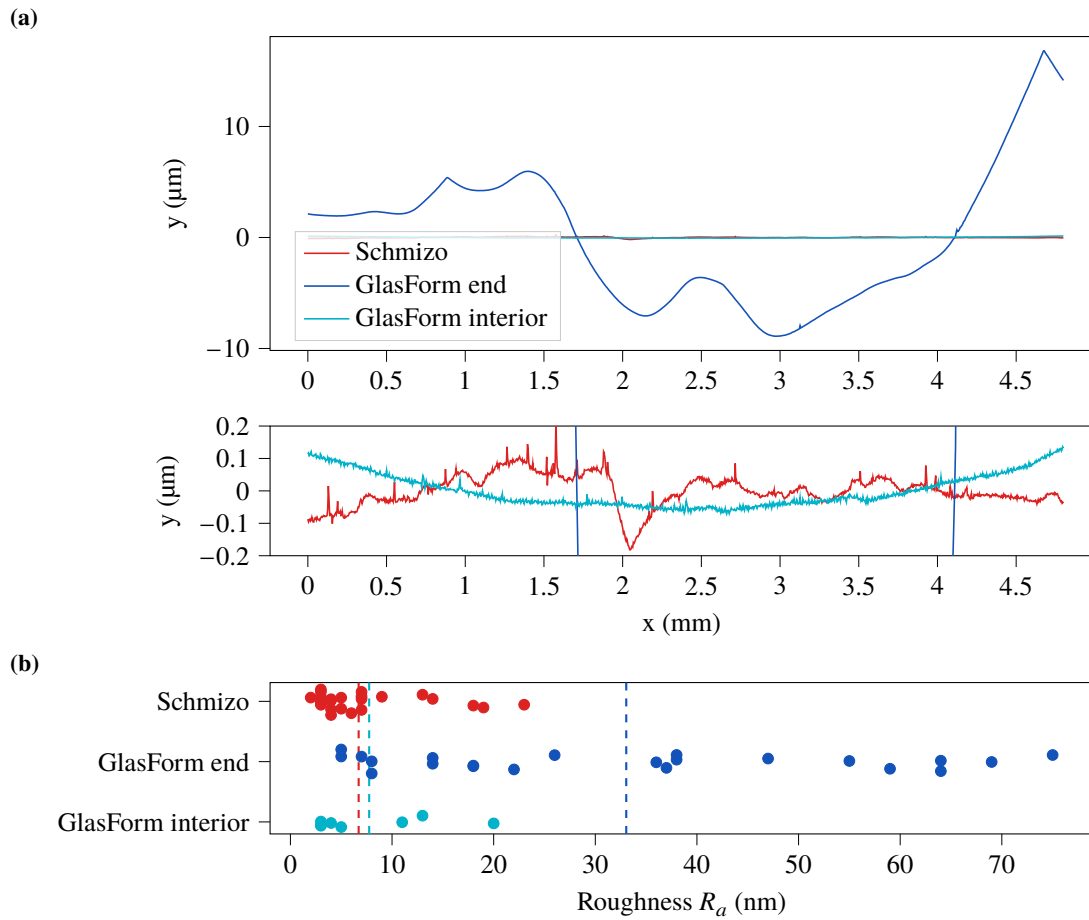
Figure 6.14: Roughness measurement of the uncoated, 595 mm GlasForm and Schmizo guides. (a) Some examples of the measured profiles of the GlasForm and Schmizo guides. (b) The distribution of roughness values, R_a, from the scans. The mean roughnesses of the Schmizo guide and the interior of the GlasForm guide are similar, but that of the GlasForm guide is much larger. The dashed lines indicate the mean roughnesses.

(a)

(b)
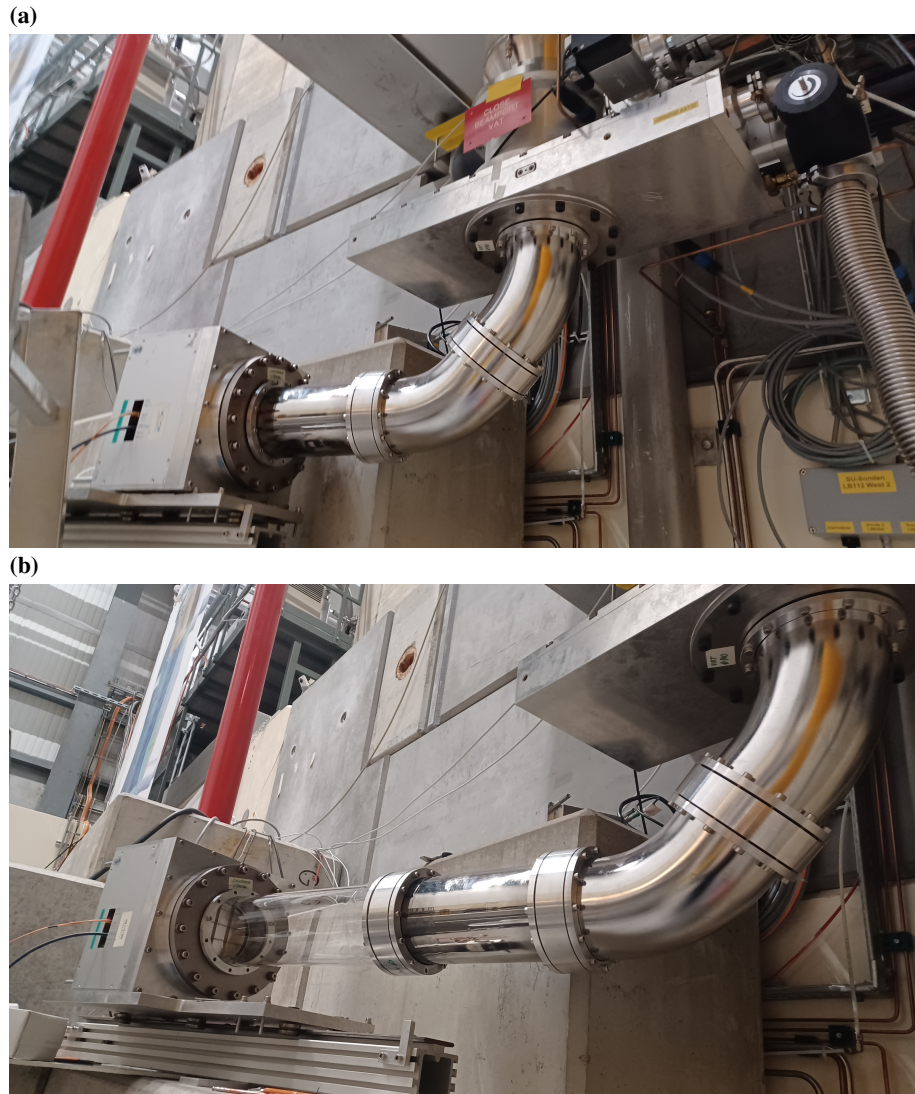
Figure 6.15: The transmission measurement setup on beamport West-2, in December 2023. (a) The reference setup. (b) A test guide is installed (500 mm uncoated). UCN fall through the beamport valve (the wide, rectangular object), through the steel bend (comprised of two 45° turns), and into the test guide and reference guide, which are mounted horizontally.

To determine the points $\{Q_S\}$ for the spline, we chose the minimum and maximum cumulative beam in the data and evenly spaced points in between. We chose the spacing such that the number of parameters minimises the Bayesian Information Criterion [105], [115] of the fit, to determine the number of points so that the spline accurately captures the data without overfitting. We chose the largest possible spacing that is still below $120 \cdot 1.8 \cdot 8$ mA s.

6.6.4   Results

The results of the experiment are shown in Figure 6.16. The fit accurately captures the data. The fitted transmission values and the effects of the status of NLK West-1 and n2EDM are shown in Table 6.4, where for consistency r_W and r_S are represented as a transmission, $T = 1 - r$. The fit error is 2 % for the transmissions, and 1 % for r_W and R_S. The value for SB45-1 is not reliable, since there were large fluctuations in UCN output that are poorly understood.

The West-2 beam guide is attached to the UCN source storage volume two metres above the West-1 and South beam guides. UCN must have an energy of at least 330 neV, measured at the deuterium surface, to reach beamline West-2. The UCN fall one metre before reaching the beamport [37], and another 40 cm in the steel bends (Figure 6.15). This acts as an energy filter. UCN that enter the reference guide, must have a minimum kinetic energy of 140 neV. Since the UCN have to pass through a steel bend just before they enter the reference guide, the UCN are less collimated than at West-1 or South. Because of both factors, the measured transmission is lower in the setup at West-2 than in the setup at West-1.

The guide provided by GlasForm had a slightly higher, but comparable transmission to HG4-3, which was also provided by GlasForm. The Schmizo guide, which is visually smoother, had a better transmission than both GlasForm guides, which supports our hypothesis.

To make a crude estimate of what the transmission of the GlasForm and Schmizo guides would be if it had been measured at beamport West-1, we observe that $\log(T_{W2})/\log(T_{W1})$ is about 1.57 for both the uncoated 500 mm guide, as for HG4-3. If this ratio holds for the GlasForm and the Schmizo guides, their transmission, had they been measured at West-1, would be 0.93 and 0.96, respectively.

6.6.5   Conclusion

The UCN transmission of the Schmizo guide was better than that of the GlasForm guide. Both guides had been coated in the same production cycle, eliminating variation between coating productions as a possible explanation for differences in UCN transmission. This suggests that by using the smoother glass tubes from Schmizo for UCN guide production we could significantly improve UCN transmission.

We are basing this on UCN transmission of a single UCN guide, made with a glass tube produced by Schmizo, we currently do not have enough evidence to state that all glass tubes produced by Schmizo will provide better guides. Nevertheless, we decided that it is worthwhile to replace six of the horizontal guides in n2EDM (HG1, HG3 and HG4, top and bottom) with UCN guides made from Schmizo glass tubes. We will then learn if this leads to a significant increase of UCN in the precession chambers.

6.7   Determining the Fermi potential using cold neutron reflectometry

During the coating process of the UCN guides, we attach standard microscopy slides to the ends of the UCN guides, as witness samples. We assume that these slides receive the same coating as the guides, and use them for diagnostics of the coating process. The slides are made of blank float glass, with a surface roughness of $R_a < 1$ nm, which we measured using the DekTak profilometer. The dimensions are 75 mm×25 mm×1.0 mm. The slides are attached to the outside of the tubes using scotch tape. Typically, one sample is attached on
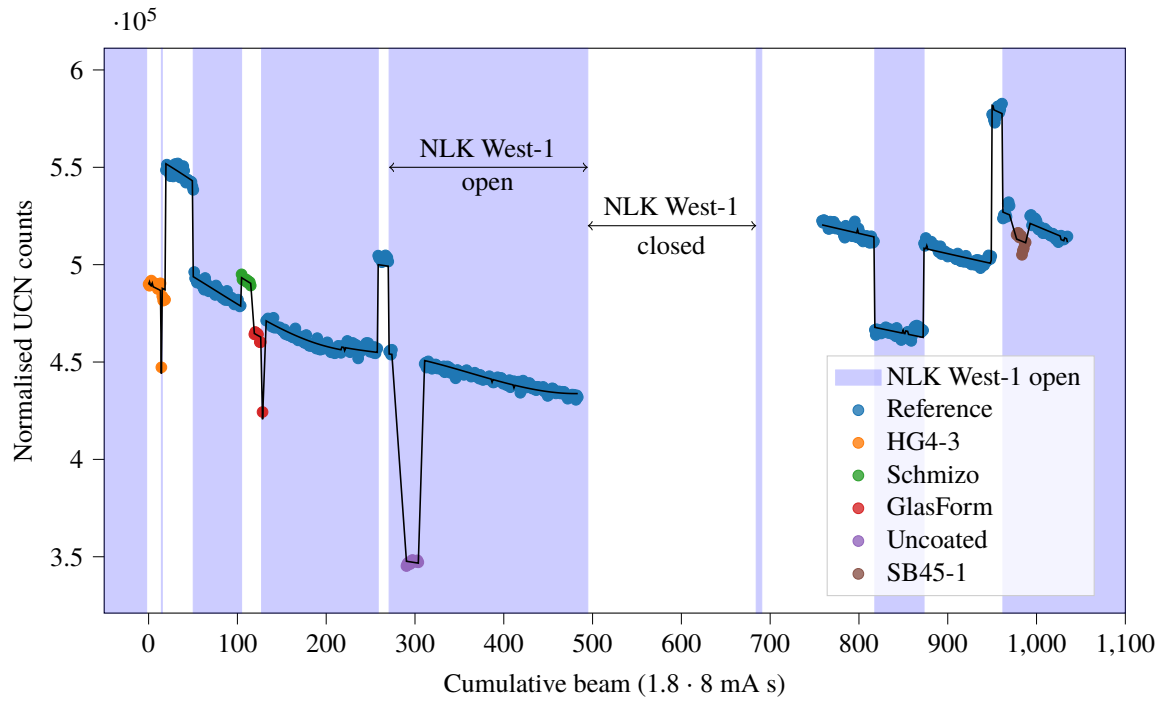
Figure 6.16: The UCN counts from 12 to 200 s normalised by the integrated beam current during the pulse (in standard pulses, $1.8 \cdot 8$mA s) measured in December 2023. From the UCN counts in the reference setup, the effect of the opening and closing of NLK West-1 is clearly visible. This effect also explains the measurements of HG4-3, and the GlasForm guide, which otherwise would look like outliers. The black line represents the fit.

Table 6.4: Transmission of guides with smoother and rougher ends.

Name	Length (mm)	Transmission West-2	Transmission West-1
HG4-3	595	0.879	0.922
GlasForm (rougher)	595	0.891	
Schmizo (smoother)	595	0.938	
Uncoated	500	0.768	0.844
SB45-1	450	0.976	0.968
NLK West-1 open[a]		0.910	
n2EDM taking data[a]		0.997	

[a] When τSPECT [35] was taking data, NLK West-1 was open. When n2EDM was taking data, BPV was opening and closing. These factors reduced the count rate in West-2. The numbers in the table are not transmission values, but they are multiplied with the count rate as factors, like the transmission values.
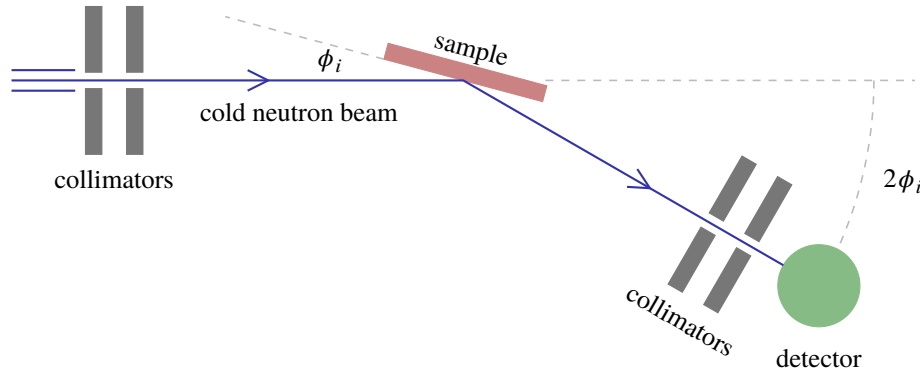
Figure 6.17: A sketch of the Narziss cold neutron reflectometry instrument, including the rotation angles of the sample table and the detector, used in the text.

the side facing the sputter head inside the coating machine (side B for Berg, because it faces the hills next to PSI), and one sample facing away from the sputter head (side A for Aare, since it faces the Aare). A small piece of tape ($0.5\,\text{cm}^2$) is stuck to the face of the slide facing inward. This piece of tape is removed after coating, creating a sharp edge between the coated and uncoated surface of the glass. A surface profilometer can then be used to measure the coating thickness at this edge.

Using cold neutron reflectometry, we measured the Fermi potential of the NiMo coating on the witness samples. This was done using the Narziss [116] polarised cold neutron reflectometry facility at one of SINQ beamlines at PSI. Narziss provides a beam of monochromatic neutrons with a wavelength of 4.97 Å [38] and an energy of 3.312 meV.

A drawing of the reflectometry setup is shown in Figure 6.17. The sample is cleaned using a cloth wetted with isopropyl alcohol. The sample is mounted on a rotating sample table (also cleaned with isopropyl alcohol) in the beamline, using a suction system. A neutron detector is placed behind the sample on a rotating arm with a vertical rotation axis. During a measurement, we vary the rotation angle of the sample table, ϕ_s, and the rotation angle of the detector, ϕ_d, such that $\phi_d = 2\phi_s$.

6.7.1 Analysis

Whether neutrons are reflected or transmitted depends on the momentum component perpendicular to the surface,[6]

$$p_\perp(\phi_s) = |\mathbf{p}| \sin(\phi_s). \tag{6.12}$$

At small angles ($p_\perp^2/2m_n < V_F$), the neutrons are totally reflected (neglecting neutron capture, see Figure 1.1 in the introduction). At angles larger than the critical angle, neutrons start entering the surface. The momentum inside the sample, $q_\perp$, is given by

$$q_\perp(\phi_s, V_F) = \sqrt{p_\perp(\phi_s)^2 - 2m_n V_F}. \tag{6.13}$$

The fraction of reflected neutrons (neglecting neutron capture) is given by the piecewise function

$$R(\phi_s; V_F) = \begin{cases} \left(\frac{p_\perp - q_\perp}{p_\perp + q_\perp}\right)^2 & \Leftarrow p_\perp^2 \geq 2m_n V_F \\ 1 & \Leftarrow p_\perp^2 < 2m_n V_F \end{cases}, \tag{6.14}$$

---

[6]We are working with grazing angles, not angles of incidence. So $\phi = 0$ is parallel to the surface and $\phi = \frac{\pi}{2}$ is perpendicular to the surface.
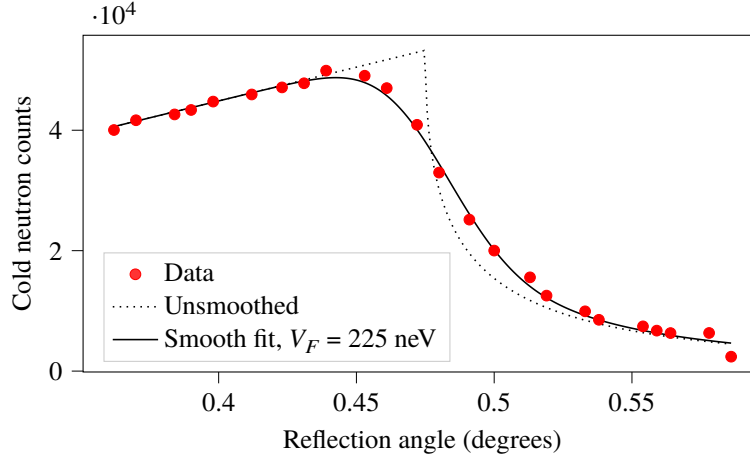
Figure 6.18: The observed cold neutron counts during a reflectometry measurement of a sample. Gaussian smoothing has to be applied to the fit function to obtain a better fit. The parameters for the exact formula (C) are the ones that were fitted to the smooth fit function ($\tilde{C}$). Fitting the exact formula directly would overestimate the Fermi potential by 9 neV.

where the dependency of $p_\perp$ and $q_\perp$ on ϕ_s and V_F is not explicitly written. The observed neutron count rate is then

$$C(\phi_s; V_F, A) = A \sin(\phi_S) R(\phi_s, V_F),\qquad(6.15)$$

where A is a scale parameter, and $\sin(\phi_s)$ is a geometrical correction, proportional to the cross-section of the sample in the beam.

Due to the finite width of the collimators (1 mm) and the surface roughness of the sample, the phase space of the beam is small, but not point-like. The shape of the curve around the critical angle is not as sharp as in the piecewise formula. We can account for this by taking the convolution of the fit function with a Gaussian to obtain a smoothed fit function:

$$\tilde{C}(\phi_s; V_F, A, \sigma) = \int_{-\infty}^{+\infty} d\phi \; \frac{e^{-\frac{\phi^2}{2\sigma^2}}}{\sigma\sqrt{2\pi}} \; C(A, \phi_s - \phi, V_F).\qquad(6.16)$$

In this fit function, ϕ_s is the independent variable, A and σ are nuisance parameters, and V_F is, of course, the parameter of interest.

An example fit is shown in Figure 6.18. In addition to the data point and the fit ($\tilde{C}$), the exact formula (C) is shown to demonstrate why this Gaussian smoothing is required to obtain a good fit.

The coating thickness usually varies between 250 and 600 nm, but is definitely under 1 μm. Assuming that the coating thickness is sloped by 1 μm over a length of 25 mm, this tilts the sample by 0.12 mrad. This creates a contribution of 4.3 neV to the uncertainty. This could have been prevented by recalibrating the setup for every sample, which was not done for time reasons, since differences of 4.3 neV are not relevant.

6.7.2 Results and interpretation

For selected guides, the measured Fermi potentials are shown in Table 6.5. All measured samples had a Fermi potential between 200 and 230 neV. There was no relation between the Fermi potential measured for the witness samples and the UCN transmission of the guide.

Strikingly, the samples attached to the UCN guide on side A, facing away from the sputter head, always have a lower Fermi potential than the samples on side B, facing the sputter head. The neutron optical potential

Table 6.5: Fermi potentials measured in the witness samples using cold neutron reflectometry.

| UCN guide | V_F (neV) | |
	side A[a]	side B[b]
HG3-2	208	229
HG1-TI	202	220
HG1	216	-
HG4-1	208	-
HG4-2	205	218
HG1-3	204	223
SB45	202	219

[a] 'Aare' side, facing away from the sputter head when the guide is mounted on the trolley in the coating machine.

[b] 'Berg' side, facing the sputter head.

of a material depends linearly on its density:

$$V_F = \frac{2\pi\hbar^2}{m_n} \sum_i b_i n_i \tag{6.17}$$

where b_i is the scattering length of a specific isotope, and n_i is the number density of this isotope in the substance [30], [36]. It is therefore likely that the density of the NiMo deposited on side A has a 8 % lower density than the NiMo deposited on side B. A possible explanation is that the samples on the B side receive nickel and molybdenum ions with steeper angles, compared to the sample on side A. Given this finding, it is recommended that vertical guides are installed with the B side (with higher V_F) pointing downwards.

6.8 The effect of surface parameters on UCN transmission in simulation

To investigate which properties of the interior surfaces of the UCN guides are most likely to explain the large variation in specific transmission, we simulated the UCN guide transmissions using MCUCN. In the simulation, we implemented the transmission setup as it had been installed at beamport West-1 (Figure 6.1). During all simulations, the length of the reference guide was set to 300 mm.

6.8.1 time-of-arrival distribution

At first, we simulated the reference setup, and a 615 mm guide with standard surface parameters, to verify that the time-of-arrival distributions make sense (Figure 6.19). These are very comparable to the time-of-arrival distributions obtained in the measurement (Figure 6.2).

6.8.2 Variation of surface parameters

The parameters and their default values

The default values used in MCUCN for the surface parameters of NiMo-coated guides in the UCN source, are $V_F = 220$ neV, $d = 0.05$ and $\eta = 3{\times}10^{-4}$. The Fermi potential, V_F, and the loss coefficient, η, can be
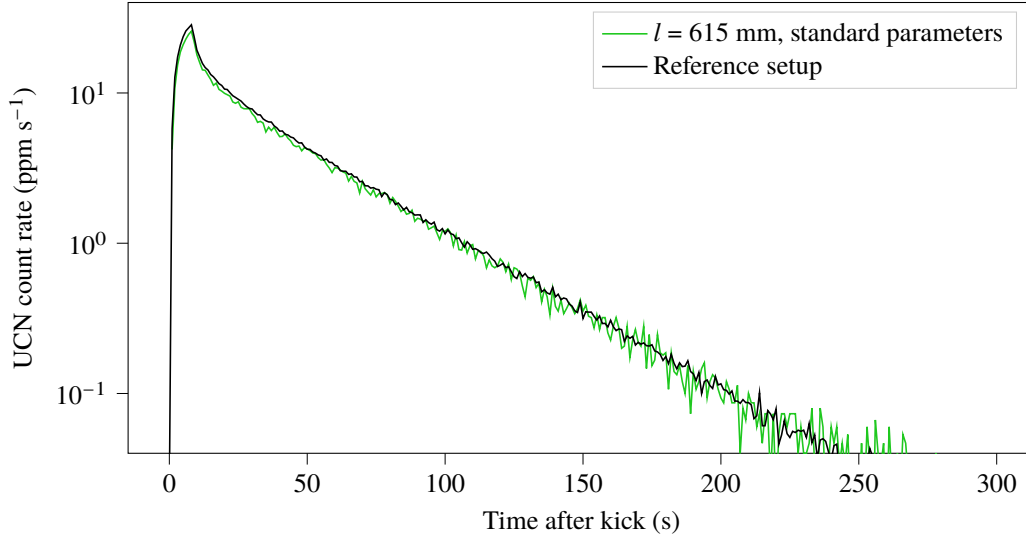
Figure 6.19: The simulated time-of-arrival distributions for the reference setup and the standard guide with l = 615 mm, V_F = 220 neV, and η = 3×10^{-3}. The count rate is defined as the fraction of the total number of simulated trajectories that arrives in the detector per second. The number of simulated trajectories was 1.5×10^8 for the standard 615 mm guide, and 1×10^9 for the setup. Therefore, the statistical noise is smaller for the reference setup.

calculated from the neutron capture and (in)elastic scattering cross-sections for nickel and molybdenum for $2200\,\mathrm{m\,s^{-1}}$ neutrons [30], [85].

The loss coefficient, η, is the ratio of the imaginary and real parts of the Fermi potential. It represents losses due to inelastic scattering and neutron capture in the material. In our simulations, we use an effective loss coefficient, which includes contributions from pinholes in the coating and gaps between guides. It is therefore higher than the value calculated from cross-sections (3×10^{-4} versus 1.2×10^{-4} [30]). The loss coefficient and diffusive scattering fraction have been fitted to experiments at the PSI UCN source [101], [117].

Results

We simulated UCN transmission through a 615 mm long guide, which is between the lengths of HG1 and HG4. We varied the surface parameters V_F, d and η around these default values for NiMo guides. We also varied the length of the guide. The UCN transmission as a function of these parameters, broken down by time interval, is shown in Figure 6.20.

For l = 615 mm and default surface parameters, the transmission is (0.939 ± 0.004) for UCN counted between 12 and 200 seconds after the proton beam pulse (the uncertainty is statistical only). This corresponds to a specific transmission of 0.902, which is in the middle of the range we found experimentally (Table 6.2).

The Fermi potential is linearly dependent on the density of a material [30], [36]. Variations in the deposition of NiMo during the coating process can thus affect the Fermi potential. As illustrated in Figure 6.20a, the dependence of the transmission on V_F is small around V_F = 220 neV. The guide would need to have a Fermi potential of 160 neV to explain a specific transmission of 0.84 (the lowest we measured). From neutron reflectometry (Section 6.7) we know that the Fermi potential of the guides, including the ones with a low specific transmission, was at least 200 neV. Hence, variation in Fermi potential alone cannot explain the observed variation in transmission.

The transmission has a small dependence on the loss parameter around its default value, η = 3×10^{-4}.
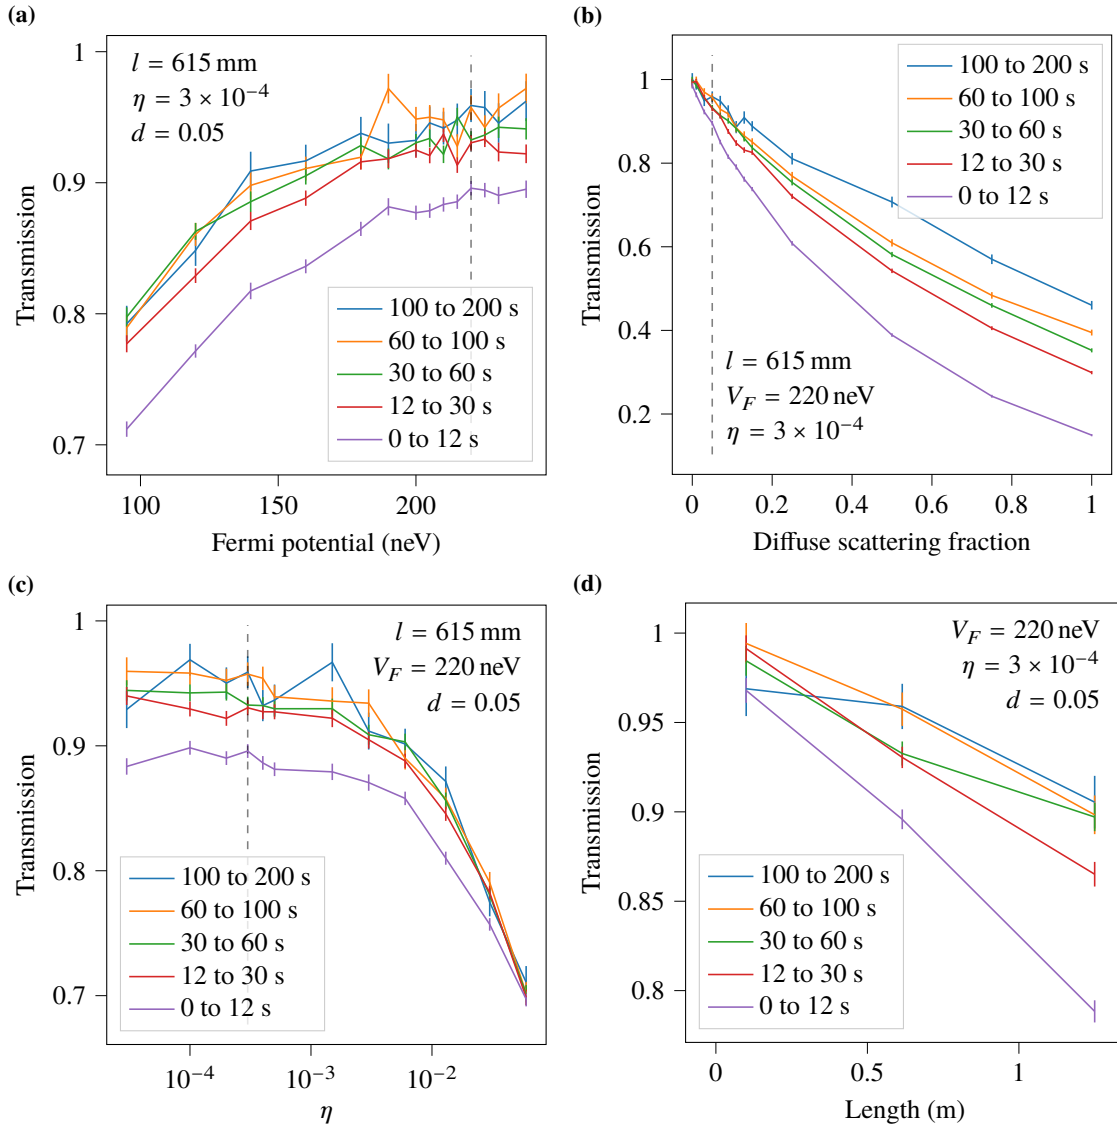
Figure 6.20: The simulated UCN transmission of a UCN guide as a function of (a) Fermi potential, (b) diffuse scattering fraction, (c) the loss coefficient and (d) the length of the guide. In each plot, the parameters that were not varied were set to their default values: $V_F = 220\,\text{neV}$, $d = 0.05$, $\eta = 3 \times 10^{-4}$. The dashed line indicates the default value of the varied parameters. The dimensions of the simulated guide were $l = 615\,\text{mm}$ and inside diameter $= 130\,\text{mm}$, the length was varied in (d).
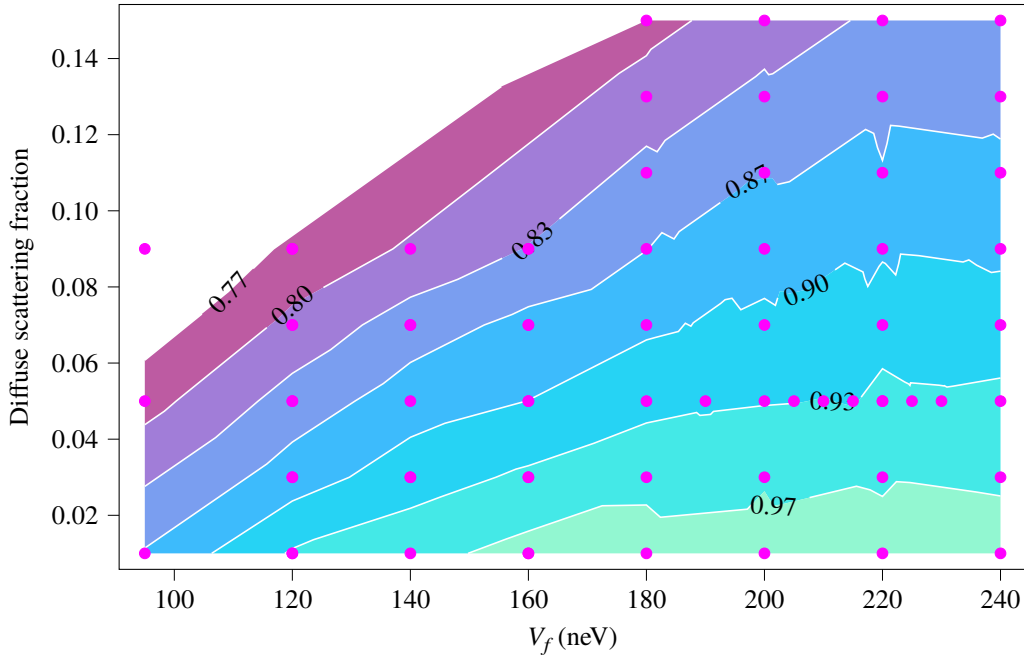
Figure 6.21: Transmission (12 to 200 s) as a function of Fermi potential and diffuse scattering fraction. $\eta = 3 \times 10^{-4}, l = 615$ mm. The magenta markers show the simulated parameters. The contours are constructed using cubic interpolation.

To explain a specific transmission of 0.84, η would have to be around twenty times higher than the default value for NiMo-coated guides in MCUCN, which has been experimentally validated [101], [117]. Such a large deviation is implausible and can be excluded.

The transmission depends strongly on the diffuse scattering fraction over its entire domain, including around the default value, $d = 0.05$. Almost perfect transmission is obtained for $d = 0$. The observed variation in the specific transmission (0.842 to 0.955, see Table 6.2) can be explained by varying d from 0.02 to 0.09.

6.8.3 Variation of diffuse scattering fraction and Fermi potential

To constrain the parameter space consisting of diffuse scattering fraction and Fermi potential to the region that is likely to represent the surface parameters of our NiMo, we expanded the surface parameters of the simulated guide in a grid using $L = 615$ mm, $V_F \in [95, 240]$ and $d \in [0.01, 015]$. The transmission from 12 to 200 s, as a function of V_F and d, is shown in Figure 6.21.

The effects of V_F and d on the simulated UCN transmission are not independent. For $V_F > 220$, the transmission does not increase as a function of V_F. This is expected because several metres of NiMo-coated UCN guides in the UCN source remove almost all UCN with energies above 220 neV. For lower V_F, the dependence of the transmission on d increases. UCN with energies above V_F can be reflected at small grazing angles. With a larger diffuse scattering fraction, there is more mixing of the momentum components, making collisions at angles close to the surface normal more likely. The UCN spectrum is thus cleaned more efficiently of UCN with energies above V_F when d is higher.

We constrained the parameter space to the region that best represents the surface parameters of the measured UCN guides. First, we used the observed distribution of specific transmissions (see Table 6.2) to calculate the expectation value (0.930) and standard deviation (0.027) for the transmission of a 615 mm guide

(a)   **(b)**
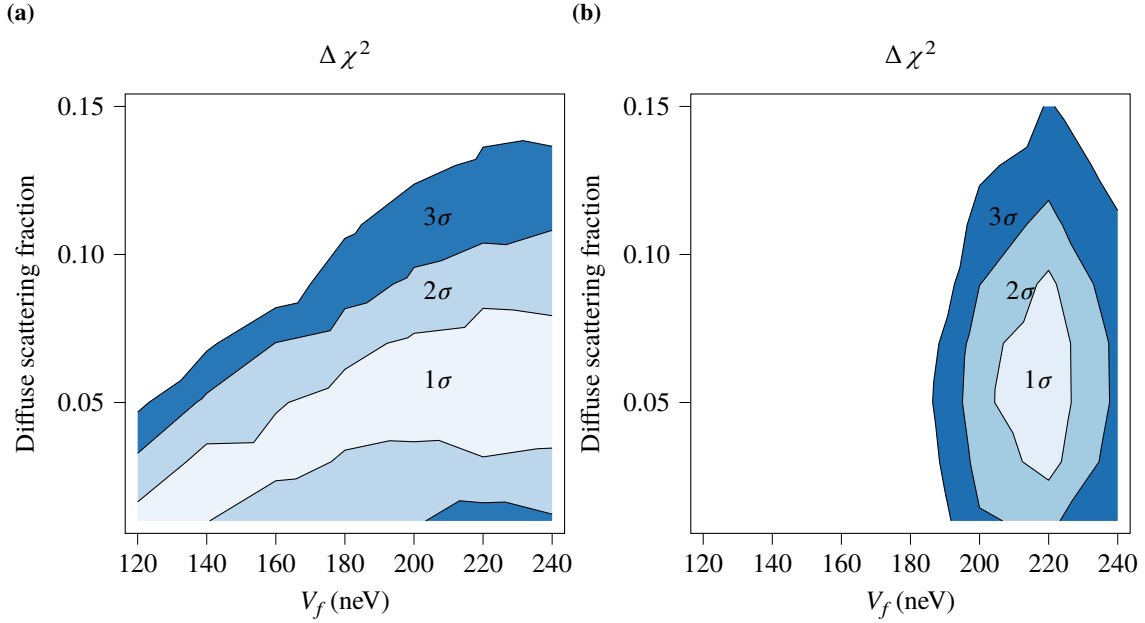
Figure 6.22: (a) A $\Delta\chi^2$ plot in (V_F, d) space, using the results of the transmission measurements to constrain the parameter space. The transmission varies less than 1 standard deviation from the experimental mean in the 1σ region. (b) $\Delta\chi^2$ is calculated based on the experimentally determined means for both transmission and Fermi potential. The 1σ region represents a 68 % confidence interval, based on the empirical distributions for V_f and transmission.

(the length used in the simulations). Defining

$$\chi^2(V_F, d) = \left(\frac{T(V_f F, d) - \mu_T}{\sigma_T} \right)^2, \tag{6.18}$$

we can partition the (V_F, d)-space into regions, based on the likelihood that a guide with surface parameters V_F and d has a transmission in the observed range (see Figure 6.22a).

The Fermi potentials of the guides were measured to be (216 ± 9) neV (mean $\pm$ standard deviation) using reflectometry (Section 6.7). We can use this to apply further constraints:

$$\chi^2(V_F, d) = \left(\frac{T(V_F, d) - \mu_T}{\sigma_T} \right)^2 + \left(\frac{V_F - \mu_{V_F}}{\sigma_{V_F}} \right)^2. \tag{6.19}$$

We again partition the parameters space, based on the likelihood that a guide has both a transmission and a Fermi potential within the observed ranges (see Figure 6.22b).

Based on transmission measurements, MCUCN simulations, and cold neutron reflectometry of witness samples, we conclude that the best explanation for differences in UCN transmission is that some guides on average scatter UCN more diffusely than others.

6.9   Testing the rough end hypothesis in simulation

To test whether the observed roughness or waviness in the first and last 50 mm of the UCN guides can adequately explain the variation in UCN transmission, we simulated transmission through a 10 cm guide, varying the diffuse scattering fraction from 0 to 1. The results are shown in Figure 6.23.

(a)

(b)
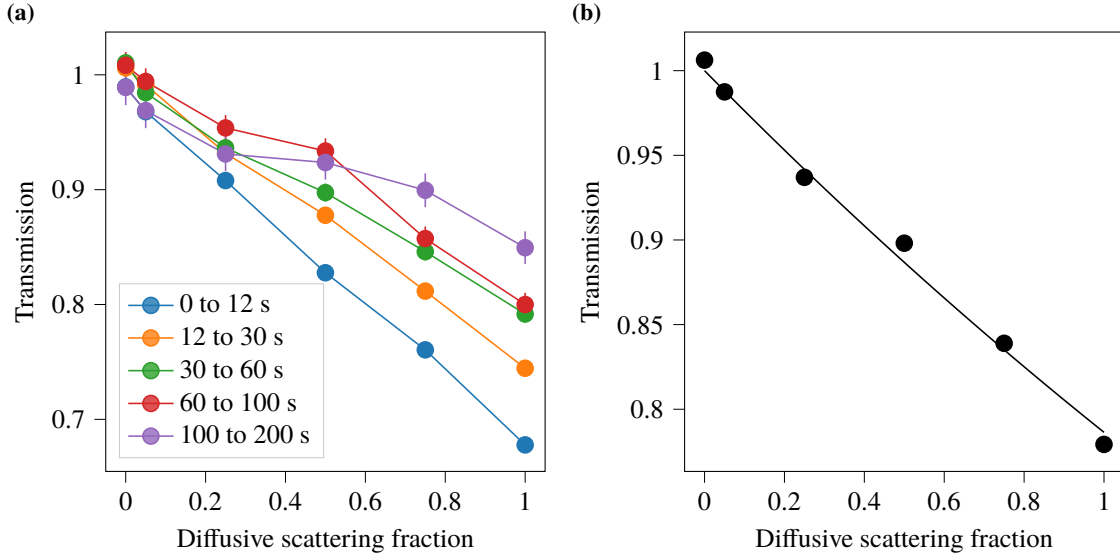
Figure 6.23: Simulated transmission through a 10 cm section of guide, as a function of the diffuse scattering fraction. (a) The simulated transmission, broken down by time interval. (b) The transmission from 12 to 200 s, with an exponential fit.

A simple exponential curve of the form

$$T(d; \tilde{T}) = \tilde{T}^{-d} \tag{6.20}$$

was fitted to the data. The fit parameter $\tilde{T}$ represents the transmission for perfect diffuse reflection. For UCN counted between 12 and 200 s after the proton beam pulse, $T_d = (0.786 \pm 0.005)$ (statistical uncertainty only). This is lower than any measured value.

Under the hypothesis that the central section of all guides has the same surface parameters (and thus specific transmission), and that variation in transmission is caused by variation in the roughness of the ends, we can use the fit to estimate the diffuse scattering fraction in the end of the guides necessary to explain the transmission.

For a guide with transmission T and length l, the specific transmission in the middle segment (without the last 50 mm on either end) is assumed to be the highest observed specific transmission, $T_0 = 0.955$. Thus, we estimate the transmission of the middle segment to be

$$T_e = T_0^{\frac{l-10\,\text{mm}}{1\,\text{m}}}. \tag{6.21}$$

Assuming multiplicativity of transmission, the transmission in the 50 mm on both ends is

$$T_e = \frac{T}{T_e}. \tag{6.22}$$

Using the fit from Equation 6.20, the diffusive scattering fraction of the end parts is

$$d_e = \frac{\log T_e}{\log A}. \tag{6.23}$$

For all measured guides, the result is shown in Table 6.6. This shows that a variation in diffusive scattering fraction in the ends of the guides from 0.05 to 0.33 can completely explain the variation in UCN transmission between the guides.

Table 6.6: Estimated diffuse scattering fraction in the ends of the UCN guides

Name	Length	Transmission	Diffusive scattering fraction in ends
HG1-21	0.625	0.964	0.052 ± 0.024
HG1-3	0.625	0.958	0.078 ± 0.024
HG1Ti	0.625	0.927	0.214 ± 0.025
HG3	1.250	0.944	0.019 ± 0.033
HG3-2	1.250	0.917	0.140 ± 0.033
HG4-1	0.595	0.903	0.329 ± 0.026
HG4-2	0.595	0.912	0.288 ± 0.026
HG4-3	0.595	0.922	0.243 ± 0.025
SB45-1	0.450	0.969	0.064 ± 0.022
U29.8-1	0.298	0.978	0.055 ± 0.021
U30-1	0.300	0.977	0.058 ± 0.021
U40-1	0.400	0.966	0.086 ± 0.022
U40-2	0.400	0.966	0.086 ± 0.022

6.10   Discussion

We measured the UCN transmission of most guides that were built into n2EDM. The specific transmission of these guides varied from 0.842 to 0.955. The large variation in transmission, as well as the number of guides with poor transmission, were surprising.

The obtained transmission values were reproducible within a standard deviation of 0.007. The uncertainty due to imperfect mounting for a single joint was estimated to be 0.004, by simulating gaps between guides and misalignment of guides using aluminium and steel rings. This effect is contained in the empirically determined reproducibility. The variation in the obtained transmission values is much larger than the uncertainty in the measurements. The variation in transmission reflects differences in the surface properties of the UCN guides.

By visual inspection, we determined the uncoated surface area in the guides to be smaller than $5\,\text{cm}^2$. This cannot explain a difference in transmission greater than 0.008, according to simulations and transmission measurements.

We determined that the Fermi potential of all guides is larger than $200\,\text{neV}$, by cold neutron reflectometry. UCN transmission of a $615\,\text{mm}$ long guide does not significantly depend on the Fermi Potential in the range $V_F \in [200, 240]\,\text{neV}$, according to MCUCN simulations. Variation of V_F is thus also ruled out as an explanation for the differences in transmission.

We also found that the UCN transmission was not sensitive to changes in the loss coefficient, up to a value ten times higher than the value experimentally validated for the UCN guides in the UCN source.

The remaining hypothesis is that the differences in transmission can be explained by differences in roughness of the UCN guides. This could be either intrinsic to the glass guides or the coating. We did not find that the coating process does not measurably increase the surface roughness using profilometry.

The guides with lower transmissions were made using a manufacturing process that can create rough sections at the ends of the guides. By varying the diffuse scattering parameter in simulation, we found that roughness in the ends of the guides can explain the observed differences in transmission.

We coated two UCN guides, produced by GlasForm and Schmizo, together, so that the coating process was equal for both guides. We found that the guide manufactured by Schmizo, which was visually smoother, had a higher transmission.

Based on these results, we have replaced the UCN guides in n2EDM with UCN guides based on glass

tubes from Schmizo, with a smaller surface waviness.

CHAPTER 7

DETECTION OF MAGNETIC DIPOLES IN THE UCN GUIDES

7.1 Overview

Local inhomogeneities in the magnetic field can offset the readings of the n2EDM caesium magnetometer array. This leads to an incorrect reconstruction of the field gradients, $G_{l,m}$, potentially causing a false EDM effect. The n2EDM systematic error budget imposes stringent limits on the remaining magnetic dipoles of any component installed inside the MSR [24], [72]. These limits depend on the position of the component; the closer it is to the ground electrode, the stricter the limits become. The lowest tolerance for magnetic dipoles is in the vicinity of the caesium magnetometer cells.

We scanned for magnetic dipoles in all components that are built into the MSR and checked whether their magnitudes are below the limit. This was done using the magnetic gradiometer, developed and built by our group [118], [119]. Components are also measured after degaussing to know the dipole strength when the components are built in. There is a risk that a spark in the high voltage system of n2EDM magnetises components installed in the experiment. To ensure the components cannot be magnetised to a problematic degree, we also measure them after magnetising them with a hand held, permanent 30 mT magnet.

Of all the glass UCN guides in the n2EDM experiment, the 595 mm long guides (labelled HG4) are the closest to the caesium magnetometer array and the precession chambers. D. Pais determined that the maximum allowed dipole strength in these guides is $|\mathbf{m}| = 15\,\mathrm{nA\,m^2}$ [72]. We produced five guides of this type. Using the methods described in this chapter, we found that four of these satisfy the imposed limit.

The production of UCN guides involves coating glass tubes with nickel-molybdenum. This exposes the tubes to a relatively strong magnetic field of several milliteslas, and adds new material. To determine whether the coating process introduces new magnetisable domains and whether the magnetic fields in the coating machine are strong enough to magnetise the existing magnetisable areas, we measured a UCN guide before and after coating. We found that while coating introduces new magnetisable areas, it does not magnetise the existing or newly created magnetisable areas to a measurable degree.

For comparison, a non-magnetic stainless steel guide was measured. The dipole strength of the guide is five orders of magnitude higher than what we measured in the good NiMo-coated glass guides.
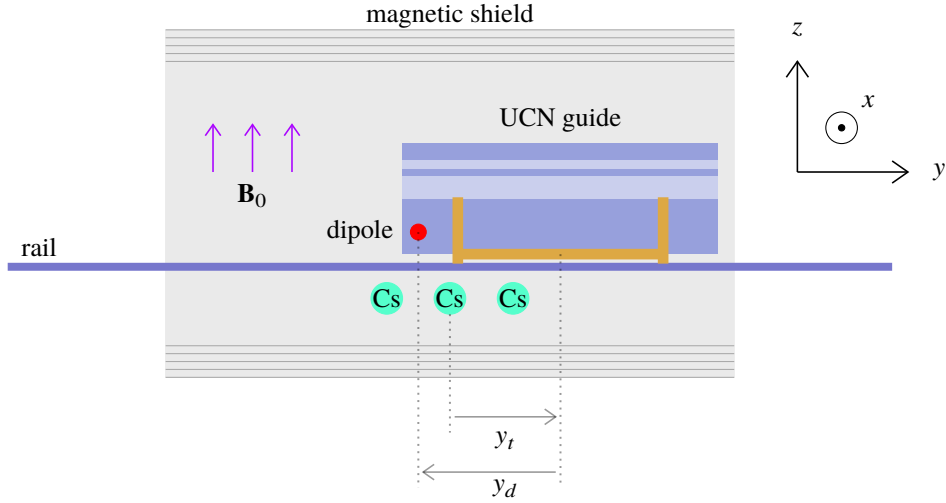
Figure 7.1: A sketch of the setup for measuring UCN guides in the gradiometer, including the coordinate system (not to scale). The caesium magnetometer cells are indicated by 'Cs'. The displacement of the centre of the trolley (and the UCN guide) from the centre of the middle caesium cell along the y-axis is indicated by y_t. The longitudinal position of a dipole relative to the centre of the guide is indicated by y_d. The saddle coils that generate the $\mathbf{B}_0$ field are not shown.

7.2   Setup and methods

7.2.1   The magnetic gradiometer

Our group has developed a magnetic gradiometer, an apparatus which can detect magnetic dipoles in samples. The development and commissioning were led by V. Kletzl, who describes the apparatus in detail in her PhD thesis [118]. We use this apparatus to ensure that the sum of all the dipoles we build into the experiment does not exceed the limit derived from our budget for systematic effects. Samples are driven through the gradiometer on a plexiglass trolley. An array of three stationary caesium vapour magnetometer cells (Cs magnetometers) [120] records the magnetic field as the sample passes over them. The trolley and the Cs magnetometer array are enclosed in a five-layer mu-metal shield. Inside the shield, saddle coils generate a vertical magnetic field, $\mathbf{B}_0$, with a field strength of 2.6 μT in a cylindrical volume (l = 1950 mm, ⌀ = 500 mm).[1] A sketch of the gradiometer is shown in Figure 7.1.

The quantity of interest is the difference between the magnetic fields measured between the first and the third Cs magnetometers, which are spaced 172 mm apart. The Cs magnetometers measure the absolute magnetic field strength. This is the sum of the background field, the $\mathbf{B}_0$-field, and the field generated by the sample. A simulated signal produced by a dipole aligned with $\mathbf{B}_0$ is shown in Figure 7.2.

The caesium magnetometers are not sensitive to horizontal magnetic field components, because the absolute magnetic field at their positions is dominated by the vertically oriented $\mathbf{B}_0$. A horizontal field component of 1 nT (which is large in this context) would change the absolute field by approximately 0.2 pT, which is around the noise level of the gradiometer. This allows us to simplify the analysis by assuming that the absolute field strength is equal to the vertical field component. This allows us to add and subtract absolute field strengths, for example, to perform background subtraction. This is required because the trolley itself is not entirely non-magnetic, and produces a signal when it is driven over the Cs magnetometers. Therefore, in conjunction with every set of measurements, a background measurement is done, where the signal of the empty trolley is recorded without a sample. This background signal is subtracted from sample measurements.

---

[1]For reference, the geomagnetic field in Switzerland is approximately 47 μT [61].
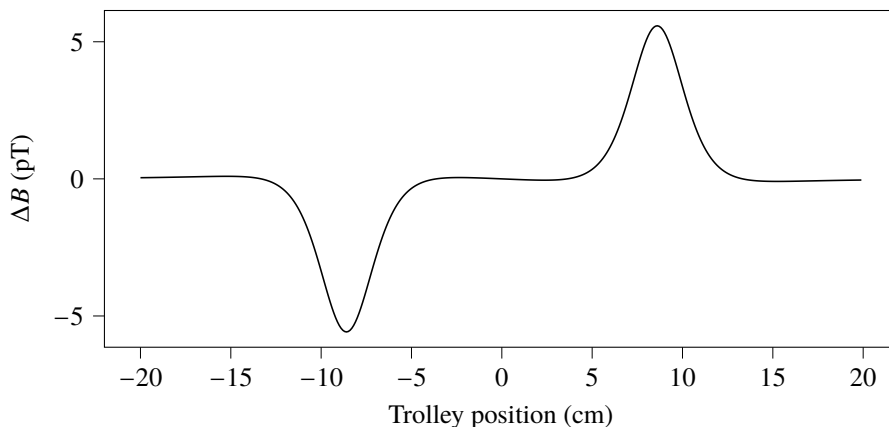
Figure 7.2: A simulated signal in the gradiometer, caused by a $1\,\mathrm{nA\,m^2}$ dipole pointing downwards, placed in the centre of the trolley at a height of 5 mm above the surface of the trolley. The dipole passes over the centre of the Cs magnetometers with a distance of 33 mm. The signal, $\Delta|B| \approx B_{z,2} - B_{z,1}$, becomes negative as the dipole is driven over the first caesium cell, and positive as it is driven over the second cell.

Although all components inside the MSR must be degaussed before installation, currents created by sparks between the high voltage electrode and grounded components could magnetise components. To estimate how strongly components can be magnetised, samples are normally magnetised using a 30 mT permanent magnet. The strength of this magnet was chosen such that it is significantly larger than any field in the experiment while being smaller than the oscillating field produced by the demagnetiser that we use to degauss components. The sample is then placed on the centre of the trolley, with the magnetised side facing down, aligning the dipoles with $\mathbf{B}_0$ for the highest sensitivity. After this measurement, components are degaussed and measured again.

7.2.2   The wooden UCN guide mount

We built a mounting structure to securely place UCN guides on the gradiometry trolley, shown in Figure 7.3. The structure consists of two rectangular surfaces of multiplex wood, with a semicircle of 144 mm cut out (slightly larger than the outer diameter of the UCN guides, 140 mm), joined together by four wooden rods. The rods are glued to the surfaces using wood glue (polyvinyl acetate). The wood glue was not measured in the gradiometer, but a minimal amount was used. The distance between the surfaces (350 mm) is such that they gently pinch the plexiglass surface of the trolley, so that the structure cannot slide as the trolley moves. When a UCN guide is mounted on the structure, its lowest point is 3 mm above the trolley surface. To limit the background signal introduced by the wooden UCN guide mount, it was designed so that the amount of material close to the longitudinal axis of the trolley is small.

To make the wooden mounting structure as non-magnetic as possible, all the components were measured in the gradiometer. They were sanded down until the signal strength, defined as $\max(\Delta B) - \min(\Delta B)$, was below 35 pT for all individual parts. The assembled structure causes a smaller signal than the individual parts, since all material is further away from the Cs magnetometers.

The parts were degaussed. After degaussing, all the parts were remeasured. This time, the signal of all parts was indistinguishable from the background. The wooden mounting structure was glued together. Afterwards, it was verified that the signal of the assembled structure is still indistinguishable from the background.

Figure 7.3: The wooden UCN guide mount, placed on top of the gradiometer trolley.

7.2.3   Procedure for measuring UCN guides

Background measurement

The wooden mounting structure is placed on the trolley, without a UCN guide. It is measured with at least ten repetitions. This measurement will henceforward be referred to as the background measurement.

UCN guide positioning

The UCN guides are relatively large objects. The magnetic field strength of dipoles decays with the third power of distance. Care has to be taken that every point of the guide passes within a small enough distance of the Cs magnetometers. The guide is measured multiple times. After every measurement, the guide is rotated by 30°, so that a new section of the guide is facing downwards.

To do this consistently, we define a coordinate system on the guide in the following way. We always position the guide with its centre above the centre of the trolley. We choose which end of the guide enters the gradiometer first, and which one last. We mark the end that enters last (the end that is pointing away from the gradiometer when the trolley is in the parking position) with a piece of tape, on which a line is drawn. When the line on the tape points upwards, the guide is said to be in position 12. If one attaches a clock face to the marked side of the guide, the angular position corresponds to the hour pointing upwards (see Figure 7.4).

The guides were scanned in twelve positions, rotated by 30° between scans, with six to eight repetitions per position.

Magnetisation procedure

To put an upper limit on the dipole strengths that can be reached if a UCN guide is magnetised, we magnetise it by rubbing the 30 mT hand-held permanent magnet over the entire outside surface of the guide. The UCN guide is measured in twelve different angular orientations. Every time it is measured at a new angle, we take it off the trolley, and magnetise the side that will face down, by making long, slow strokes with the permanent magnet.

Figure 7.4: A UCN guide on the gradiometer trolley. A clock face is attached to the marked side (which enters the gradiometer last), defining the angular positions. The position shown in this figure is position 8. Since the paper clock face produces a magnetic signal, it was removed before the measurement.

Degaussing procedure

All n2EDM components that should go into the MSR are degaussed before installation. For small components, this is typically done using a Wallmag Dem-4 Easy Logic Self Adjusting Demagnetizer[2], a type of plate demagnetiser. Objects are moved back and forth, from left to right, above the surface of the demagnetiser, with increasing distance to the surface.

7.3 Analysis

7.3.1 Data preparation

The gradiometer records the magnetic field strength measured by the Cs magnetometers and the position of the trolley with 10 ms intervals. It stores these values in a CSV file, along with the timestamp. Since samples are usually measured multiple times, the repetition number is stored in an additional column. For the purpose of detecting dipoles, the quantity of interest is the magnetic field difference between the first and the third cell, $\Delta\mathbf{B}$.

Data cleaning

Sometimes, glitches occur in the communication between the computer and the gradiometer. In such cases, one data point is generated with corrupted values for the position and magnetic field values. These points are detected using the fact that their position is not between the positions of the previous and next point, and are removed from the data. This also discards the first and last points in one repetition, which does not affect the measurement. This does not remove the magnetic field minima and maxima.

Magnetic drift correction

Ramping of the magnet in the neighbouring superconductor test facility (SULTAN [71]) causes slow, linear drifts of the external magnetic field, changing the background contribution to the field gradient inside the

---

[2]https://www.walmagmagnetics.com/

gradiometer. To correct for magnetic field drifts, for every repetition, a line is fitted to the $\Delta \mathbf{B}$ points outside the region of interest (defined by the length of the sample), and subtracted from $\Delta \mathbf{B}$.

Binning

The speed of the trolley is not constant. The trolley accelerates at the beginning and decelerates at the end of a repetition, causing unequal steps in trolley position. To facilitate the comparison of different files, the data points of all repetitions in one measurement are averaged into 0.5 cm bins.

Background subtraction

We take a new background measurement every time the UCN guide mount is placed on the trolley. We measure the signal produced by the trolley with the UCN guide mount without a guide. The resulting data are treated in the same way as the data from the sample measurements and then subtracted from the sample measurements.

Magnetic field drift correction is done on individual measurements. It is not necessary to measure the background before and after a sample measurement.

7.3.2   Dipole fitting

Fit function

In the gradiometer coordinate system, the z-axis is vertical, the y-axis is along the direction of movement of the trolley, and the x-axis is perpendicular to both. The origin is defined as the centre of the central Cs magnetometer. The position of the trolley is at $y_t = 0$ when the centre of the trolley surface is on the z-axis, directly above the central Cs magnetometer.

The distance from a dipole in the sample to one of the Cs magnetometers can be written as:

$$\mathbf{r}_{1,2}(\mathbf{r}_t, \mathbf{r}_d) = \mathbf{r}_d + \mathbf{r}_t - \mathbf{r}_{\text{Cs}1,2}, \tag{7.1}$$

where $\mathbf{r}_d$ is the position of the dipole when the trolley is in the null position, $\mathbf{r}_t$ is the displacement of the trolley from the null position, and $\mathbf{r}_{\text{Cs}1,2}$ are the caesium cell positions. Since the trolley only moves along the y-axis, and the Cs magnetometers are on the y-axis with opposite y-coordinates, this can be rewritten as

$$\mathbf{r}_{1,2}(y_t, \mathbf{r}_d) = \mathbf{r}_d + \hat{y}\left(y_t \mp y_{\text{Cs}}\right), \tag{7.2}$$

where $\hat{y}$ is the unit vector in the y direction, y_t is the trolley displacement and y_{Cs} is the semidistance between the caesium cell centres (86 mm).

The vertical component of the magnetic field measured by the Cs magnetometers can be written as

$$B_{z1,2}(y_t, \mathbf{r}_d, \mathbf{m}) = \frac{\varepsilon_0}{4\pi}\left(\frac{3z_d(\mathbf{r}_{1,2} \cdot \mathbf{m})}{|\mathbf{r}_{1,2}|^5} - \frac{m_z}{|\mathbf{r}_{1,2}|^3}\right), \tag{7.3}$$

where ε_0 is the vacuum permittivity, z_d is the z-component of $\mathbf{r}_d$, $\mathbf{m}$ is the magnetic dipole moment and m_z is the vertical component of the magnetic dipole moment [118].

The fit function for the signal is then,

$$\Delta B(y_t, \mathbf{r}_d, \mathbf{m}) = B_{z2}(y_t, \mathbf{r}_d, \mathbf{m}) - B_{z1}(y_t, \mathbf{r}_d, \mathbf{m}). \tag{7.4}$$

When fitting the signal, we treat y_t as the independent variable, and $\mathbf{r}_d$ and $\mathbf{m}$ as the fit parameters.

Table 7.1: Bounds for dipole fits during the measurements of UCN guides in the gradiometer.

Parameter	Lower bound	Upper bound
$|\mathbf{m}|$	0	$1\,\mu\mathrm{A\,m^2}$
$\theta_{d,1}$	0	$360°$
$\theta_{d,2}$	0	$360°$
$\phi_{d,1}$	0	$360°$
y_d	$-\frac{1}{2}$ guide length	$+\frac{1}{2}$ guide length
Δy_d	$-12\,\mathrm{mm}$	$+12\,\mathrm{mm}$

From Equation 7.3, one can see that m_x only appears in the inner product $\mathbf{r}_{1,2} \cdot \mathbf{m}$. That is, the shape of the signal does not depend on m_x at all if the dipole is in the lateral centre. For small displacements, the effect of m_x on the fit shape is limited and can be absorbed into m_z. If we include m_x as an independent variable during fitting, the fit algorithm has difficulties with estimating the fit parameters and the uncertainties. Normally we omit m_x.

Fit coordinates

The primary goal of the fit is to find the strength of dipoles in the guide, the secondary goal is to determine the location. The signal strength depends on both. If the location is misdetermined, so is the dipole moment. Therefore, the location parameters have to be constrained to the physical extent of the guide.

Since the guide is a cylindrical object, it is more natural to give bounds on dipole positions in cylindrical coordinates. We will describe the position of dipoles with the coordinates (R_d, ϕ_d, y_d). The conversion back to Cartesian coordinates is given by

$$\begin{pmatrix} x_d \\ y_d \\ z_d \end{pmatrix} = \begin{pmatrix} R_d \sin(\phi_d) \\ y_d \\ z_c + R_d \cos(\phi_d) \end{pmatrix},$$ (7.5)

where z_c is the height of the centre of the guide above the Cs magnetometers (108 mm). These coordinates are relative to the UCN guide. See Figure 7.1.

Fitting two curves simultaneously

Relatively strong dipoles are visible in scans of multiple consecutive positions. By fitting the two scans in which the signal strength from one dipole is the strongest simultaneously, one can put tighter constraints on ϕ_d. This reduces the error in the estimate of the dipole moment. We use that $|\mathbf{m}|, R_d$ and y_d are the same in both scans. Between the scans, the tube was rotated by 30 degrees, therefore $\phi_{d,2} = 30° - \phi_{d,1}$.

Since between the scans, the guide was magnetised again, the direction of the dipole moment can change. Therefore, the angle of $\mathbf{m}$ in the y-z plane, θ_d, is fitted separately for both dipoles. In addition, small errors in the longitudinal placement of the guides can have a large effect on the fit, therefore Δy_d is added as a nuisance parameter.

All fit parameters, and their bounds, are summarised in Table 7.1.

7.4 The horizontal UCN guides closest to the precession chambers

All guides that we produced of the type HG-4 (length: 595 mm, outer diameter: 140 mm) which are installed closest to the stack, were measured in the gradiometer using the procedure described earlier in this chapter.

Table 7.2: The strongest dipoles found by scanning the HG4 guides in the gradiometer.

UCN guide	Status	$|\mathbf{m}|$ (nA m^2)	Angular position	Horizontal dipole position y_d (cm)
HG4-1	magnetised	5.9 ± 0.5	2	-27
		2.9 ± 0.3	9.5	16
	degaussed	-	-	-
HG4-2	magnetised	120 ± 15	6	24
		5.5 ± 0.5	2	-28
	degaussed	1.2 ± 0.1	6	24
HG4-3	magnetised	5.4 ± 0.2	4	-28
		2.7 ± 0.4	7	23
	degaussed	-	-	-
HG4-5	magnetised	2.0 ± 0.2	11.5	25
		0.7 ± 0.1	2	28
HG4-6	magnetised	1.7 ± 0.2	7.5	22
		1.6 ± 0.3	1	-28

When we measured a guide for the first time, we magnetised the bottom of the guide using the handheld 30 mT magnet before measuring in every angular position. The guides were degaussed before the second series of measurements. For guides HG4-5 and HG4-6, the dipoles were so weak even when magnetised that we did not remeasure them after degaussing.

Plots of the scans of HG4-1, HG4-2 and HG4-3 are shown in Figure 7.5. The strongest two dipoles found in every guide are shown in Table 7.2. In HG4-1 and HG4-3, the worst dipoles are 5.9 and 2.9 nA m^2, respectively, well below the limit of 15 nA m^2 [72]. In both HG4-1 and HG4-3, the worst dipoles are on the side of the guide where the marking was made. This is after magnetisation with the 30 mT magnet, which serves as a worst case scenario. HG4-2 exceeds this limit by an order of magnitude when magnetised, the strongest dipole is 120 nA m^2.

When the guides were degaussed, no dipoles could be fitted above 0.5 pT. The exception is the very strong dipole in HG4-2. It was still visible after degaussing, with a net dipole strength two orders of magnitude smaller.

The UCN guides HG4-1, HG4-2 and HG4-3 were made from glass tubes provided by GlasForm. As described in Chapter 6, it was decided to replace these guides, as they contain wrinkles on the inner surface, reducing their UCN transmission. We made two new UCN guides from glass tubes provided by Schmizo, HG4-5 and HG4-6. These two new guides were first measured magnetised. The signal was significantly lower than from the first three guides. Since in guides with stronger dipoles (HG4-1 and HG4-3), no dipole could be observed or fitted after degaussing, we decided it was not necessary to measure HG4-5 and HG4-6 after degaussing.

7.5   The effect of coating on the presence of dipoles

SB45-2 is the straight guide for the second switchbox (i.e. the guide that transports UCN from the SCM to the precession chambers). This was the first guide we coated after the gradiometer had been commissioned, it was an opportunity to measure a guide before and after coating. From these measurements, we can learn whether the ring magnet in the sputter head of the coating machine is strong enough to magnetise the guides and whether the coating process itself can introduce magnetisable particles.

In total, the guide was measured four times, using the same procedure described in section 7.2.3. Every
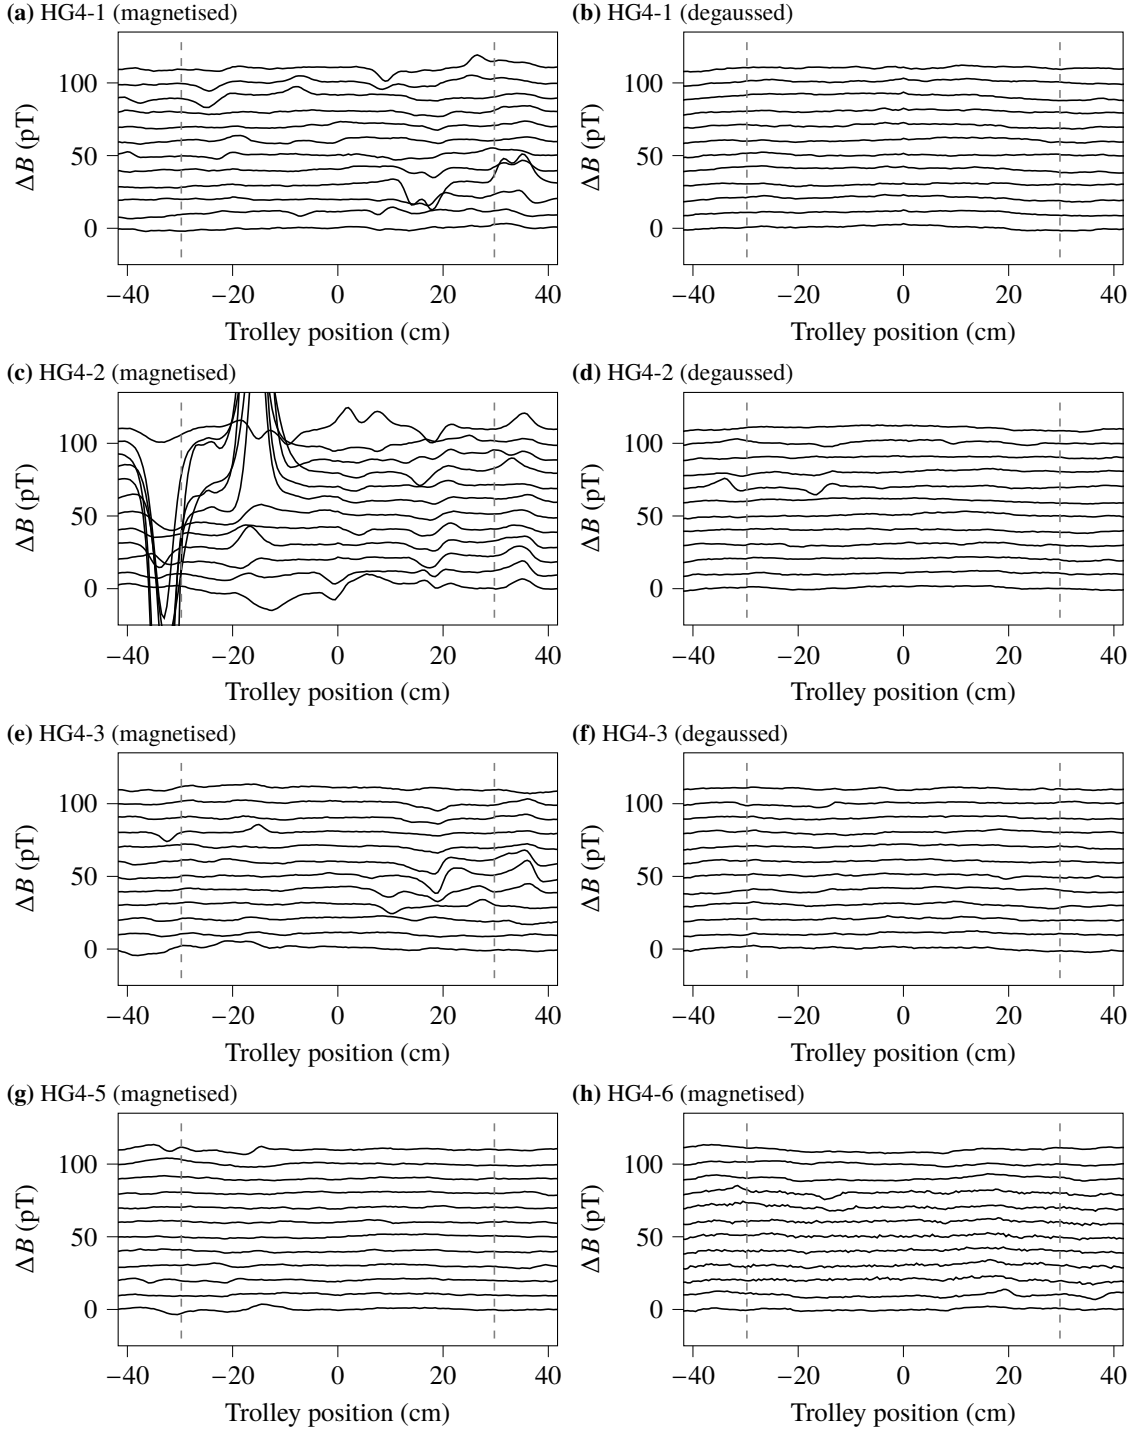
Figure 7.5: Gradiometer scans of the horizontal UCN guides closest to the precession chamber, HG4. The guides were measured using the magnetisation procedure, degaussed, and remeasured. The background was subtracted. The guides were measured at 12 angular positions. The lines are offset by 10 pT, with Position 0 (12) at the bottom and Position 11 at the top. The horizontal axis represents the trolley position y_t relative to the centre of the caesium magnetometers. A negative y_t (on the left) corresponds to the end of the guide that enters the gradiometer first. At the vertical, dotted lines, the end of the guide crosses the centre of the setup. A dipole at position y_d produces a peak and a valley at $y_t = -y_d \pm 86$ mm (Figure 7.1). The scale in (c) is the same as in the other plots, the difference between the peak and the valley of the scan is 1.1 nT.

time it was measured, all twelve angles were scanned. The procedure was as follows:

- The guide magnetised using the 30 mT magnet and measured in the gradiometer.

- The guide was degaussed using the procedure described in section 7.2.3, and measured again.

- The guide was coated. After coating, the guide was measured in the gradiometer without magnetising or degaussing it. This was to learn whether coating magnetises the guide.

- The guide was measured in the gradiometer, it was again magnetised before every measurement using the 30 mT magnet. This was to learn whether coating introduces new magnetisable areas.

7.5.1   Magnetic field strength in coating machine

A Hall sensor-based Lakeshore Model 410 gaussmeter[3] was attached to an uncoated guide, which was placed on the trolley inside the coating machine. The trolley was moved to scan the field components of the ring magnet in the sputter head inside the coating machine. The radial and longitudinal fields were measured. By argument of symmetry, there cannot be an azimuthal field. During coating, the electric discharge current produces additional magnetic fields. The measured field of the ring magnet is thus a minimum estimate. The measured fields are shown in Figure 7.6.

Based on this measurement, UCN guides are exposed to a magnetic field of at least (2.6 ± 0.2) mT during coating.

7.5.2   Results

There were small but detectable dipoles present before coating (Figure 7.7a). The fitted dipole moments were (2.8 ± 0.2) nA m^2, and (1.9 ± 0.2) nA m^2 which could be completely eliminated by degaussing (Figure 7.7b).

After coating, no dipoles were visible in the scan (Figure 7.7c). There are some oscillations in the signal to which a dipole moment of (0.56 ± 0.09) nA m^2 can be fitted.

By magnetising the guide with the 30 mT magnet, dipoles larger than those visible before coating were revealed, with dipole strengths of (5.5 ± 0.2) nA m^2 and (3.3 ± 0.3) nA m^2 7.7d).

The measurements of the SB45-2 guide, before and after coating, show that while some dipoles were inherent to the glass, up to (2.8 ± 0.2) nA m^2, the coating process introduced new magnetisable areas in a UCN guide, up to (5.48 ± 0.14) nA m^2 in this case.

During coating, the 2.6 mT magnetic field that the guides are exposed to for hours is insufficient to magnetise these areas. This is reassuring. The magnetic fields inside the MSR which will be below 10 μT, or the 50 μT Earth magnetic field, are not strong enough to magnetise the guides to a problematic degree.

It is more difficult to judge whether sparks, between the HV electrode and grounded components, can magnetise the guides.

7.6   Measuring the field strength of a non-magnetic stainless steel guide

Previous nEDM experiments have used non-magnetic stainless steel UCN guides. To put the quality of our NiMo-coated glass guides into perspective, I have measured a non-magnetic stainless steel guide in the gradiometer. The guide was provided by Nocado[4]. The steel alloy used was either 304 or 316L.

---

[3]https://www.lakeshore.com/products/categories/overview/discontinued-products/discontinued-products/model-410-hand-held-gaussmeter
[4]Nocado GmbH, Daimlerstraße 12, 27793 Wildeshausen, Germany, https://nocado.com/
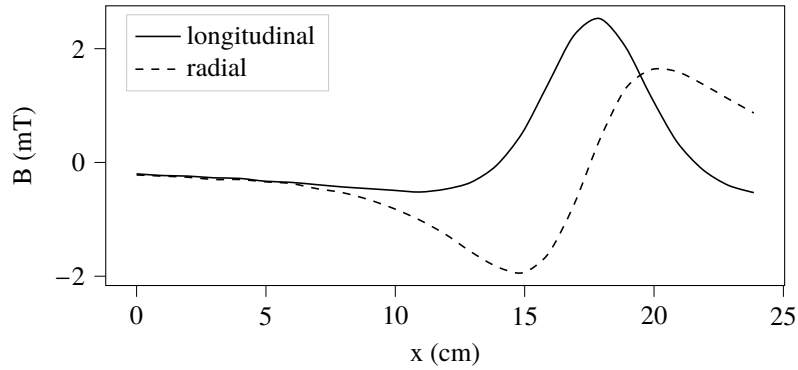
Figure 7.6: The magnetic field in the sputter coating machine, measured with a gaussmeter while the sputter coating machine was turned off. The strongest field the guide is exposed to is at least 2.6 mT.

(a) SB45-2 (magnetised)

(b) SB45-2 (degaussed)

(c) SB45-2 (coated)

(d) SB45-2 (magnetised after coating)
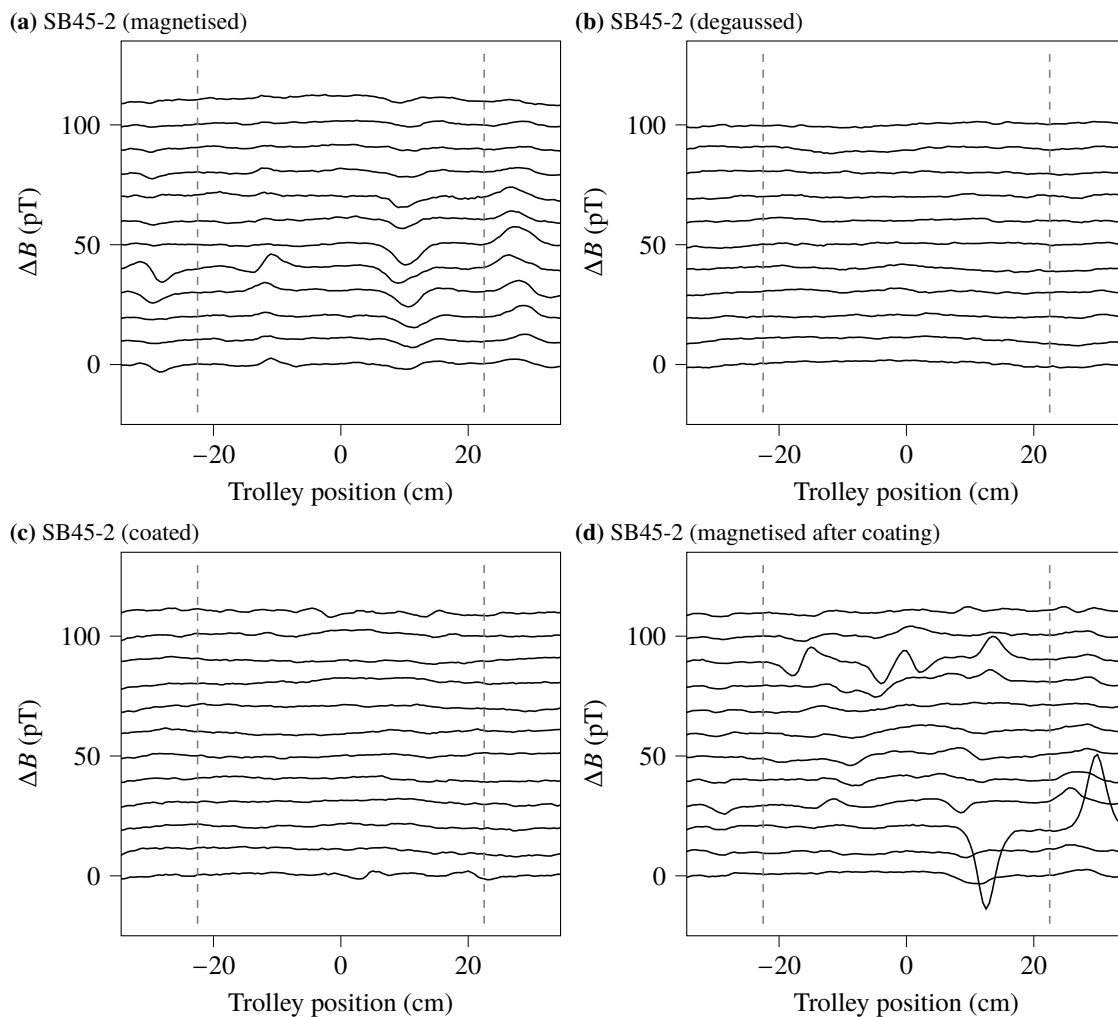
Figure 7.7: Gradiometer scans of the diagonal switchbox guide, SB45-2. (a) First, the guide was scanned using a normal magnetisation procedure. (b) The guide was degaussed and scanned again. Due to a software bug, the last line was not saved properly. (c) The guide was scanned directly after coating to see if the ring magnet had magnetised the guide. (d) The guide was measured using the magnetisation procedure to see if new magnetisable areas had been introduced.

The outer diameter of the guide is 76.2 mm, the wall thickness 1.60 mm, and the length 287 mm. This corresponds to about 108 cm^3 of steel.

7.6.1   Measurement and analysis

Degaussed steel guide   The guide was degaussed using the same procedure as for the glass guides. When the steel guide was scanned in the gradiometer, it saturated the caesium magnetometer cells. The distance between the guide and the cells had to be increased dramatically. This was accomplished by mounting it on an 11 cm thick cardboard box. First, the box was scanned without the guide, as a background measurement. The box contains several dipoles and has a peak-to-peak signal strength of 265 pT (the difference between the minimum and maximum of ΔB_z). For a normal measurement, this would be prohibitive.

The guide had a peak-to-peak signal of 5075 pT, after subtraction of the background. The guide is far removed from the Cs magnetometers; any attempt to resolve individual dipoles would amount to overfitting. In the analysis, I assumed the guide has a single dipole in its central axis. This simplification results in a poor fit, but it is the order of magnitude we are interested in. From the fit, we obtain a dipole strength of approximately 1×10^5 nA m^2.

Magnetised steel guide   After magnetising the steel guide with the 30 mT magnet, the guide was so magnetic that it was impossible to measure it, despite being elevated 11 cm above the trolley surface. The ΔB_z at the moment the Cs magnetometers were saturated was 20 nT, and it was not close to reaching the peak. Since the data are corrupted, we cannot fit a curve to the signal, but from signal strength at the time the gradiometer was saturated, we can set the lower limit of the dipole strength to 10^6 nA m^2.

Unsurprisingly, stainless steel is an unsuitable material for UCN guides in the n2EDM experiment. Whether degaussed or magnetised, stainless steel guides are at least 5 orders of magnitude more magnetic than our NiMo-coated glass guides.

7.7   Discussion

We have developed a method to measure the presence of magnetisable areas within the UCN guides, which are large objects and therefore relatively challenging to measure.

In four out of five horizontal UCN guides that we produced that should be the closest to the precession chambers, the largest dipoles were significantly smaller than 15 nA m^2, which was the limit imposed by the systematics budget [24] following the calculations of D. Pais [72]. When these guides were degaussed, no dipoles could be detected. The two UCN guides made from glass tubes from Schmizo had the smallest dipoles. The guide HG4-2 had significant magnetisable areas, and should not be used in n2EDM. However, when degaussed, the magnetisable area in this guide was also well below the limit.

We conclude that NiMo-coated borosilicate glass is an excellent choice of materials for the UCN guides in n2EDM, based on the magnetic characteristics.

We found that magnetisable areas can be present in the bulk of the glass, but can also be introduced into the guide by the coating process. The coating process itself exposes the guides to a magnetic field of at least (2.6 ± 0.2) mT. For the guide we measured after coating, the magnetic field in the coating machine was not strong enough to magnetise it to a measurable degree. We have not measured what effect this magnetic field would have on the strongly magnetisable area in HG4-2.

CHAPTER 8

OPTIMISATION OF THE PULSE SHAPES FOR N2EDM

8.1 Overview

We investigated the shifts in resonance frequency in the Ramsey cycle of n2EDM, which are caused by the cross-effect of the Hg $\frac{\pi}{2}$-pulse on the UCN spin state, as well as by a possible tilt angle between the static magnetic field and the rotating magnetic field.

In n2EDM, a vapour of ^{199}Hg is used as a co-magnetometer [64]. The $\frac{\pi}{2}$-pulse that starts the mercury precession can affect the initial UCN spin state. If the UCN spin axis is not aligned with $\mathbf{B}_0$ initially, the resonance frequency is shifted. I explain how we designed pulse shapes to minimise the $\frac{\pi}{2}$cross effects of the Hg $\frac{\pi}{2}$-pulse on the UCN spin state.

If the rotating field generated during the UCN $\frac{\pi}{2}$-pulse is not exactly horizontal, it contains an oscillating vertical field component that affects the resonance frequency. We designed a UCN pulse that is less sensitive to possible tilt of the rotating field, suppressing the frequency shift.

Using these optimised Hg and UCN $\frac{\pi}{2}$-pulses, the frequency shifts from these effects can be kept below 2×10^{-11} Hz. The frequency associated with the systematics budget for n2EDM [24] equals 4.4×10^{-9} Hz.

8.2 Mathematical formalism for spin evolution

First, we develop the tools required to study the effect of rotating transverse magnetic fields on spin systems precessing in a static magnetic field. I will present a quantum mechanical treatment of the evolution of the spin states under a transversal field pulse, adapted from [121].

Neutrons and ^{199}Hg nuclei are spin-$\frac{1}{2}$ fermions. Their spin is a two-level quantum mechanical system with two complex degrees of freedom.[1] Such a system can be represented by a spinor:

$$|\chi\rangle = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}. \tag{8.1}$$

---

[1]This equates to four real degrees of freedom. However, one degree of freedom is fixed by the normalisation, and one is the (unobservable) total phase. Since there are two observable degrees of freedom and the expected spin must evolve like the classically expected value (Ehrenfest's theorem [122]), the spin-$\frac{1}{2}$ system is isomorphic to a classical spinning object.

8.2.1   The Hamiltonian

The Hamiltonian operator of the system is:

$$\hat{H} = -\hat{\boldsymbol{\mu}} \cdot \mathbf{B} = -\gamma \; \hat{\mathbf{S}} \cdot \mathbf{B}, \tag{8.2}$$

where $\hat{\boldsymbol{\mu}}$ is the magnetic moment, γ is the gyromagnetic ratio and $\hat{\mathbf{S}}$ is the spin operator:

$$\hat{\mathbf{S}} = \frac{\hbar}{2} \begin{pmatrix} \sigma_x \\ \sigma_y \\ \sigma_z \end{pmatrix}. \tag{8.3}$$

Here σ_i are the Pauli matrices:

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{8.4}$$

This is a generic Hamiltonian for a spin in a magnetic field. In n2EDM, there are two main fields to consider: the static field $\mathbf{B}_0 = B_0 \hat{\mathbf{z}}$, and during the $\frac{\pi}{2}$-pulses, the circularly rotating transversal field

$$\mathbf{B}_c(t) = B_c \begin{pmatrix} \cos(\omega t + \phi) \\ \sin(\omega t + \phi) \\ 0 \end{pmatrix}. \tag{8.5}$$

Since the electric field is parallel with the static magnetic field, $\mathbf{B}_0$ can be thought of as an effective magnetic field, such that $-\gamma \sigma_z B_0$ includes the contribution of the electric dipole moment, without loss of generality. Substituting $\mathbf{B}_0 + \mathbf{B}_c$ for $\mathbf{B}$ and inserting the Hamiltonian in the Schrödinger equation, we obtain the following differential equation for the time evolution of the spin state:

$$\frac{d}{dt}|\chi\rangle = -\frac{i}{2}\gamma \left[\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} B_0 + \begin{pmatrix} 0 & e^{-i(\omega t + \phi)} \\ e^{+i(\omega t + \phi)} & 0 \end{pmatrix} B_c \right] |\chi\rangle. \tag{8.6}$$

After substituting ω_0 for $-\gamma B_0$ and ω_c for $-\gamma B_c$, this can be written as a system of coupled differential equations:

$$\begin{aligned} \dot{\alpha} &= -\frac{1}{2}i \left(\omega_0 \alpha + \omega_c e^{-i(\omega t + \phi)} \beta \right), \\ \dot{\beta} &= \frac{1}{2}i \left(\omega_0 \beta - \omega_c e^{+i(\omega t + \phi)} \alpha \right). \end{aligned} \tag{8.7}$$

8.2.2   Free precession

In the absence of a transversal field ($B_c = 0$), the solution is trivially

$$\alpha = \alpha_0 e^{-i\omega_0 t/2}, \quad \beta = \beta_0 e^{+i\omega_0 t/2}. \tag{8.8}$$

Another way to write this is using the time evolution operator, $\hat{U}_{FP}(t)$:

$$|\chi(t)\rangle = U_{FP}(t)|\chi_0\rangle \quad, \quad U_{FP}(t) = \begin{pmatrix} e^{-i\omega_0 t/2} & 0 \\ 0 & e^{+i\omega_0 t/2} \end{pmatrix}. \tag{8.9}$$

8.2.3   Rotating transversal field

When a rotating transversal magnetic field is applied, as is the case during a $\frac{\pi}{2}$-pulse, the solution is slightly more complex. We can simplify the system of equations in Equation 8.7 by transforming them to the reference frame co-rotating with $\mathbf{B}_c$. This is achieved using the following substitutions:

$$\alpha = \zeta e^{-i(\omega t + \phi)/2}, \quad \beta = \eta e^{+i(\omega t + \phi)/2}, \tag{8.10}$$

where ω is the angular rotation frequency of the transversal field. This is not to be confused with the precession frequency about the static field, ω_0. The system of equations now becomes:

$$\dot{\zeta} = \frac{1}{2}i\left((\omega - \omega_0)\zeta - \omega_c\eta\right),$$
$$\dot{\eta} = -\frac{1}{2}i\left((\omega - \omega_0)\eta + \omega_c\zeta\right). \tag{8.11}$$

By taking a time derivative and substituting one of the equations in the other, we obtain

$$\ddot{\eta} = -\frac{1}{4}\left((\omega - \omega_0)^2 + \omega_c^2\right)\eta, \tag{8.12}$$

which is analogous to the harmonic oscillator with the general solution:

$$\eta = A\sin\Omega t + B\cos\Omega t, \quad A, B \in \mathbb{C}. \tag{8.13}$$

Here

$$\Omega \equiv \frac{1}{2}\sqrt{(\omega - \omega_0)^2 + \omega_c^2}. \tag{8.14}$$

Inserting the solution for η in ζ (in Equation 8.11) leads us to the solution for ζ

$$\zeta = C\sin\Omega t + D\cos\Omega t, \quad \text{with}$$
$$C = -\frac{2i\Omega B - (\omega + \omega_0)A}{\omega_c},$$
$$D = \frac{2i\Omega A - (\omega - \omega_0)B}{\omega_c}. \tag{8.15}$$

Now that we have the general solution in the rotating reference frame, we can use Equation 8.10 to obtain the solutions in the laboratory frame:

$$\alpha = (A\sin\Omega t + B\cos\Omega t)\,e^{-i(\omega t + \phi)/2},$$
$$\beta = (C\sin\Omega t + D\cos\Omega t)\,e^{+i(\omega t + \phi)/2}. \tag{8.16}$$

The initial conditions give us

$$B = \alpha_0 e^{+i\phi/2},$$
$$D = \beta_0 e^{-i\phi/2},$$
$$A = \frac{-(\omega - \omega_0)\alpha_0 e^{+i\phi/2} + \omega_c\beta_0 e^{-i\phi/2}}{2i\Omega},$$
$$C = \frac{\omega_c\alpha_0 e^{+i\phi/2} + (\omega - \omega_0)\beta_0 e^{-i\phi/2}}{2i\Omega}. \tag{8.17}$$

The general solution can be written in matrix form as

$$|\chi(t)\rangle = \hat{U}_{RF}(t)|\chi_0\rangle, \tag{8.18}$$

with

$$\hat{U}_{RF}(t) = \begin{pmatrix} \left(-\frac{\omega - \omega_0}{2i\Omega}\sin\Omega t + \cos\Omega t\right)e^{-i\omega t/2} & \frac{\omega_c}{2i\Omega}\sin\Omega t e^{-i\omega t/2 - i\phi} \\ \frac{\omega_c}{2i\Omega}\sin\Omega t e^{+i\omega t/2 + i\phi} & \left(\frac{\omega - \omega_0}{2i\Omega}\sin\Omega t + \cos\Omega t\right)e^{+i\omega t/2} \end{pmatrix}. \tag{8.19}$$

The object $\hat{U}_{RF}(t, t_0)$ is the time evolution operator. The advantage of describing the effect of the rotational RF pulses on the spin state in terms of time evolution operators is that they can be calculated independently of the initial state, and can be multiplied to study the effect of consecutive pulses on any initial state. For example, the final spin state of the UCN after a Ramsey cycle can be written as:

$$|\chi_f\rangle = \hat{U}_{\text{Ramsey}}|\chi_i\rangle$$
$$= \hat{U}_{RF2}(B_0, \gamma, t_{p2}, \omega_2, B_{c2}, \phi_2)\,\hat{U}_{FP}(B_0, \gamma, T)\,\hat{U}_{RF1}(B_0, \gamma, t_{p1}, \omega_1, B_{c1}, \phi_1)\,|\chi_i\rangle, \tag{8.20}$$

where B_0 is the static magnetic field, γ is the gyromagnetic ratio, $t_{p1,2}$ are the pulse durations, T is the free precession duration, $\omega_{1,2}$ are the pulse frequencies, $B_{c1,2}$ are the transversal magnetic field strengths, and $\phi_{1,2}$ are the phases of the applied pulses. The Larmor frequencies about the static and transversal magnetic fields are calculated using the gyromagnetic ratio. Usually, ω, ω_c and t are identical for both pulses, so we can drop their indices. Because of rotational symmetry, we can set ϕ_1 to zero. We define ϕ_2 as:

$$\phi_2 = \phi_1 + (t + T)\omega + \Delta\phi, \tag{8.21}$$

where $\Delta\phi$ is a parameter of the experiment. In a classical Ramsey experiment, $\Delta\phi$ is zero, which means both pulses are in phase. $\Delta\phi$ can also be varied instead of ω to generate a fringe pattern.

8.2.4   Spin measurement

So far, we have developed the tools we need to calculate the spin state during a pulse sequence, as a function of time. The spin direction is extracted from the spinor by computing the expectation value of the spin operator:

$$\langle \hat{\mathbf{S}} \rangle(t) = \langle \chi(t) | \; \frac{\hbar}{2} \begin{pmatrix} \sigma_x \\ \sigma_y \\ \sigma_z \end{pmatrix} \; | \chi(t) \rangle. \tag{8.22}$$

The components of the spin vector are non-commutative. Therefore, one can only measure the spin along one specified axis. In a precession experiment, the spin is measured in the direction of the static field.

In n2EDM, the spin states of the UCN coming from each precession chamber after a Ramsey cycle are measured in two U-shaped spin analysers (USSA). Each USSA has two legs, each incorporating a spin flip coil, a spin analysing foil, and a UCN detector. The spin analysing foils are 25 μm thick aluminium foils, coated with a 400 nm layer of iron. The iron layer is magnetised to 1.8 to 2.0 T [65]. This adds approximately $\pm$ 120 nT to the iron Fermi potential (180 neV), which becomes 60 neV or 300 neV respectively, depending on the UCN spin projection onto the magnetisation axis. Therefore, only UCN with a specific spin state pass through the foil. In one of the legs, the spin-flip coil is turned on, so that both detectors measure different spin states.

After each Ramsey cycle, we will get the number of detected spin-up (N^+) and spin-down (N^-) UCN for each chamber. From this, we calculate the asymmetry

$$A = \frac{N^+ - N^-}{N^+ + N^-}. \tag{8.23}$$

For a population in a pure spin state $|\chi\rangle$, the asymmetry is identical to the expectation value:

$$\langle \chi | \sigma_x | \chi \rangle = \frac{2}{\hbar} \langle \hat{S}_z \rangle. \tag{8.24}$$

In reality, the UCN will not be in a pure state due to depolarisation induced by field inhomogeneities. The efficiency of the spin flippers and spin analysing foils is excellent, but not perfect. These effects are captured in a parameter called visibility, which we extract from the Ramsey fit. A detailed discussion hereof is beyond the scope of this chapter. From now on, we will assume that the neutrons are in a pure state and that the visibility of the Ramsey curve is equal to one. Unless indicated explicitly, the axis along which the asymmetry is calculated is the z-axis ($A = A_z$).

8.3   Fringe patterns

Using the time evolution matrices derived in the previous section, we can investigate some Rabi and Ramsey fringe patterns.

8.3.1 The Rabi method

The precursor of Ramsey interferometry was the Rabi resonance method, which was developed by Isodor Rabi to measure nuclear magnetic moments of atoms [17]. He won the 1944 Nobel Prize for this work. In Rabi's method, a single π-pulse is applied to a spin system, the pulse frequency is scanned to find the Larmor frequency. The Rabi method is a good starting point for understanding the Ramsey method, which reduces to the Rabi method if the duration of the free precession period is set to zero.

In Figure 8.1, the evolution of a spin system (represented by the asymmetry) is shown during Rabi pulses with varying frequencies. The spin system always starts vertically aligned (aligned with the static field). Once the transversal field rotates the spin into the horizontal plane, the precession around the vertical field becomes apparent. The final asymmetry depends on how closely the pulse frequency (ω) approaches the Larmor frequency (ω_0) about the static field. The asymmetry after the pulse, as a function of ω, follows a fringe pattern with the shape

$$
\begin{aligned}
A(\omega) &= -1 + 2\left|\begin{pmatrix} 1 & 0 \end{pmatrix} U_{RF} \begin{pmatrix} 1 \\ 0 \end{pmatrix}\right|^2, \\
&= -1 + 2|U_{RF,11}|^2, \\
&= -1 + 2\left(\cos^2 \Omega t + \left(\frac{\omega - \omega_0}{2\Omega}\right)^2 \sin^2 \Omega t\right).
\end{aligned}
\tag{8.25}
$$

The fringe pattern for a Rabi pulse, together with those of Ramsey cycles, is shown in Figure 8.2 (for the Rabi pulse, the free precession time $T = 0$).

8.3.2 The Ramsey method

Ramsey's method of separated oscillatory fields [16] is an improvement of Rabi's method. The π-pulse is split into two $\frac{\pi}{2}$-pulses, which are separated by a period of free precession. The longer the free precession time T, the more phase difference builds up between the pulse frequency ω and the Larmor frequency about $\mathbf{B}_0$, namely ω_0. A simple Ramsey experiment has the independent parameters B_0, T, t and ω. The transversal field strength, B_c, is calculated from t, and the expected ω_0, and should be much smaller than B_0.

For several T, the fringe pattern has been plotted in Figure 8.2. For $T = 0$, the Ramsey method reduces to the Rabi method. Increasing T reduces the width of the fringes. The asymmetry oscillates within a certain envelope. This envelope is determined by the pulses and does not depend on T.

The effect of varying the pulse duration, t, is shown in Figure 8.3. Shortening the pulse duration broadens the envelope of the fringe pattern. To understand why this is the case, consider the time evolution of the asymmetry in Figure 8.1, for the completely off-resonance pulse. Initially, the phase of the rotating field pulse is leading the spin phase by $\frac{\pi}{2}$ rad. The spin precesses about the transversal field. As the pulse frequency does not match ω_0, the spin direction does not remain perpendicular to the rotating field, reducing the spin-flip efficiency. After a while, the transversal field is no longer leading the spin direction, but lagging behind it. This inverts the direction of the precession. The z-component of the spin keeps oscillating around its original position for the duration of the pulse. For a pi-half pulse, $\omega_c t = \frac{\pi}{2}$. A shorter pulse requires a stronger transversal field. In off-resonance conditions, this means there are fewer oscillations with a larger amplitude, which means that a larger change in asymmetry can be reached once the pulse ends.

8.4 Side effects of rectangular pulses

An overview of the side effects of the rectangular oscillating field pulses is shown in Figure 8.4.
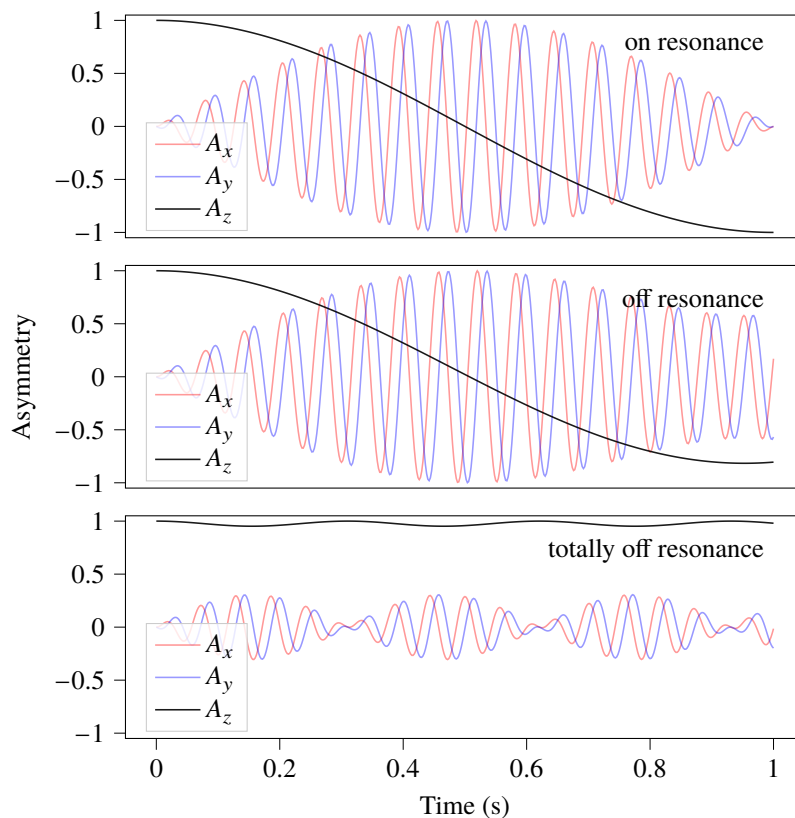
Figure 8.1: The time evolution of a spin system undergoing a 1 s Rabi pulse. The Larmor frequency about the static field equals $\omega_0 = 100$ rad s^{-1}. Top: under the resonance condition, the pulse frequency is equal to the Larmor frequency (100 rad s^{-1}). The spin, which starts parallel with the static magnetic field (the z-direction), starts shifting. As the spin tilts away from the z-axis into the xy-plane, precession is visible as an oscillation of the horizontal spin components. Middle: an off-resonance pulse is applied, with $\omega = 99$ rad s^{-1}. The spin evolution initially looks similar to the resonance spin evolution. Over time, the phase difference between the spin precession and the Rabi pulse accumulates. When the spin is not at a right angle with the transversal field, the spin-flip efficiency is reduced. The asymmetry does not reach -1 at the end of the pulse. Bottom: a 120 rad s^{-1} pulse, totally off-resonance, barely manages to push the spin from its initial axis.
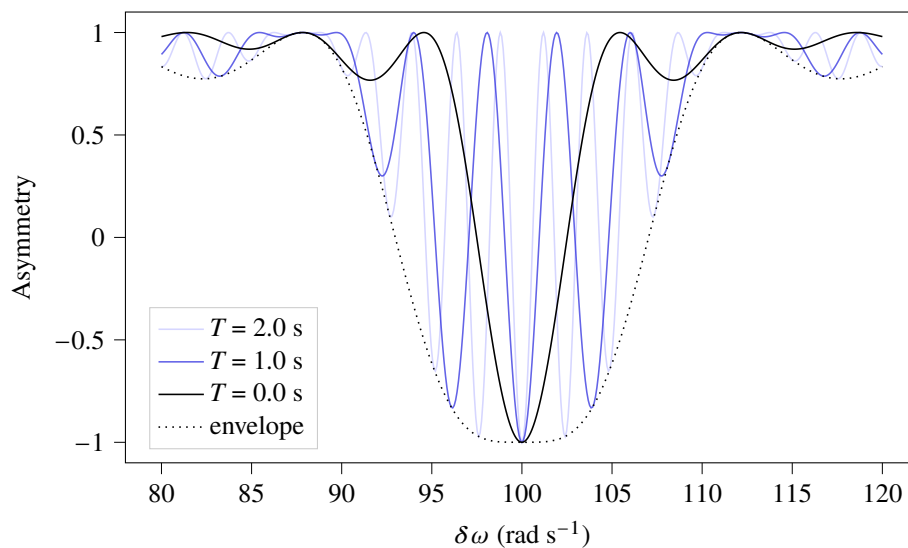
Figure 8.2: The Rabi fringe pattern ($T = 0$) and some Ramsey fringe patterns. $\omega_0 = 100$ rad s^{-1}. The Rabi pulse was a 1 s π-pulse. In the Ramsey cycles, this pulse was split into two 0.5 s $\frac{\pi}{2}$-pulses, separated by a free precession period (T) of 1 or 2 s. Increasing the length of the free precession period causes more phase to build up, leading to narrower fringes, which enables more accurate measurement of the resonance frequency. The pattern will stay within the dotted envelope.
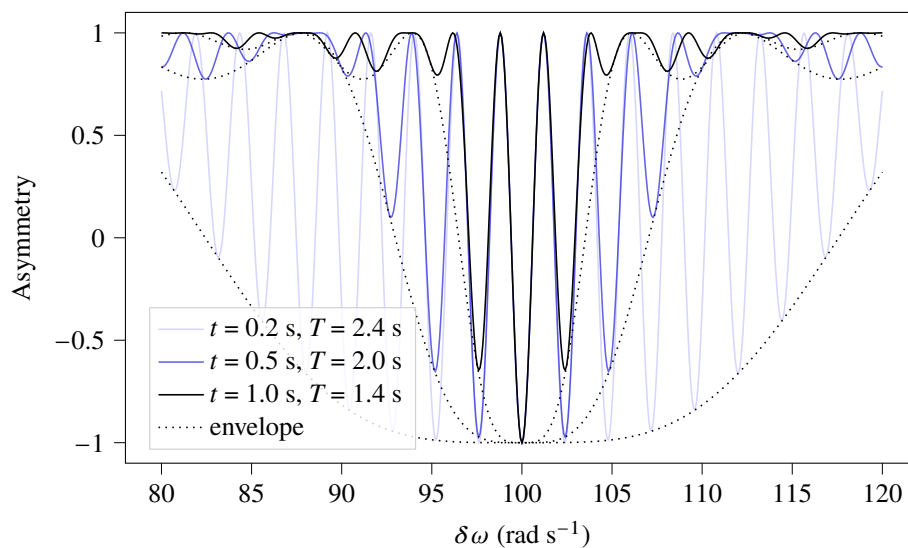
Figure 8.3: More Ramsey fringe patterns. $\omega_0 = 100$ rad s^{-1}. The pulse duration, t, is varied to demonstrate its effect on the pulse envelope. The free precession duration, T, is varied along to keep the effective precession time (explained in Section 8.5.2), and thus the width of the central fringe, equal. Shortening the pulse duration broadens the fringe pattern envelope.

```
┌──────────────┐                      ┌──────────────┐
│   Hg pulse   │                      │   UCN pulse  │
└──────────────┘                      └──────────────┘
```
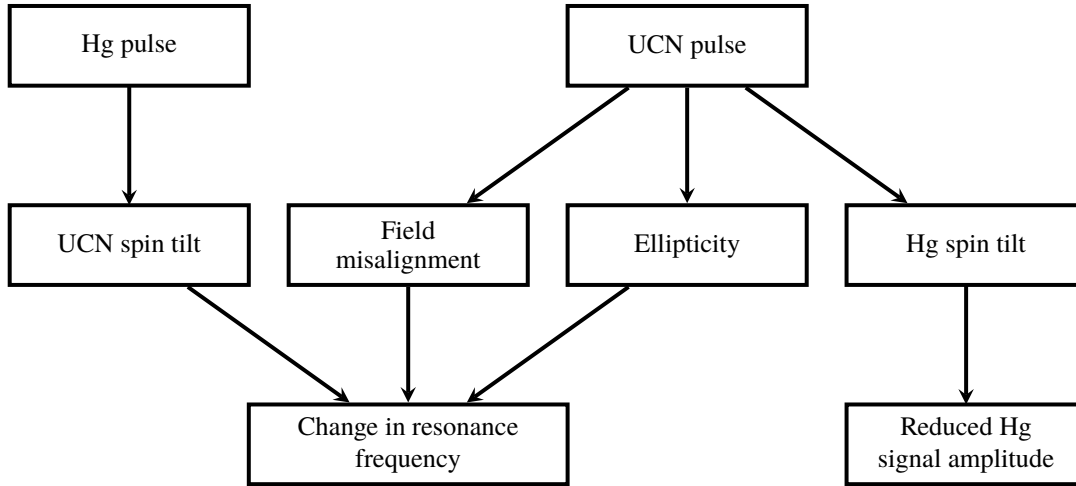

Figure 8.4: An overview of the detrimental effects of pulses.

8.4.1 Mercury co-magnetometry

In n2EDM, mercury-199 is used as a co-magnetometer [64] to monitor the magnetic field strength in the precession chamber in real time with a sensitivity of 25 fT [24]. The mercury vapour is polarised in a polarisation chamber by optical pumping, and injected in the precession chambers, after they have been filled with UCN and the UCN shutters are closed. Before the first UCN $\frac{\pi}{2}$-pulse is applied, a $\frac{\pi}{2}$-pulse is applied to the mercury. During the UCN Ramsey cycle, the mercury spin precesses in the chamber. Its spin state is probed using a laser, and the magnetic field strength is determined from its Larmor frequency.

The mercury pulse affects the UCN spin state, and vice versa. This can affect the visibility of the mercury precession signal, and hence the magnetic field estimation, as well as the estimated Larmor frequency of the neutrons, potentially leading to systematic effects. For a given mercury pulse, $\hat{U}_{RF}(B_0, \gamma_{Hg}, t_p, \omega, B_c, \phi)$, its effect on the UCN spin can be calculated simply by replacing γ_{Hg} by γ_n (see Table 8.1).

To ensure that both the UCN spin and the mercury spin precess in the horizontal plane as perfectly as possible, we need to ensure the pulses affect the polar angle ($z = r\cos(\theta)$) of the other spin as little as possible. For a given time evolution matrix, the change in polar angle depends on the initial orientation of the spin. The maximum change can be achieved when the initial orientation is vertical, for example, when $|\chi_i\rangle = (1, 0)^T$.

$$
\begin{aligned}
\Delta\theta_{\max}(\hat{U}_{RF}) &= \arccos(A_f) - \arccos(A_i), \\
&= \arccos(\langle \chi_f | \hat{A} | \chi_f \rangle) - \arccos(1), \\
&= \arccos(\langle \chi_i | \hat{U}_{RF}^\dagger \hat{A} \, \hat{U}_{RF} | \chi_i \rangle), \\
&= \arccos(|\hat{U}_{RF,00}|^2 - |\hat{U}_{RF,10}|^2).
\end{aligned}
\tag{8.26}
$$

To express this more easily in a graph, we can express this pulse response in decibel:

$$
\text{response} = 10 \log_{10}\left(\frac{\Delta\theta_{\max}(U_{RF})}{\frac{\pi}{2}\text{rad}}\right) \text{dB}.
\tag{8.27}
$$

The change in angle is first divided by $\frac{\pi}{2}$ radians, so that the intended target of the pulse has a response of 0 dB.

Using these methods, we calculated the response of $\frac{\pi}{2}$-pulses intended for either UCN or Hg, on a range of gyromagnetic ratios. In the calculation, we used the constants from Table 8.1 for the gyromagnetic ratios, and the static field. See Figure 8.5. The UCN response for the Hg pulse is about -15 dB. This means that the UCN spin can be tilted by up to $10^{-1.5} \cdot 90°$, about 2.8°.

Table 8.1: Constants used in the calculations of pulse effects.

Quantity	Symbol	Value
^{199}Hg gyromagnetic ratio	γ_{Hg}	-29.1646943(69)·2π MHz T^{-1} [124]
neutron gyromagnetic ratio	γ_n	7.5901152(62) ·2π MHz T^{-1} [124]
static field strength	B_0	1 μT
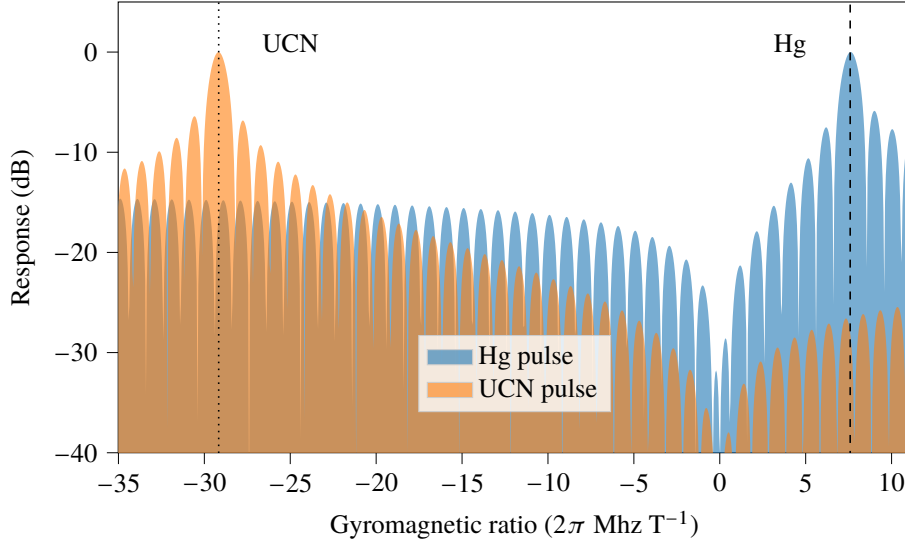
Figure 8.5: Frequency response curves for standard rectangular $\frac{\pi}{2}$-pulses. $B_0 = 1$ μT, $t = 1$ s. Near the UCN frequency, the response to the Hg pulse is about -15 dB. This means that the UCN spin can be tiled by up to 2.8° (see the text).

The goal is to develop $\frac{\pi}{2}$-pulses that suppress these side effects. In the nEDM experiment, this was done by tuning the pulse duration such that the UCN and Hg frequencies fall between the lobes of the response curve [123]. The new function generator in n2EDM can generate arbitrary pulse shapes. We can use this to reduce the side effects of the pulses.

8.4.2   Tilt of the static and rotating magnetic fields

We define the tilt angle as the angle between the static field direction and the rotation axis of the rotating magnetic field. If this tilt angle is non-zero, the rotating field produces an oscillating magnetic field component parallel to the static field. This shifts the estimated precession frequency.

8.5   Pulse modulation using window functions

The $\frac{\pi}{2}$-pulses share a lot of similarities with finite impulse response (FIR) filters. We applied a method used in the design of FIR filters to reduce spectral leakage to mitigate the side effects of the $\frac{\pi}{2}$-pulses.

In digital signal processing, FIR filters are a class of filters that remove undesired frequencies from a sampled signal, with a finite response to a Kronecker δ input (the response terminates after a given time). FIR filters work by calculating the discrete convolution of the signal with the filter coefficients.

To create a primitive FIR filter, one can define the coefficients by sampling a constant amplitude sine wave. The sampling rate must be equal to that of the signal. The pass frequency corresponds to the frequency
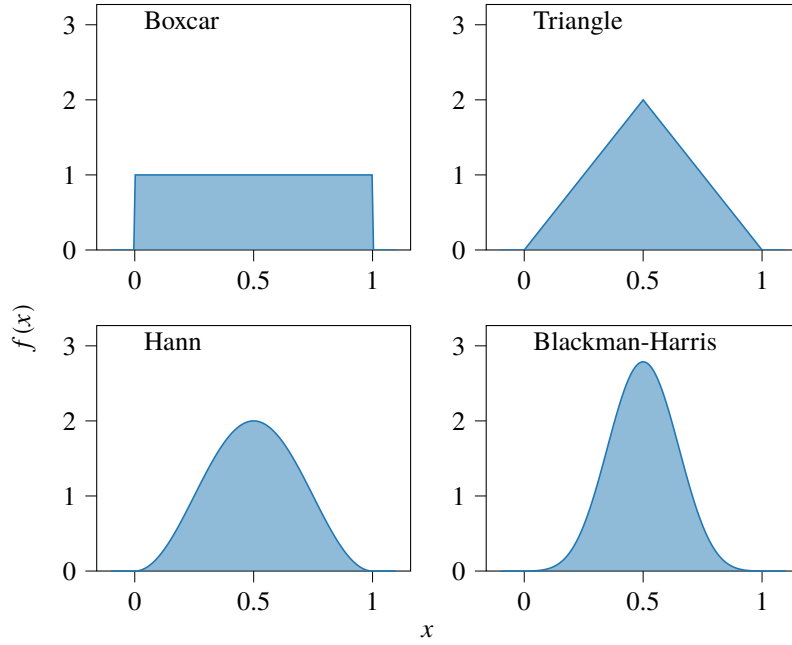
Figure 8.6: Some examples of window functions, zero outside the domain $[0, 1]$, with a total area of 1.

of the sine wave. The spectral leakage is determined by the Fourier transform of the envelope of the filter coefficients. For our primitive filter, the envelope is a boxcar function (constant on a specified domain, and zero everywhere else). The Fourier transform of a boxcar function is a sinc function:

$$\text{sinc}(x) = \begin{cases} \frac{\sin x}{x} & \Leftarrow x \neq 0 \\ 1 & \Leftarrow x = 0 \end{cases}, \tag{8.28}$$

which does not decay to zero quickly.

To modify the spectral leakage outside the target region, the filter coefficients can be modulated using window functions [125]. They are also called apodisation functions since they remove the 'feet' of the frequency response function.

In this work, we require window functions to satisfy the conditions

$$f(x) = 0 \quad \forall \quad x \notin [0, 1], \tag{8.29}$$

$$\exists f_{\max} \in \mathbb{R} : \ \forall x \in [0, 1] : \ |f(x)| \leq f_{\max}, \tag{8.30}$$

and

$$\int_0^1 dx \, f(x) = 1. \tag{8.31}$$

Any bounded, integrable function that is zero outside a finite domain can be rescaled to meet these requirements. Many window functions that were found in other sources were rescaled in this work.

8.5.1   Families of Window functions

All window functions are valued 0 outside the domain $[0, 1]$. This will not be explicitly written. Examples of window functions are shown in Figure 8.6.

Table 8.2: Unnormalised coefficients for common cosine windows.

Name	a_0	a_1	a_2	a_3	a_4
Hann	0.5	0.5			
Hamming	0.54	0.46			
Nutall	0.355768	0.487396	0.144232	0.012604	
Blackman-Nutall	0.3635819	0.4891775	0.1365995	0.0106411	
Blackman-Harris	0.35875	0.48829	0.14128	0.01168	
Flat top	0.21557895	0.4663158	0.277263158	0.083578947	0.006947368

Boxcar or rectangular window   This is the simplest window function.

$$f(x) = \begin{cases} 1 & \Leftarrow 0 < x < 1 \\ \frac{1}{2} & \Leftarrow x = 0 \text{ or } x = 1 \end{cases}. \tag{8.32}$$

B-spline windows   These windows are k-fold convolutions of the rectangular window. In general:

$$f^{(k)}(x) = A \left(\prod_{i=1}^{k} \int dt_i \right) \cdot \delta \left(kx - t_1 \right) \cdot \prod_{i=1}^{k} f_{\text{boxcar}} \left(t_{i-1} - t_i \right), \tag{8.33}$$

where A is a normalisation factor, and k in the term kx in the delta function ensure the function remains zero outside $[0, 1]$. The second-order B-spline window is the triangular window:

$$f(x) = \begin{cases} 4x & \Leftarrow x \le \frac{1}{2} \\ 4 - 4x & \Leftarrow x > \frac{1}{2} \end{cases}, \tag{8.34}$$

The fourth-order B-spline window is the Parzen window

$$x' = \left| x - \frac{1}{2} \right|, \tag{8.35}$$

$$f(x') = \begin{cases} 1 - 6(2x')^2 + 6(2x')^3 & \Leftarrow x' \le \frac{1}{4} \\ 2(1 - 2x')^3 & \frac{1}{4} < x' \le \frac{1}{2} \end{cases}. \tag{8.36}$$

Cosine windows   This family of windows has the general form:

$$f(x) = \sum_{j=0}^{n-1} (-1)^j a_j \cos(2\pi j x). \tag{8.37}$$

In the literature, the coefficients are usually chosen such that the maximum value of the function is 1. Since the only term that does not vanish under the integral is a_0, the functions can be trivially normalised to have an area of 1 by the transformation $a_j \to a_j/a_0$. For some commonly used windows, the coefficients are listed in Table 8.2.

Properties of windowed pulses

The peak and the full width at half maximum (fwhm) of the window functions are shown in Table 8.3. The peak of the window function indicates how much larger the maximum rotating field will be, compared to a rectangular pulse with the same duration. The width of the pulse can be seen as an effective pulse duration, which is relevant for the time averaging of field inhomogeneities.

Table 8.3: Descriptive parameters of window functions.

Window function	f_{max}	FWHM	t_{eff}/t_p
Boxcar	1	1	0.6366
Triangle	2	0.5	0.5953
Parzen	2.6667	0.3612	0.5669
Hann	2	0.5	0.5844
Hamming	1.8519	0.5277	0.5933
Nutall	5.1099	0.1806	0.5603
Flat top	4.6387	0.2328	0.5252
Blackman-Nutall	2.7504	0.3475	0.5650
Blackman-Harris	2.7875	0.3432	0.5641

8.5.2　Effective precession time

During a Ramsey experiment, the central fringe can be fitted using the function

$$A(\omega) = A_{offset} - \alpha \cos\left((\omega - \omega_0)T_{eff} + \phi\right), \tag{8.38}$$

with

$$T_{eff} = T + 2t_{eff}. \tag{8.39}$$

T_{eff} is the effective precession time, and t_{eff} is the effective precession time during a pulse [126]. The spin system under investigation does not only build up phase during the free precession period but also during the $\frac{\pi}{2}$-pulses (the Rabi experiment is exclusively based on this effect).

The effective precession time during a pulse is proportional to the time integral over the projection of the spin in the plane perpendicular to the static field, in the resonance condition.

$$t_{eff} = \int_0^{t_p} dt \sqrt{\langle \chi(t)|A_x|\chi(t)\rangle^2 + \langle \chi(t)|A_y|\chi(t)\rangle^2} = \int_0^{t_p} dt \ \sin(\theta(t)). \tag{8.40}$$

It is convenient to express this integral in terms of the polar angle θ, since on resonance, θ evolves directly proportionally with B_c, and thus ω_c:

$$\left(\frac{d\theta}{dt}\right)(t) = \omega_c(t) = \frac{\pi}{2}\frac{1}{t_p}f\left(\frac{t}{t_p}\right). \tag{8.41}$$

During the first pulse[2] where initially $\theta = 0$, θ evolves as

$$\theta(t) = \frac{\pi}{2t_p} \int_0^t dt' f\left(\frac{t'}{t_p}\right). \tag{8.42}$$

The integral in Equation 8.40 becomes larger if a larger fraction of the spin flip occurs in the first moments of the pulse. If we consider symmetric ($f(\frac{1}{2} + x) = f(\frac{1}{2} - x)$), non-negative ($f(x) \geq 0 \ \forall \ x$) window functions, we find that t_{eff} is higher if more weight of the window function is concentrated in the extremities: because of the sine under the integral, increasing θ from 0 to $\frac{\pi}{4}$ contributes more to t_{eff} than increasing it from $\frac{\pi}{4}$ to $\frac{\pi}{2}$. The window function that maximises t_{eff} under these conditions is

$$f(x) = \frac{1}{2}\delta(x) + \frac{1}{2}\delta(x - 1), \quad t_{eff} = \frac{1}{2}\sqrt{2}t_p. \tag{8.43}$$

If we require that window functions are monotonically non-decreasing for $x < \frac{1}{2}$ and non-increasing for $x > \frac{1}{2}$, the boxcar window maximises t_{eff}. For the boxcar window, $t_{eff} = \frac{2}{\pi}t_p$, making $T_{eff} = T + \frac{4}{\pi}t_p$ for a classical Ramsey experiment [126].

---

[2]For the second pulse, the transformation $\theta \to \theta + \frac{\pi}{2}$ has to be carried out. This has the same effect as the transformation $t' \to t_p - t'$ within the window function. Since all our window functions were symmetric, t_{eff} is equal for both pulses.

The delta function $\delta(x - \frac{1}{2})$ minimises t_{eff} with a value of $\frac{1}{2}t_p$. It follows that $t_{\text{eff}}/t_p \in [\frac{1}{2}, \frac{2}{\pi}]$ for all window functions that fulfil these three conditions. Out of all the window functions we tested, the flat top window is the only window that violated these conditions, since it is negatively valued for some t.

For all windows we evaluated, t_{eff}, in units of t_p, was calculated numerically and tabulated in Table 8.3. Since the sensitivity of a Ramsey experiment depends on the fringe width, and thus T_{eff}, it is favourable for window functions to have a higher t_{eff}. The variation, however, is small. For the experimental parameters $T = 180$ s and $t_p = 2$s, the ratio of the largest to the smallest T_{eff} is 1.0024.

8.5.3   Spin evolution during a windowed pulse

Rectangular pulse product

When we 'window' a pulse, we modulate B_c using a window function:

$$B_c(t) = B_{c,\text{eff}} f\left(\frac{t}{t_p} \right). \tag{8.44}$$

The time evolution operators for the $\frac{\pi}{2}$-pulses, derived in Section 8.2.3, are valid for a constant B_c. It is difficult to find an exact solution in case of varying B_c, but it can be approximated by discretising time, assuming that B_c is constant.

First, we choose a number of steps for the discretisation, n, and calculate the time step, $\Delta t = t_p/n$.

The time is discretised using the sequence $(t_i) = (i\Delta t)_{i<n, i\in\mathbb{N}}$. At every t_i, the time evolution operator $\hat{U}_i$ is calculated, which evolves the state from t_i to t_{i+1}. Each U_i is the exact solution for the time evolution during a pulse with a duration of Δt, with a constant field evaluated at $t = t_i + \frac{1}{2}\Delta t$, where the phase of the rotating field is updated from the initial phase by adding ωt_i:

$$\hat{U}_i = \hat{U}_{RF}(B_0, \gamma, \Delta t, \omega, B_c(t_i + \tfrac{1}{2}\Delta t), \phi_0 + \omega t_i). \tag{8.45}$$

The time evolution operator for the windowed pulse, $\hat{U}_W$, is then obtained by multiplying the operators in the descending order of their indices:

$$\hat{U}_W = \hat{U}_{n-1}\hat{U}_{n-2}...\hat{U}_0. \tag{8.46}$$

We will refer to this method to calculate the time evolution operators for windowed pulses as the rectangular-pulse-product method (RPP method).

Numerical integration using Runge-Kutta method

The rectangular pulse product method has been compared to the Runge-Kutta method (RK45) [127], as implemented in the SciPy function `integrate.solve_ivp` [128]. Runge-Kutta methods are a class of methods to solve real or complex ordinary differential equations of the general form:

$$\frac{d}{dt}y = g(t,y), \tag{8.47}$$

where g is a function of the state, y, and the time, t. In our case, this is the Schrödinger equation, as in Equation 8.6. As the initial state, we can either choose a spin state, or we can choose the 2×2 identity matrix. The latter has the advantage that the solution is a time evolution operator.

The errors of the RPP method and RK45 are compared in Figure 8.7. Both methods converge on the same number to high precision, for very fine time steps ($n > 10^5$). At coarse discretisation ($n < 10^3$), the RPP method benefits from the fact that it is based on valid exact solutions. It has smaller errors than RK45,
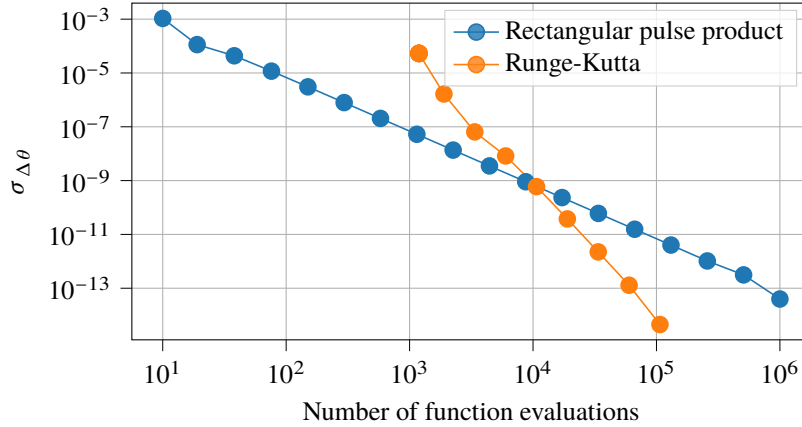
Figure 8.7: The error of the rectangular pulse product method and the Runge-Kutta method, for the computation of the time evolution operators for windowed pulses, for a given pulse shape. The error of the RPP method scales with n^{-2}, while the error for RK45 scales with $n^{-4.5}$. It is not possible to improve the errors beyond approximately 10^{-15}, due to numerical noise. For few evaluations, the RPP method is more accurate than RK45. Around 10^4 function evaluations, RK45 performs better. For an equal number of function evaluations, the execution time is significantly longer for RK45. Apart from the number of function evaluations, the error also depends on the window function, the pulse duration and the frequency.

and unlike RK45, the unitarity of the operators is guaranteed. The error of the RK45 method improves faster over time (with $\mathcal{O}(n^{-4.5})$) than the RPP method ($\mathcal{O}(n^{-2})$). The cross-over point is somewhere around $n = 10^4$. However, the evaluation time was significantly higher for RK45 than for the multiplicative method (respectively 29.2 s and 0.5 s for $n = 600008$).

The shown error curves are for one specific pulse. For different pulses, the offsets, and the crossing point of the two curves may be different. In particular, for pulses that effect a minimal change in θ, the absolute error is lower.

In practice, the RPP method will be used, because of its speed, and robust behaviour around 10^3 function evaluations.

8.6 Suppressing Hg and UCN pulse cross effects.

8.6.1 Side effects for pulses with common window functions

We now have the tools to evaluate the effect of a windowed mercury $\frac{\pi}{2}$-pulse on a UCN spin state, or vice versa. Using the RPP method, the time evolution operator for a pulse with a given target (Hg or UCN), duration and window function is calculated. From this operator, the change in θ is calculated. For pulse durations from 0 to 3 seconds, the UCN spin response to the Hg pulse is plotted in Figure 8.8, for selected window functions. The pulse duration affects the change in θ. For all window functions we tested and for reasonable pulse durations (about 0.5 to 5 s), the UCN response to the Hg pulse is significantly lower than when we use an unmodulated (boxcar) pulse. Of the shown windows, the Parzen window decays the quickest with increasing pulse duration.

The responses of the UCN spin to a mercury pulse, are also shown in Table 8.4, for pulses with a duration of 1, 2 or 3 s. If we naively calculate the pulse effect, we might find ourselves between two lobes of the response curve. Therefore, we created a function that computes the worst case $\Delta\theta$, around $t_p \pm 0.05$s, using the L-BFGS-B algorithm, for a given window function and a given pulse duration, $R_{max}(t,f)$. The L-BFGS-
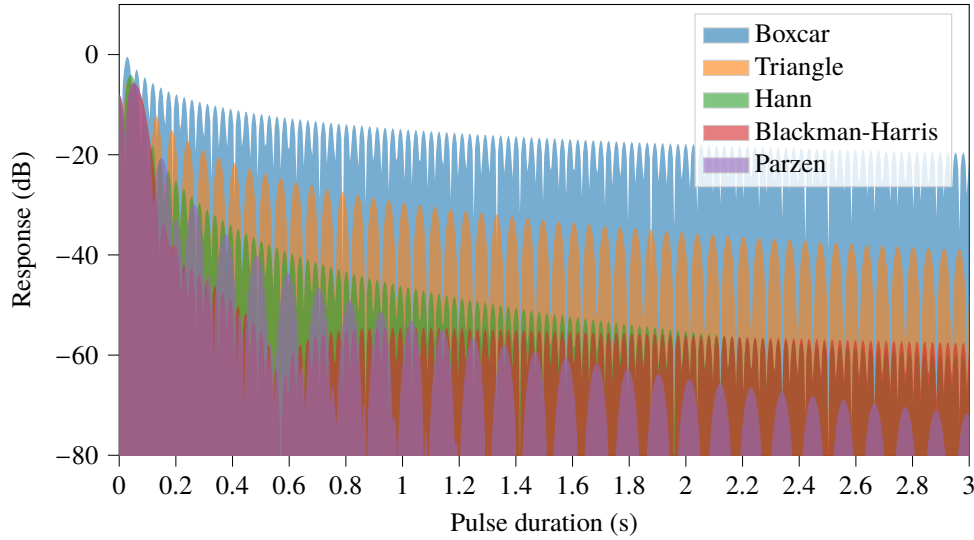
Figure 8.8: The effect of a mercury $\frac{\pi}{2}$-pulse on the UCN spin, as a function of pulse duration. A response of 0 dB means θ can be rotated by $\frac{\pi}{2}$ rad. The boxcar window is the worst. The two tested B-spline windows (triangle, Parzen) get progressively better. Interestingly, their lobes are wider. The cosine windows have lobes equal in shape to those of the boxcar window. What is interesting, is the dip in response around 0.55 s for the Blackman-Harris window.

B algorithm is a modified version of the BFGS algorithm, that works with box constraints [129].

The response of the Hg spin to the UCN pulse is 11.7 dB weaker. This is because the induced $\Delta\theta$ depends on the gyromagnetic ratio. Neutrons have an absolutely larger gyromagnetic ratio, therefore, the UCN $\frac{\pi}{2}$-pulse utilises a weaker field:

$$\gamma_n B_{c,n} t_p = \gamma_{Hg} B_{c,Hg} t_p = \frac{\pi}{2}. \tag{8.48}$$

Therefore, the mercury precession frequency around the transversal field in the UCN pulse is lower:

$$\omega_{c,Hg \text{ in UCN pulse}} = \left(\frac{\gamma_{Hg}}{\gamma_n}\right)^2 \omega_{c,n \text{ in Hg pulse}}. \tag{8.49}$$

8.6.2 Optimisation of cosine windows

When the mercury $\frac{\pi}{2}$-pulse is modulated using the Blackman-Harris window, there is a minimum in the UCN response around $t_p = 0.55$s. This raises the question if we can create a cosine window with optimised parameters to move this minimum in the response curve to the desired pulse duration.

We used the following algorithm to minimise the UCN response to the Hg $\frac{\pi}{2}$-pulse:

- Use the coefficients from the Blackman-Harris window for an initial guess (a_j).

- Optimise the coefficients using the Broyden-Fletcher-Goldfarb-Shanno (BFGS [130]–[133]) gradient descent algorithm:

  - Create a cosine window function, f, using the coefficients (a_j).
  - Calculate $R_{max}(t_p, f) = R_{max}(t_p, (a_j))$.
  - Calculate the gradient, $\nabla_{(a_j)} R_{max}(t_p, (a_j))$.

Table 8.4: Rotation of the UCN initial spin state caused by the mercury pulse and the associated frequency shift. The free precession duration used in the calculation was 180 s.

Window function	UCN response to Hg pulse (-dB), t_p =			Frequency shift (Hz), t_p =		
	1.0 s	2.0 s	3.0 s	1.0 s	2.0 s	3.0 s
Boxcar	14.8	17.8	19.5	4.5×10^{-5}	2.3×10^{-5}	1.5×10^{-5}
Triangle	29.4	35.4	39.0	1.6×10^{-6}	3.9×10^{-7}	1.7×10^{-7}
Parzen	53.0	64.7	71.6	6.8×10^{-9}	4.6×10^{-10}	9.4×10^{-11}
Hann	46.0	55.1	60.4	3.4×10^{-8}	4.2×10^{-9}	1.2×10^{-9}
Hamming	23.1	26.1	27.8	6.7×10^{-6}	3.4×10^{-6}	2.3×10^{-6}
Nutall	19.0	22.1	23.9	1.7×10^{-5}	8.4×10^{-6}	5.6×10^{-6}
Flat top	42.3	45.0	46.7	8.0×10^{-8}	4.3×10^{-8}	2.9×10^{-8}
Blackman-Nutall	45.1	47.9	49.6	4.2×10^{-8}	2.2×10^{-8}	1.5×10^{-8}
Blackman-Harris	54.4	55.9	57.5	5.0×10^{-9}	3.5×10^{-9}	2.4×10^{-9}
Cosine, optimised for 1.0 s	72.9*	61.7	62.7	7.0×10^{-11}	9.2×10^{-10}	7.3×10^{-10}
Cosine, optimised for 2.0 s	58.4	79.7*	70.4	2.0×10^{-9}	1.5×10^{-11}	1.2×10^{-10}
Cosine, optimised for 3.0 s	57.6	68.7	89.6*	2.4×10^{-9}	1.8×10^{-10}	1.5×10^{-12}

* indicates the best response for a given t_p.
The Hg response to the UCN $\frac{\pi}{2}$-pulses is 11.7 dB weaker than the UCN response
to the Hg $\frac{\pi}{2}$-pulses, due to the difference in gyromagnetic ratio.

  – If $|\nabla_{(a_j)} R_{\max}(t_p, (a_j))| < \epsilon$, where ϵ is a predefined tolerance level, terminate the loop.
  – Generate a new guess in the direction of the negative gradient.

 • Return the window function with the coefficients (a_j) that minimise $R_{\max}(t_p, (a_j))$.

The response of the UCN spin to a mercury $\frac{\pi}{2}$-pulse, modulated using cosine windows optimised for pulse durations of 1, 2 and 3 seconds, is shown in Figure 8.9. We succeeded in moving the dip in the response curve to the desired pulse duration.

At the pulse duration it was designed for, the optimised cosine window outperforms all other pulses we investigated (see Table 8.4), for any chosen pulse duration. For 1 s mercury $\frac{\pi}{2}$-pulses, the worst case response was −72.9 dB. This means that at worst, the polar angle of the UCN spin state would rotate by 5.1×10^{-8} rad. For longer pulse durations, the effect is even smaller. The performance of the pulses in the experiment depends on how accurately the designed field can be produced. The pulse design is not the limiting factor.

8.6.3   Hg pulse induced shift in the resonance frequency

A deviation of the initial spin from the vertical axis affects the reconstructed resonance frequency in a Ramsey experiment. We calculated the shift in the fitted nEDM as a function of the initial state to verify that the $\Delta\theta$ induced by the Hg $\frac{\pi}{2}$ pulse is small compared to the design-sensitivity of n2EDM ($\sigma_{\text{sys}} = 0.6 \times 10^{-27}\, e \cdot$ cm, $\sigma_{\text{stat}} = 1.1 \times 10^{-27}\, e \cdot$ cm [24]).

We can calculate the effect of a specific perturbation of the initial state on the final state. For this calculation, we take the limit of $t \to 0$. This simplifies Equation 8.19 in the following way:

$$\omega_c t = \frac{\pi}{2}, \tag{8.50}$$

$$\lim_{t\to 0} \Omega t = \frac{\pi}{4}, \tag{8.51}$$

$$\lim_{t\to 0} \sin \Omega t = \lim_{t\to 0} \cos \Omega t = \frac{1}{2}\sqrt{2}, \tag{8.52}$$
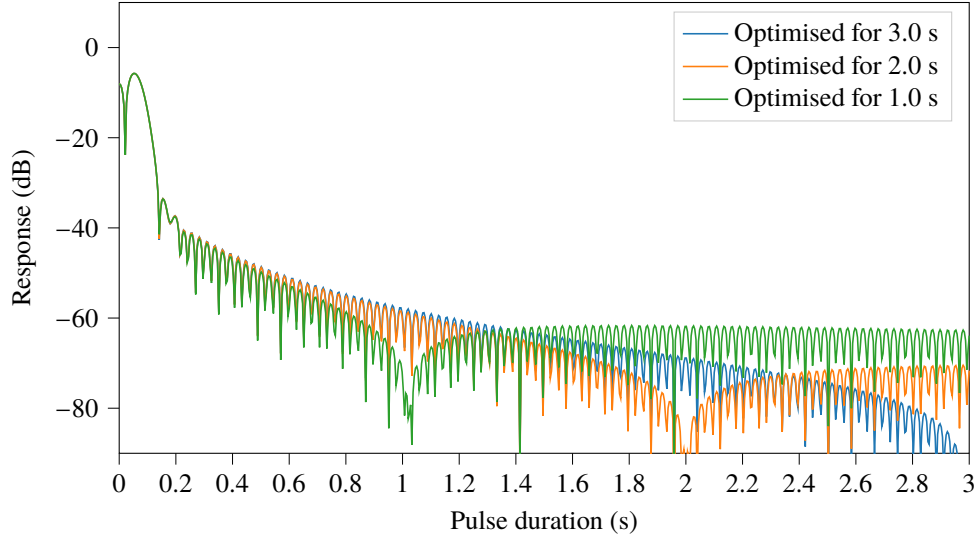
Figure 8.9: The effect of mercury $\frac{\pi}{2}$-pulses modulated with a cosine window with parameters optimised for a specific pulse duration. Using the optimisation algorithm, we can move around the dip in the response function to a desired pulse duration.

$$\lim_{t \to 0} \frac{\omega - \omega_0}{2\Omega} = 0, \quad \lim_{t \to 0} \frac{\omega_c}{2\Omega} = 1, \tag{8.53}$$

$$\lim_{t \to 0} e^{\pm i\omega t} = 1. \tag{8.54}$$

We can use Equation 8.20 to calculate the time evolution matrix for the entire Ramsey experiment, with $\phi_1 = 0$ (at the start of the first pulse, the rotating field points in the x direction), and $\phi_2 = \omega T$. The probability that the spin at the end of the Ramsey sequence is up, P_{up}, is then given by

$$P_{up} \left| \left\langle (1,0) \left| \hat{U}_{\text{Ramsey}} \left| \begin{pmatrix} \alpha \\ \beta \end{pmatrix} \right\rangle \right| \right.^2 . \tag{8.55}$$

Using $A = 2P_{up} - 1$, we find

$$A = (|\beta|^2 - |\alpha|^2) \cos \Delta \omega T + 2\text{Re}(\alpha \beta^*) \sin \Delta \omega T, \tag{8.56}$$

where $\Delta \omega \equiv \omega - \omega_0$. We can use the polar and azimuthal angle, θ and ϕ, to express the initial spin state as

$$\begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} e^{+\frac{i\phi}{2}} \sqrt{\frac{1}{2} + \frac{1}{2} \cos \theta} \\ e^{-\frac{i\phi}{2}} \sqrt{\frac{1}{2} - \frac{1}{2} \cos \theta} \end{pmatrix} \tag{8.57}$$

and rewrite Equation 8.56 as

$$A = -\cos \theta \cos \Delta \omega T + \sin \theta \sin \Delta \omega T \cos \phi. \tag{8.58}$$

In the initial state $(1, 0)$ $(\theta = 0)$, this reduces to the Ramsey pattern

$$A = -\cos \Delta \omega T. \tag{8.59}$$

For $\theta > 0$, and $\phi = \pm \pi$ (the initial spin deviates in the direction of the y-axis), the visibility of the pattern is reduced with factor $\cos \theta$:

$$A = -\cos \theta \cos \Delta \omega T. \tag{8.60}$$

This results in a reduction in UCN statistics, but not in a systematic error. However, for $\theta > 0$ and $\phi = 0$ (the initial spin deviates towards the x-axis)

$$A = -\cos\Delta\omega T\cos\theta - \sin\Delta\omega T\sin\theta$$
$$= -\cos\left(T\left(\Delta\omega + \frac{\theta}{T}\right)\right). \tag{8.61}$$

This means that a tilt of the UCN spin state from the initial state can at best reduce the visibility with factor $\cos\theta$, and at worst shift the resonance frequency by $\delta\omega_0 = \theta/T_{\text{eff}}$. If the phase offset between the Hg pulse and the UCN pulse is always the same, this shift is constant. If the phase offset between the pulses is randomised, this shift is also random and can be averaged out.

For an effective precession time of $T_{\text{eff}} = (180 + \pi)$s, the change in angular frequency is $5.402\,56 \times 10^{-6}$ rad s^{-1} per radian, and the change in frequency equals $1.516\,74 \times 10^{-8}$ Hz per degree. Using

$$d_n = \frac{\hbar\omega_0}{2E}. \tag{8.62}$$

The frequency shift is shown for every window function applied to the mercury shift in Table 8.4.

The total systematics budget of n2EDM corresponds to a frequency shift of $\pm 4.4 \times 10^{-9}$ Hz, assuming $\mathbf{B}_0 = 1\,\mu\text{T}$ and $E = 15\,\text{kV cm}^{-1}$. With optimised pulses, the frequency shift associated with the mercury pulse is several orders of magnitude below that.

Since we took the limit $t \to 0$ to obtain this result, we numerically verified that the result holds for $t = 2\,$s with the exact solution for the boxcar window.

8.7   Orthogonality of the static and rotating fields

So far, we have assumed that the static and orthogonal fields are perfectly orthogonal. In reality, this orthogonality is limited by the precision of the installation of the coils. If $\mathbf{B}_0 \cdot \mathbf{B}_c \neq 0$, the resonance frequency is shifted from the true Larmor frequency. This becomes a systematic error if the tilt angle is correlated with the sign of the electric field. We calculated the induced shift in frequency, and the associated nEDM shift, due to a small tilt angle between the static field, and the rotation axis of the rotating field, θ_{tilt}. In the previous nEDM experiment conducted by our collaboration, the tilt angle was below $0.6°$ [134]. The goal is to design a window function that for a $1°$ tilt keeps the frequency shift significantly below the frequency shift corresponding to the design systematic uncertainty of n2EDM ($6 \times 10^{-28}\, e \cdot \text{cm}$, corresponding to 4.4×10^{-9} Hz [24] for $E = 15\,\text{kV cm}^{-1}$).

We evaluated this using Runge-Kutta, and by using the RPP method.

The main effect is due to the oscillation of the static field during the pulse.

We keep $\mathbf{B}_0$ aligned with the z-axis, and modify $\mathbf{B}_c$:

$$\mathbf{B}_c = \begin{pmatrix} \cos(\omega t + \phi)\cos\theta_{\text{tilt}} \\ \sin(\omega t + \phi) \\ \cos(\omega t + \phi)\sin\theta_{\text{tilt}} \end{pmatrix}, \tag{8.63}$$

$$\Delta\omega = \frac{\pi\tan\theta}{\omega_0 t T}\sin\left(\frac{1}{2}\omega_0 T\right)\cos\left(\frac{1}{2}\omega_0(T + 2t) - \phi\right). \tag{8.64}$$

For the other window functions, we will calculate the frequency shift numerically.

8.7.1   Methods

Runge-Kutta

Using the Runge-Kutta method [127], we can numerically solve the time derivative of the spin state during a pulse:

$$
\frac{d}{dt}|\chi\rangle = -\frac{i}{2}\gamma \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} B_0|\chi\rangle,
$$
$$
-\frac{i}{2}\gamma \begin{pmatrix} -\cos(\omega t + \phi)\sin\theta_{\text{tilt}} & \cos(\omega t + \phi)\cos\theta_{\text{tilt}} - i\sin(\omega t + \phi) \\ \cos(\omega t + \phi)\cos\theta_{\text{tilt}} + i\sin(\omega t + \phi) & \cos(\omega t + \phi)\sin\theta_{\text{tilt}}e^{+i(\omega t + \phi)} \end{pmatrix} \begin{matrix} \\ 0 \end{matrix} Bf(t)|\chi\rangle.
$$
$$
(8.65)
$$

Rectangular pulse product method

Alternatively, we can use the RPP method using the following approximations:

$$
\omega_0 \quad \rightarrow \omega_0 + \omega_c + \cos(\omega t + \phi)\sin\theta_{\text{tilt}},
$$
$$
\omega_c(t) \rightarrow \omega_c(t)\sqrt{\sin^2(\omega t + \phi)\cos^2\theta_{\text{tilt}} + \cos^2(\omega t + \phi)}.
$$
$$
(8.66)
$$

This only gives good results if $dt \ll 1/\omega_0$.

Obtaining the frequency shift

Using either the Runge-Kutta or the RPP method, the asymmetry, A, of the UCN spin state is calculated at the end of the Ramsey sequence. $A(\omega)$ is calculated at 5 equally spaced points in the interval $[\omega_0 - 0.35/T, \omega_0 + 0.35/T]$. We find ω_{min} that minimises $A(\omega)$ using a cosine fit. The frequency shift equals $\omega_{\text{min}} - \omega_0$.

8.7.2   Results

The frequency shift predicted using the Runge-Kutta method was within 10 % of the shift predicted using the RPP method. Like before, the RPP method is preferred, because it is less computationally intensive.

For small angles, the frequency shift is linear in θ, corresponding $\tan\theta$ term in Equation 8.64. The frequency shift oscillates with the initial phase of the pulse, ϕ. In the approximate analytical result, ϕ is encapsulated in the term $\cos(\frac{1}{2}\omega_0(T + 2t) - \phi)$. To find the amplitude of the oscillation, it is not necessary to vary ϕ, if we vary t.

We calculated the frequency shift for different window functions and pulse durations using the RPP method. We used $T = 180\,\text{s}$, $B_0 = 1\,\mu\text{T}$, $\theta_{\text{tilt}} = 1°$, and $\phi = 0$. The results are shown in Figure 8.10.

We constructed the generalised cosine window that minimises the frequency shift for a pulse duration of 2 s. The optimisation method used is similar to the one described in Section 8.6.2. At the design pulse duration, this window function again outperforms other window functions that we tested. This includes the Parzen Window, which has excellent performance at most pulse durations. The optimised window function generates a frequency shift on the order of 1×10^{-12} Hz for a tilt angle of $1°$, which is four orders of magnitude below the shift associated with the systematics budget of n2EDM (4.4×10^{-9} Hz).
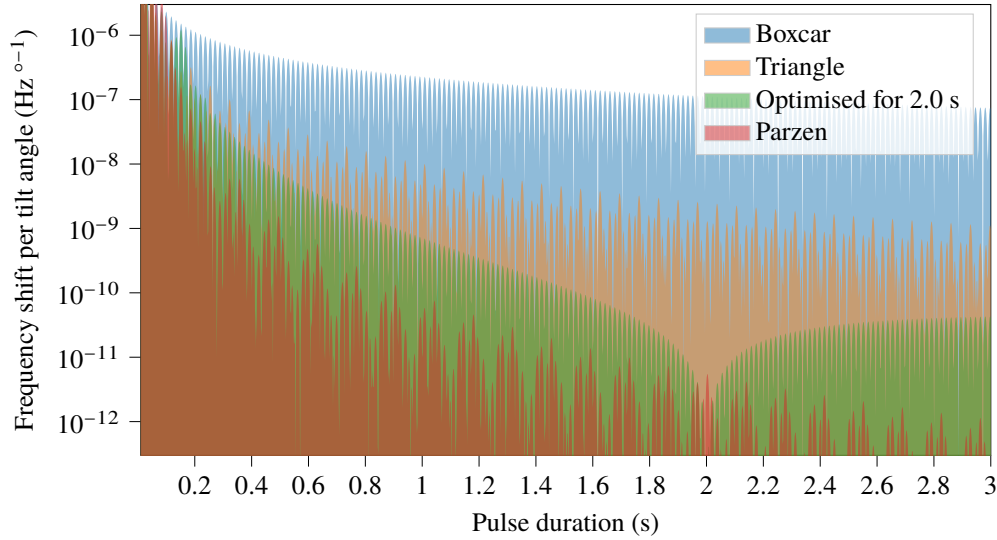
Figure 8.10: The shift in the estimated UCN precession frequency, calculated for a tilt angle between the rotating and static field, as a function of the pulse duration. The shift was calculated from the spin time evolution during an entire Ramsey cycle. The shift is linear in tilt angle. The last window function was optimised to minimise the frequency shift for a pulse duration of 2 s. The systematics budget for n2EDM ($6 \times 10^{-28}\, e \cdot \mathrm{cm}$) corresponds to a shift of 4.4×10^{-9} Hz.

8.8   Discussion

We have developed a versatile set of tools to calculate the evolution of spin in a Ramsey experiment. We studied several effects induced by ideal $\frac{\pi}{2}$-pulses, with different window functions. We found that by applying window functions, in theory, we can suppress the side effects from the Hg $\frac{\pi}{2}$-pulse on the UCN spin state by six orders of magnitude compared to rectangular pulses (see Table 8.4). The effect induced by the UCN $\frac{\pi}{2}$-pulse on the Hg spin state can be suppressed by the same factor. In the real experiment, the suppression factor will depend on how accurately the designed rotating field pulses can be produced. It also depends on how homogeneous the static and magnetic fields can be produced.

By designing cosine window functions with optimised parameters, we can ensure that the lobes in the UCN response curves to the Hg $\frac{\pi}{2}$-pulse are suppressed for the design Hg $\frac{\pi}{2}$-pulse duration. This is more robust to small field fluctuations. With an optimised window function for the Hg pulse, the frequency shift induced by the Hg $\frac{\pi}{2}$-pulse can be kept below 2×10^{-11} Hz, for a pulse duration of 2 s.

For the UCN pulses, the window functions should be optimised to minimise the effects of a possible rotation between the static and rotating magnetic fields. With an optimised pulse, this effect can be kept below 5×10^{-12} Hz.

In the previous nEDM experiment at PSI, the pulse durations were set to 2 s. For n2EDM, the pulse duration has not been decided definitively. There are several arguments for longer or shorter pulse durations.

The n2EDM sensitivity per cycle [24],

$$\sigma_{\mathrm{cycle}} = \frac{\hbar}{2\alpha E T_{\mathrm{eff}}\sqrt{N}}, \tag{8.67}$$

depends on the effective precession time, T_{eff}, and the number of detected UCN at the end of the cycle, N, and the visibility of the Ramsey pattern, α, which is smaller than 1, due to imperfect spin transport, depolarisation in the chamber, and imperfect spin analysis. Longer storage times lead to a larger depolarisation in the chamber, and larger loss of UCN, reducing both α and N. Since the contribution of the UCN $\frac{\pi}{2}$ pulses to T_{eff}

is 34 to 48 % smaller than the contribution of free precession time (while the Hg pulses do not contribute at all), it makes sense to spend the storage time budget on free precession instead of on the pulses, as much as reasonably possible. In addition, if there are absolute errors in the produced field during a pulse, the relative error in the field reproduction is larger for longer pulses.

On the other hand, for longer pulse durations, a larger suppression of side effects can be achieved (see table 8.3. Since the UCN move slowly, during a longer pulse they sample the field in the precession chamber more evenly, so that the pulse effect is less sensitive to field inhomogeneities. This is especially relevant for window-modulated pulses, where the majority of the spin flip happens in a small part of the pulse, evidenced by the FWHM of the pulse windows (Table 8.3).

As a possible compromise, one could use a short Hg pulse and longer UCN pulses. The Hg atoms travel much faster than the UCN (178 versus 5 m s^{-1}), and average out the field inhomogeneities in a shorter time span. The dip in the response curve for the optimised window function limits the UCN response to the Hg pulse, even at slight inhomogeneities of the transversal field. The longer UCN pulses allow the UCN to sample the field effectively, while still contributing to the effective precession time.

The ideal pulse durations will have to be decided on with experimental data. Then, the ideal window function can be calculated using the developed tools.

It is possible to design a combined $\frac{\pi}{2}$-pulse that acts on the Hg and UCN spin simultaneously. The time this would save (one pulse duration) could be used to make the precession time slightly longer, without additional loss of UCN. However, such a pulse sequence is more complicated to understand and more difficult to tune correctly, making it more difficult to control for systematic effects. Therefore, we decided not to pursue this option.

CHAPTER 9

SUMMARY AND OUTLOOK

This thesis describes various efforts to improve the statistical sensitivity of the n2EDM experiment. The statistical sensitivity depends on the number of UCN that are counted in the detectors after each cycle of the experiment. This is influenced by the UCN production and extraction rates in the UCN source moderator, and by the UCN transport and storage properties of the UCN system in the n2EDM apparatus. In addition, the statistical sensitivity depends on the final polarisation of the neutrons. To preserve the neutron polarisation and prevent frequency shifts in the Ramsey cycle, there are strict limits on the magnetic dipole contamination allowed in the central part of the apparatus.[1] The pulses in the Ramsey sequence need to be tuned to avoid shifts in the neutron resonance frequency, to limit statistical and systematic uncertainties.

The dissertation contains two parts. In part I, work on the UCN source is presented. In part II, work on the n2EDM experiment is presented.

Part I: The UCN source at PSI

We described how we measured the concentration of para-deuterium and hydrogen deuteride in the UCN source, using Raman spectroscopy. The molar hydrogen deuteride (HD) concentration is $(0.210 \pm 0.002_{\mathrm{sys}} \pm 0.016_{\mathrm{stat}})$ %. Our value is within 2σ of the previous result of 0.18 % by N. Hild [84]. This confirms the previous result with a more accurate reference mixture and shows that the HD concentration has not significantly increased since 2018. At this concentration, the neutron capture rate in ^{1}H is approximately $4.5\,\mathrm{s}^{-1}$. This is comparable to the $J = 1 \rightarrow 0$ upscattering rate in para-deuterium at a para-deuterium fraction of 0.6 %, which is typical in the UCN source after several months of operation. This is smaller than the capture rate in deuterium $(6.92\,\mathrm{s}^{-1})$ and the thermal upscattering rate at the effective temperature of the UCN source moderator ($12.8\,\mathrm{s}^{-1}$ at 6.3 K). UCN losses in the moderator are dominated by neutron capture in deuterium, which cannot be avoided, and thermal upscattering. Based on the findings, it is not warranted to replace the UCN source deuterium with even higher-purity deuterium, which would be technically difficult. Lowering the deuterium temperature would decrease the loss rate more.

During UCN source operation, radiation-induced heating causes sublimation and desublimation of the solid deuterium moderator surface, creating a frost layer on the surface. This rough layer backscatters UCN, reducing UCN yield. The UCN yield can be recovered by a thermal cycling procedure called conditioning. This involves gradual heating and cooling of the surface, causing the frost layer to sublimate and desublimate in a controlled way, removing the frost layer. This was previously done using electric heaters near the lid of the moderator vessel. We investigated a new method that relies on the radiation heating from proton beam pulses to perform conditioning. The duration of conditioning with pulses is shorter (1 h instead of 2 h to 3 h). UCN production is not interrupted, but temporarily reduced. We found that conditioning with pulses is reliable, and is a suitable mode of operation for the UCN source during long-term experiments. The optimal

---

[1]The limit to prevent frequency shifts is several orders of magnitude lower (stricter) than to prevent UCN depolarisation.

conditioning frequency is the one that maximises UCN statistics for n2EDM (or other experiments). This depends on whether the UCN produced during UCN source conditioning are used. After freezing the deuterium moderator, the optimal frequencies are about 4× per day and 2× per day, if the UCN produced during conditioning are used or not used in the n2EDM analysis, respectively. After intensive UCN source operation, the ortho-deuterium fraction increases above 99 %, and the optimal conditioning frequency approaches 2× or 1× per day.

In 2022, we obtained a record UCN yield of 5.71×10^7 UCN during production pulses, normalised to 2.2 mA of beam current, with a deuterium mass of 5.28 kg in the moderator vessel. We filled the moderator vessel with this mass of deuterium by heating the converter to extract part of the deuterium adsorbed in the oxisorb, as was done in 2021 [48]. Later in the year, we filled the moderator vessel with a record mass of 5.68 kg of deuterium. This resulted in a high UCN yield but did not break the record set with a mass of 5.28 kg. Since the pressure rise in the moderator vessel during pulses was higher with the large deuterium mass, we suspect that the filling level was too high compared to the cooling channels, so that the deuterium surface was not cooled adequately. We expect even higher UCN yields and thus improved UCN statistics for n2EDM with a new moderator vessel with adapted cooling channels, which allows for a larger deuterium mass to be cooled effectively.

We compared the UCN intensities of beamports South and West-1 at the PSI UCN source. The UCN source intensities at beamports West-1 and South are stable over the long term: the relative standard deviation of the UCN count ratio between these beamports for the interval from 12 to 200 s after a proton beam kick is 0.14 % over the long term, and 0.10 % for shorter measurements with slope correction. This means beamport South can be used to monitor the UCN intensity at beamport West-1 to a high precision. The UCN intensity ratio of West-2 compared to the other beamports is less stable. The relative standard deviation is 0.7 % when compared over a short number of pulses (< 50), and 0.4 % with slope correction.

We found that the intensity in beamport South is (81.6 ± 0.2) % of the intensity at beamport West-1, with the caveat that this was measured while a guide section in beamline South was broken and temporarily replaced by a copper sheet. The difference between the UCN intensities is the largest in the first twelve seconds after the pulse. According to mcucn simulations, this effect can be explained by the fact that we only open the Villigen flap of the flapper valve during the pulse. Most of the affected neutrons are too energetic to be stored. Keeping one of the flaps closed reduces the UCN yield at beamport South by about 28 % for realistic n2EDM filling times, regardless of which flap remains closed. The flapper valve will be replaced together with the rest of the UCN source deuterium insert during the service of the UCN source, improving UCN yields.

Part II: The n2EDM experiment at PSI

The UCN transmission of the glass UCN guides for n2EDM was measured using a direct transmission setup. The best UCN guides had a transmission of more than 95 %, normalised to a guide length of 1 m. We found significant differences in UCN transmission between UCN guides. Using cold neutron reflectometry, we verified that the Fermi potential of the witness samples, which are attached to the guides during coating with NiMo, is close to 220 neV. Mcucn simulations demonstrated that the variation in the Fermi potential within the range given by reflectometry cannot explain the differences in transmission. We did find grooves on the interior of some UCN guides, which were created during the manufacturing process of the tubes. According to simulation and measurement, these grooves could explain differences in UCN transmission between guides. Following these results, we replaced the 2.5 m horizontal UCN guides leading to each of the n2EDM precession chambers with groove-less tubes.

The horizontal UCN guides that transport UCN from and to the precession chambers are close to the chamber and the caesium magnetometer array. These guides are not allowed to have magnetic dipoles stronger than 15 nA m². Using the gradiometer, we scanned the UCN guides that are installed closest to the precession chambers and checked for the presence of magnetic dipoles. We found that all guides except one satisfied this strict requirement, even after they had been magnetised with a 30 mT permanent magnet.

Further, we found that the (around 3 mT) magnetic field in the coating machine is not strong enough to significantly magnetise UCN guides. The guides we produce are sufficiently non-magnetic that they do not significantly dephase neutron spins or affect the field reconstruction by the caesium magnetometer array.

We investigated the shifts in the resonance frequency in n2EDM Ramsey cycles caused by several factors. The Hg pulse tilts the UCN spin state by a small amount. This effect causes a shift in the resonance frequency of the neutrons, which depends on the phase between the UCN and Hg $\frac{\pi}{2}$-pulses. A possible tilt between the static and rotating magnetic fields can also induce a frequency shift. By optimising the UCN and Hg $\frac{\pi}{2}$-pulses, we can keep both of these frequency shifts below 2×10^{-11} Hz. The effect of ellipticity of the rotating field was not investigated.

The commissioning of the n2EDM experiment is still ongoing. Many systems worked out of the box, and the first Ramsey cycles with neutrons were measured in 2023. Data taking to measure the nEDM is expected to commence in 2025. We are confident that the experiment will search for the nEDM with unprecedented accuracy and that this work contributed to its success.

Acknowledgments

The opportunity to pursue a PhD degree at ETH Zurich was kindly granted to me by Professor Klaus Kirch. My deepest gratitude goes to him for his outstanding support throughout this work, for his stimulating insights and questions, and for his dedication to his students.

This work would not have been possible without my co-supervisor, Bernhard Lauss, whom I would like to thank for his excellent guidance over the last four years, his enthusiasm, and for being a pleasant colleague to work with.

I would like to extend my gratitude towards Professor Kazimierz Bodek for his constructive feedback on my thesis, and for his participation in my examination.

This work was made possible by SNF grant 188700, which I gratefully acknowledge.

I would like to thank all my colleagues at the UCN group at PSI. Everyone I have worked with is responsible for the good atmosphere and safe environment at work, which is not to be taken for granted. I thank Michael Meier, Luke Noorda and Nadine Kohler for their excellent technical work on the experiment and interesting discussions. Philipp Schmidt-Wellenburg for his introduction to many of the topics at the start of my PhD and for stimulating discussions about Physics. Manfred Daum, who contributed to the good atmosphere in the office, and shared his experience in physics and his knowledge about the history of the PSI with me. Geza Zsigmond for sharing his knowledge about the UCN source and about MCUCN. He was always there to answer my questions. Dieter Ries and Georg Bison, for their input in the programming and installation of DAQ components. Efrain Segarra for being a very helpful and supportive colleague.

I thank my fellow PhD students for the work together the interesting discussions about physics and other topics around the coffee machine, and outside-work activities such as hiking. Duarte Pais and Pin-Jung Chiu for their warm welcome when I started my PhD, and for introducing me to the topics. Ingo Rienäcker, who answered a lot of my questions about the UCN source. Victoria Kletzl commissioned the gradiometer which was used in this work and helped me translate the abstract to German. Thanks also go to Lea Segner, who helped me with measuring some UCN guides in the gradiometer. Wenting Chen for creating a good atmosphere in our office. Nathalie Ziehl for our joyful discussions. Tim Hume for proofreading this thesis and making it a more pleasant read.

I extend my warm gratitude to the Spallation Source Operation group, who operate the UCN source. In particular, I thank Bertrand Blau and Pascal Erismann, who have taught me a lot about the UCN source.

I would like to express my gratitude to Paul Sander, who assisted in measuring the HD concentration in the UCN source deuterium.

The PhD & Postdoc Association[2] and music club at PSI organise many events and provide opportunities to socialise. They made my life at PSI a lot more *gezellig*.

Finally, I would like to thank my family for their practical and emotional support. My parents have always encouraged me to undertake things and explore my interests, and have enabled my education and supported my decisions.

---

[2]https://www.psi.ch/en/ppa

APPENDICES

Appendix A

Personal contributions

Chapter 1: Introduction

I motivated the need for the n2EDM experiment and introduced UCN physics based on the available literature. I explained the design and workings of the UCN source and the n2EDM experiment based on the referenced publications of the UCN group. The UCN source had been built and operational before I joined the collaboration. The largest share of the design of the n2EDM experiment had been finished and construction was underway when I joined the collaboration.

Chapter 2: The UCN source deuterium composition

I introduced the relevant physics, based on referenced literature. The Raman spectroscopy setup in our group was used earlier by Nicolas Hild and described in more detail in his PhD thesis [84].

When we took deuterium samples from the UCN source, Pascal Erisman operated the valves belonging to the UCN source while I collected the sample and transferred it to the spectroscopy setup, where I prepared and performed the measurements.

While the methods used to analyse the Raman spectra were originally inspired by Nicolas Hild's method, I wrote the analysis code from scratch and implemented new methods for background subtraction. Using the new background subtraction, I was able to take more transition lines into account, improving the sensitivity of the analysis.

I performed the measurements and analysis for the determination of the ortho-to-para conversion rates in the sample bottle and the condenser vessel.

The new HD calibration mixture used for the determination of the hydrogen deuteride concentration in the UCN source deuterium was created by the student under my supervision, Paul Sander, with my extensive help. The presented analysis to determine the HD concentration in the calibration mixture and in the UCN source was designed and performed by me.

Chapter 3: Conditioning of the UCN source solid deuterium moderator

The BSQ Group is in charge of the operation of the UCN source, and therefore of performing the conditioning procedure. Typically, Pascal Erisman (BSQ) initiates conditioning at the request of the UCN Group.

I performed all data analysis and interpretation in this chapter. It was my idea to analyse the mirror neutron dataset from 2021 to investigate how the optimal conditioning frequency changes, and I performed the entire analysis and interpreted the results.

Chapter 4: Solid deuterium mass in the moderator vessel and UCN intensity

The method to transfer more deuterium from the to the moderator was used before, and described in [48].

The UCN source was operated by the BSQ group. I prepared the setup required for measuring the UCN yield, with the help of a colleague. I calculated the filling level of the moderator vessel as a function of height and analysed the pressure data.

Chapter 5: A comparison of the UCN intensities at the UCN source beamports

With the help of colleagues, I prepared the setup to measure the UCN rates at the beamports. I analysed the results, calculating the ratio of UCN intensities at the beamports, and the efficiency ratio of the Cascade detectors. The MCUCN simulation code was written by Geza Zsigmond [76]. I adapted the code to simulate the UCN intensities at the beam ports with different flapper valve states and analysed the simulation results.

Chapter 6: UCN transmission of the UCN guides

Many UCN guides for n2EDM were sputter coated with NiMo by me, after I prepared the guide tubes for coating, which required a multistep cleaning procedure. I contributed extensively to the operation of the coating machine. Various errors occurred during the coating process. I identified and solved these errors. I was in charge of all UCN transmission measurements. I constructed the experimental setup and supervised students who helped with the measurements. I planned the measurements and managed the communication with the UCN source operation group. I built the measurement setup for all UCN transmission measurements, aided by various colleagues. I analysed the results of the transmission measurements.

I built the setup to measure the interior surface roughness of the guides using the handheld profilometer and measured the surface profiles.

I measured the Fermi potential of the witness samples using cold neutron reflectometry at the Narziss facility and analysed the results.

I adapted the MCUCN code to be able to exactly simulate the transmission experiments that I conducted at beamport West-1.

Chapter 7: Detection of magnetic dipoles in the UCN guides

The determination of the maximum allowed magnetic dipoles in the n2EDM experiment is explained in the PhD thesis of Duarte Pais [72]. The gradiometer of the UCN group was commissioned primarily by Victoria Kletzl [118].

I designed the wooden UCN guide holder and the measurement method. I performed the measurements of UCN guides using the gradiometer. Further, I designed the analysis method to fit magnetic dipole strengths to the signal measured from UCN guides

Bernhard Lauss sputter coated the UCN guide that I scanned before and after sputter coating. I measured the magnetic field strength in the coating machine and interpreted the result.

Chapter 8: Optimisation of the pulse shapes for n2EDM

The derivation of spin evolution is taken from Daniel John Robert May's PhD thesis [121]. I developed all the tools in Python to calculate the spin evolution during Ramsey cycles with rectangular and windowed pulses, the cross-effects of the Hg $\frac{\pi}{2}$-pulse on the UCN spin state and the frequency shift induced by a tilt in the magnetic fields. The idea to reduce pulse side effects using window functions was introduced to me by Philipp Schmidt-Wellenburg. It was my idea to optimise the parameters of existing Window functions to minimise the side effects. I performed all calculations and analyses. The window functions I designed will be implemented in n2EDM.

APPENDIX B

CRYOBOX

This appendix contains a schematic of the helium and deuterium systems of the UCN source, as designed by B. Blau, taken from [135].

Kryobox

zur kontrollierten Fortluft

zum Spülen / Evakuieren

zum Spülen / Evakuieren

zum KAII-Puffer

Kälteanlage KA II =2QKK30

Phasen-separator

Kryopumpe

Konden-sation

Kon-version

sD₂ Moderator

Servicebox (He Schutzgas)

Deuterium zur Ventilbox

zur Abgas-ringleitung

He-Dichtungskon-trolle =2QRM81

He-Dichtungskon-trolle =2QRM81

D₂ Gas Proben-entnahme

80K Schild UCN Tank

ca. 63K ca. 11.8bar

60K/12.3bar

4.4K/1.3bar

4.5K/4bar

1.3bar

4bar

Cryobox_V4-20

Version 4-20, 28.01.2013, BB34

Bibliography

[1] A. D. Sakharov, "Violation of CP invariance, C asymmetry, and baryon asymmetry of the Universe", *Soviet Physics Uspekhi*, vol. 34, no. 5, p. 392, May 1991. DOI: 10.1070/PU1991v034n05ABEH002497.

[2] G. R. Farrar and M. E. Shaposhnikov, "Baryon asymmetry of the Universe in the minimal standard model", *Phys. Rev. Lett.*, vol. 70, pp. 2833–2836, 19 May 1993. DOI: 10.1103/PhysRevLett.70.2833.

[3] M. Gavela-Legazpi, P. Hernandez, J. Orloff, and O. Pène, "Standard model CP-violation and baryon asymmetry", *Modern Physics Letters A*, vol. 9, Jan. 1994. DOI: 10.1142/S0217732394000629.

[4] E. P. Wigner, "On the matrices which reduce the Kronecker products of representations of S. R. groups", in *The Collected Works of Eugene Paul Wigner: Part A: The Scientific Papers*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1993, pp. 608–654, ISBN: 978-3-662-02781-3. DOI: 10.1007/978-3-662-02781-3_42.

[5] M. Pospelov and A. Ritz, "Electric dipole moments as probes of new physics", *Annals of Physics*, vol. 318, no. 1, pp. 119–169, 2005, Special Issue, ISSN: 0003-4916. DOI: https://doi.org/10.1016/j.aop.2005.04.002.

[6] J. Engel, M. J. Ramsey-Musolf, and U. van Kolck, "Electric dipole moments of nucleons, nuclei, and atoms: The standard model and beyond", *Progress in Particle and Nuclear Physics*, vol. 71, pp. 21–74, 2013, Fundamental Symmetries in the Era of the LHC, ISSN: 0146-6410. DOI: 10.1016/j.ppnp.2013.03.003.

[7] T. E. Chupp, P. Fierlinger, M. J. Ramsey-Musolf, and J. T. Singh, "Electric dipole moments of atoms, molecules, nuclei, and particles", *Rev. Mod. Phys.*, vol. 91, p. 015 001, 1 Jan. 2019. DOI: 10.1103/RevModPhys.91.015001.

[8] C. Abel, S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison, K. Bodek, V. Bondar, M. Burghoff, E. Chanel, Z. Chowdhuri, P.-J. Chiu, B. Clement, C. B. Crawford, M. Daum, S. Emmenegger, L. Ferraris-Bouchez, M. Fertl, P. Flaux, B. Franke, A. Fratangelo, P. Geltenbort, K. Green, W. C. Griffith, M. van der Grinten, Z. D. Gruji ć, P. G. Harris, L. Hayen, W. Heil, R. Henneck, V. Hélaine, N. Hild, Z. Hodge, M. Horras, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaidic, K. Kirch, A. Knecht, P. Knowles, H.-C. Koch, P. A. Koss, S. Komposch, A. Kozela, A. Kraft, J. Krempel, M. Ku źniak, B. Lauss, T. Lefort, Y. Lemière, A. Leredde, P. Mohanmurthy, A. Mtchedlishvili, M. Musgrave, O. Naviliat-Cuncic, D. Pais, F. M. Piegsa, E. Pierre, G. Pignol, C. Plonka-Spehr, P. N. Prashanth, G. Quéméner, M. Rawlik, D. Rebreyend, I. Rienäcker, D. Ries, S. Roccia, G. Rogel, D. Rozpedzik, A. Schnabel, P. Schmidt-Wellenburg, N. Severijns, D. Shiers, R. Tavakoli Dinani, J. A. Thorne, R. Virot, J. Voigt, A. Weis, E. Wursten, G. Wyszynski, J. Zejma, J. Zenner, and G. Zsigmond, "Measurement of the permanent electric dipole moment of the neutron", *Phys. Rev. Lett.*, vol. 124, p. 081 803, 8 Feb. 2020. DOI: 10.1103/PhysRevLett.124.081803.

[9] R. Crewther, P. Di Vecchia, G. Veneziano, and E. Witten, "Chiral estimate of the electric dipole moment of the neutron in quantum chromodynamics", *Physics Letters B*, vol. 88, no. 1, pp. 123–127, 1979, ISSN: 0370-2693. DOI: 10.1016/0370-2693(79)90128-X.

[10] C. Alexandrou, J. Finkenrath, L. Funcke, K. Jansen, B. Kostrzewa, F. Pittler, and C. Urbach, "Ruling out the massless up-quark solution to the strong CP problem by computing the topological mass contribution with lattice QCD", *Physical Review Letters*, vol. 125, no. 23, Dec. 2020, ISSN: 1079-7114. DOI: `10.1103/physrevlett.125.232001`.

[11] R. D. Peccei and H. R. Quinn, "CP conservation in the presence of pseudoparticles", *Phys. Rev. Lett.*, vol. 38, pp. 1440–1443, 25 Jun. 1977. DOI: `10.1103/PhysRevLett.38.1440`.

[12] T. Mannel and N. Uraltsev, "Loop-less electric dipole moment of the nucleon in the standard model", *Phys. Rev. D*, vol. 85, p. 096 002, 9 May 2012. DOI: `10.1103/PhysRevD.85.096002`.

[13] M. Kobayashi and T. Maskawa, "CP-Violation in the Renormalizable Theory of Weak Interaction", *Progress of Theoretical Physics*, vol. 49, no. 2, pp. 652–657, Feb. 1973, ISSN: 0033-068X. DOI: `10.1143/PTP.49.652`.

[14] F. Wilczek, "Problem of strong P and T invariance in the presence of instantons", *Physical Review Letters*, vol. 40, no. 5, pp. 279–282, Jan. 1978. DOI: `10.1103/PhysRevLett.40.279`.

[15] N. J. Ayres, G. Bison, K. Bodek, V. Bondar, T. Bouillaud, E. Chanel, P.-J. Chiu, B. Clement, C. B. Crawford, M. Daum, C. B. Doorenbos, S. Emmenegger, M. Fertl, P. Flaux, W. C. Griffith, P. G. Harris, N. Hild, M. Kasprzak, K. Kirch, V. Kletzl, P. A. Koss, J. Krempel, B. Lauss, T. Lefort, P. Mohanmurthy, O. Naviliat-Cuncic, D. Pais, F. M. Piegsa, G. Pignol, M. Rawlik, I. Rienäcker, D. Ries, S. Roccia, D. Rozpedzik, P. Schmidt-Wellenburg, N. Severijns, B. Shen, K. Svirina, R. T. Dinani, J. A. Thorne, S. Touati, A. Weis, E. Wursten, N. Yazdandoost, J. Zejma, N. Ziehl, and G. Zsigmond, "Search for an interaction mediated by axion-like particles with ultracold neutrons at the PSI", *New Journal of Physics*, vol. 25, no. 10, p. 103 012, Oct. 2023. DOI: `10.1088/1367-2630/acfdc3`.

[16] N. F. Ramsey, "A molecular beam resonance method with separated oscillating fields", *Phys. Rev.*, vol. 78, pp. 695–699, 6 Jun. 1950. DOI: `10.1103/PhysRev.78.695`.

[17] I. I. Rabi, "Space quantization in a gyrating magnetic field", *Phys. Rev.*, vol. 51, pp. 652–654, 8 Apr. 1937. DOI: `10.1103/PhysRev.51.652`.

[18] C. Abel, N. J. Ayres, G. Ban, G. Bison, K. Bodek, V. Bondar, E. Chanel, P.-J. Chiu, M. Daum, S. Emmenegger, *et al.*, "nEDM experiment at PSI: Data-taking strategy and sensitivity of the dataset", in *EPJ Web of Conferences*, EDP Sciences, vol. 219, 2019, p. 02 001. DOI: `10.1051/epjconf/201921902001`.

[19] R. Golub and J. M. Pendlebury, "The electric dipole moment of the neutron", *Contemporary Physics*, vol. 13, no. 6, pp. 519–558, 1972. DOI: `10.1080/00107517208228016`.

[20] C. Baker, Y. Chibane, M. Chouder, P. Geltenbort, K. Green, P. Harris, B. Heckel, P. Iaydjiev, S. Ivanov, I. Kilvington, S. Lamoreaux, D. May, J. Pendlebury, J. Richardson, D. Shiers, K. Smith, and M. van der Grinten, "Apparatus for measurement of the electric dipole moment of the neutron using a cohabiting atomic-mercury magnetometer", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 736, pp. 184–203, 2014, ISSN: 0168-9002. DOI: `10.1016/j.nima.2013.10.005`.

[21] J. H. Smith, E. M. Purcell, and N. F. Ramsey, "Experimental limit to the electric dipole moment of the neutron", *Phys. Rev.*, vol. 108, pp. 120–122, 1 Oct. 1957. DOI: `10.1103/PhysRev.108.120`.

[22] W. B. Dress, P. D. Miller, J. M. Pendlebury, P. Perrin, and N. F. Ramsey, "Search for an electric dipole moment of the neutron", *Phys. Rev. D*, vol. 15, pp. 9–21, 1 Jan. 1977. DOI: `10.1103/PhysRevD.15.9`.

[23] I. Altarev, Y. Borisov, A. Brandin, A. Egorov, V. Ezhov, S. Ivanov, V. Lobashov, V. Nazarenko, G. Porsev, V. Ryabov, A. Serebrov, and R. Taldaev, "A search for the electric dipole moment of the neutron using ultracold neutrons", *Nuclear Physics A*, vol. 341, no. 2, pp. 269–283, 1980, ISSN: 0375-9474. DOI: `10.1016/0375-9474(80)90313-9`.

[24] N. J. Ayres, G. Ban, L. Bienstman, G. Bison, K. Bodek, V. Bondar, T. Bouillaud, E. Chanel, J. Chen, P.-J. Chiu, B. Clément, C. B. Crawford, M. Daum, B. Dechenaux, C. B. Doorenbos, S. Emmenegger, L. Ferraris-Bouchez, M. Fertl, A. Fratangelo, P. Flaux, D. Goupillière, W. C. Griffith, Z. D. Grujic, P. G. Harris, K. Kirch, P. A. Koss, J. Krempel, B. Lauss, T. Lefort, Y. Lemière, A. Leredde, M. Meier, J. Menu, D. A. Mullins, O. Naviliat-Cuncic, D. Pais, F. M. Piegsa, G. Pignol, G. Quéméner, M. Rawlik, D. Rebreyend, I. Rienäcker, D. Ries, S. Roccia, K. U. Ross, D. Rozpedzik, W. Saenz, P. Schmidt-Wellenburg, A. Schnabel, N. Severijns, B. Shen, T. Stapf, K. Svirina, R. T. Dinani, S. Touati, J. Thorne, R. Virot, J. Voigt, E. Wursten, N. Yazdandoost, J. Zejma, and G. Zsigmond, "The design of the n2EDM experiment", *The European Physical Journal C*, vol. 81, no. 6, Jun. 2021. DOI: 10.1140/epjc/s10052-021-09298-z.

[25] D. Wurm, D. H. Beck, T. Chupp, S. Degenkolb, K. Fierlinger, P. Fierlinger, H. Filter, S. Ivanov, C. Klau, M. Kreuz, E. Lelièvre-Berna, T. Lins, J. Meichelböck, T. Neulinger, R. Paddock, F. Röhrer, M. Rosner, A. P. Serebrov, J. T. Singh, R. Stoepler, S. Stuiber, M. Sturm, B. Taubenheim, X. Tonon, M. Tucker, M. v. d. Grinten, and O. Zimmer, "The PanEDM neutron electric dipole moment experiment at the ILL", *EPJ Web Conf.*, vol. 219, p. 02 006, 2019. DOI: 10.1051/epjconf/201921902006.

[26] O. Zimmer and F. Piegsa, "SuperSUN - superfluid source of ultra-cold neutrons", in *ILL2020 Vision - Future directions in neutron science*, Specific Nuclear Reactors and Associated Plants, France, 2011. [Online]. Available: http://inis.iaea.org/search/search.aspx?orig_q=RN:45037233.

[27] S. Sidhu, W. Schreyer, S. Vanbergen, S. Kawasaki, R. Matsumiya, T. Okamura, and R. Picker, "Estimated performance of the TRIUMF ultracold neutron source and electric dipole moment apparatus", *EPJ Web Conf.*, vol. 282, p. 01 015, 2023. DOI: 10.1051/epjconf/202328201015.

[28] F. M. Piegsa, "New concept for a neutron electric dipole moment search using a pulsed beam", *Phys. Rev. C*, vol. 88, p. 045 502, 4 Oct. 2013. DOI: 10.1103/PhysRevC.88.045502.

[29] E. Chanel, Z. Hodge, D. Ries, I. Schulthess, M. Solar, T. Soldner, O. Stalder, J. Thorne, and F. M. Piegsa, "The pulsed neutron beam EDM experiment", *EPJ Web Conf.*, vol. 219, p. 02 004, 2019. DOI: 10.1051/epjconf/201921902004.

[30] A. Steyerl, *Ultracold Neutrons*. World Scientific, 2020, ISBN: 9811212703.

[31] B. Lauss and B. Blau, "UCN, the ultracold neutron source – neutrons for particle physics", *SciPost Phys. Proc.*, p. 004, 2021. DOI: 10.21468/SciPostPhysProc.5.004.

[32] F. M. Gonzalez, E. M. Fries, C. Cude-Woods, T. Bailey, M. Blatnik, L. J. Broussard, N. B. Callahan, J. H. Choi, S. M. Clayton, S. A. Currie, M. Dawid, E. B. Dees, B. W. Filippone, W. Fox, P. Geltenbort, E. George, L. Hayen, K. P. Hickerson, M. A. Hoffbauer, K. Hoffman, A. T. Holley, T. M. Ito, A. Komives, C.-Y. Liu, M. Makela, C. L. Morris, R. Musedinovic, C. O'Shaughnessy, R. W. Pattie, J. Ramsey, D. J. Salvat, A. Saunders, E. I. Sharapov, S. Slutsky, V. Su, X. Sun, C. Swank, Z. Tang, W. Uhrich, J. Vanderwerp, P. Walstrom, Z. Wang, W. Wei, and A. R. Young, "Improved neutron lifetime measurement with UCNτ", *Phys. Rev. Lett.*, vol. 127, p. 162 501, 16 Oct. 2021. DOI: 10.1103/PhysRevLett.127.162501.

[33] A. T. Yue, M. S. Dewey, D. M. Gilliam, G. L. Greene, A. B. Laptev, J. S. Nico, W. M. Snow, and F. E. Wietfeldt, "Improved determination of the neutron lifetime", *Phys. Rev. Lett.*, vol. 111, p. 222 501, 22 Nov. 2013. DOI: 10.1103/PhysRevLett.111.222501.

[34] A. Czarnecki, W. J. Marciano, and A. Sirlin, "Neutron Lifetime and Axial Coupling Connection", *Phys. Rev. Lett.*, vol. 120, no. 20, p. 202 002, 2018. DOI: 10.1103/PhysRevLett.120.202002. arXiv: 1802.01804 [hep-ph].

[35] J. Auler, M. Engler, K. Franz, J. Kahlenberg, J. Karch, N. Pfeifer, K. Roß, C.-F. Strid, N. Yazdandoost, E. Adamek, S. Kaufmann, C. Schmidt, P. Blümler, M. Fertl, W. Heil, and D. Ries, *τSPECT: A spin-flip loaded magnetic ultracold neutron trap for a determination of the neutron lifetime*, 2023. arXiv: 2311.00712 [physics.ins-det].

[36] R. Golub, D. Richardson, and S. Lamoreaux, *Ultra-Cold Neutrons*. Boca Raton: Taylor & Francis, 2008, ISBN: 9780203734803. DOI: 10.1201/9780203734803.

[37] G. Bison, B. Blau, M. Daum, L. Göltl, R. Henneck, K. Kirch, B. Lauss, D. Ries, P. Schmidt-Wellenburg, and G. Zsigmond, "Neutron optics of the PSI ultracold neutron source: characterization and simulation", *Eur. Phys. J. A*, vol. 56, no. 2, p. 33, 2020. DOI: `10.1140/epja/s10050-020-00027-w`. arXiv: `1907.05730 [physics.ins-det]`.

[38] K. Bodek, M. Daum, R. Henneck, S. Heule, M. Kasprzak, K. Kirch, A. Knecht, M. Kuźniak, B. Lauss, M. Meier, G. Petzoldt, M. Schneider, and G. Zsigmond, "Storage of ultracold neutrons in high resistivity, non-magnetic materials with high Fermi potential", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 597, no. 2, pp. 222–226, 2008, ISSN: 0168-9002. DOI: `10.1016/j.nima.2008.09.018`.

[39] I. Altarev, F. Atchison, M. Daum, A. Frei, E. Gutsmiedl, G. Hampel, F. J. Hartmann, W. Heil, A. Knecht, J. V. Kratz, T. Lauer, M. Meier, S. Paul, Y. Sobolev, and N. Wiehl, "Direct experimental verification of neutron acceleration by the material optical potential of solid $^2H_2$", *Physical Review Letters*, vol. 100, no. 1, p. 014 801, Jan. 2008, ISSN: 0031-9007. DOI: `10.1103/PhysRevLett.100.014801`.

[40] National Institute of Standards and Technology, *Fundamental physical constants*, Web Page, Accessed on July 02, 2024, National Institute of Standards and Technology, 2024. [Online]. Available: `https://physics.nist.gov/cgi-bin/cuu/`.

[41] B. Lauss, "Startup of the high-intensity ultracold neutron source at the Paul Scherrer Institute", *Hyperfine Interactions*, vol. 211, no. 1–3, pp. 21–25, Mar. 2012, ISSN: 1572-9540. DOI: `10.1007/s10751-012-0578-7`.

[42] B. Lauss and U. P. Team, "A new facility for fundamental particle physics: The high-intensity ultracold neutron source at the Paul Scherrer Institute", *AIP Conference Proceedings*, vol. 1441, no. 1, pp. 576–578, Sep. 2012, ISSN: 0094-243X. DOI: `10.1063/1.3700622`.

[43] J. Grillenberger, C. Baumgarten, and M. Seidel, "The High Intensity Proton Accelerator Facility", *SciPost Phys. Proc.*, p. 002, 2021. DOI: `10.21468/SciPostPhysProc.5.002`.

[44] D. Anicic, M. Daum, G. Dzieglewski, D. George, M. Horvat, G. Janser, F. Jenni, I. Jirousek, K. Kirch, T. Korhonen, R. Künzi, A. Mezger, U. Rohrer, and L. Tanner, "A fast kicker magnet for the PSI 600 MeV proton beam to the PSI ultra-cold neutron source", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 541, no. 3, pp. 598–609, 2005, ISSN: 0168-9002. DOI: `10.1016/j.nima.2004.12.032`.

[45] B. Blau, K. Kirch, B. Lauss, D. Ries, and G. Zsigmond, *Status of the ultracold neutron source – UCN – at the Paul Scherrer Institute*, Internal nEDM communication, 2013.

[46] L. Göltl, "Characterization of the PSI ultra-cold neutron source", Dissertation number: [20350], Doctoral Thesis, ETH Zurich, Zurich, Switzerland, 2012. DOI: `10.3929/ethz-a-007581200`.

[47] D. A. Ries, "The source for ultracold neutrons at the Paul Scherrer Institute - characterisation, optimisation, and international comparison", Dissertation number: [23671], Doctoral Thesis, ETH Zurich, Zurich, Switzerland, 2016. DOI: `10.3929/ethz-a-010795050`.

[48] I. Rienäcker, "Improving ultracold neutron yields and the search for mirror neutrons at the PSI UCN source", Dissertation number: [28498], Doctoral Thesis, ETH Zurich, Zurich, Switzerland, 2022. DOI: `10.3929/ethz-b-000579840`.

[49] G. Bauer, "Physics and technology of spallation neutron sources", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 463, no. 3, pp. 505–543, 2001, Accelerator driven systems, ISSN: 0168-9002. DOI: `10.1016/S0168-9002(01)00167-X`.

[50] M. Wohlmuther and G. Heidenreich, "The spallation target of the ultra-cold neutron source UCN at PSI", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 564, no. 1, pp. 51–56, 2006, ISSN: 0168-9002. DOI: `10.1016/j.nima.2006.03.040`.

[51] H. Becker, G. Bison, B. Blau, Z. Chowdhuri, J. Eikenberg, M. Fertl, K. Kirch, B. Lauss, G. Perret, D. Reggiani, D. Ries, P. Schmidt-Wellenburg, V. Talanov, M. Wohlmuther, and G. Zsigmond, "Neutron production and thermal moderation at the PSI UCN source", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 777, pp. 20–27, 2015, ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2014.12.091`.

[52] C.-Y. Liu, A. R. Young, C. M. Lavelle, and D. Salvat, *Coherent neutron scattering in polycrystalline deuterium and its implications for ultracold neutron production*, In preparation, 2010. arXiv: `1005.1016 [nucl-th]`. [Online]. Available: `https://arxiv.org/abs/1005.1016`.

[53] F. Atchison, B. Blau, A. Bollhalder, M. Daum, P. Fierlinger, P. Geltenbort, G. Hampel, M. Kasprzak, K. Kirch, S. Köchli, B. Kuczewski, H. Leber, M. Locher, M. Meier, S. Ochse, A. Pichlmaier, C. Plonka, R. Reiser, J. Ulrich, X. Wang, N. Wiehl, O. Zimmer, and G. Zsigmond, "Transmission of very slow neutrons through material foils and its influence on the design of ultracold neutron sources", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 608, no. 1, pp. 144–151, 2009, ISSN: 0168-9002. DOI: `10.1016/j.nima.2009.06.047`.

[54] N. Sullivan, D. Zhou, and C. Edwards, "Precise and efficient in situ ortho — para-hydrogen converter", *Cryogenics*, vol. 30, no. 8, pp. 734–735, 1990, ISSN: 0011-2275. DOI: `10.1016/0011-2275(90)90240-D`.

[55] K. Bodek, B. van den Brandt, T. Bryś, M. Daum, P. Fierlinger, P. Geltenbort, M. Giersch, P. Hautle, R. Henneck, M. Kasprzak, K. Kirch, J. Konter, G. Kühne, M. Kuźniak, K. Mishima, A. Pichlmaier, D. Rätz, A. Serebrov, and J. Zmeskal, "An apparatus for the investigation of solid D_2 with respect to ultra-cold neutron sources", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 533, no. 3, pp. 491–504, 2004, ISSN: 0168-9002. DOI: `10.1016/j.nima.2004.06.157`.

[56] P. C. Souers, *Hydrogen Properties for Fusion Energy*. Univ of California Pr, 1986, ISBN: 0520055004.

[57] S. Afach, N. J. Ayres, G. Ban, G. Bison, K. Bodek, Z. Chowdhuri, M. Daum, M. Fertl, B. Franke, W. C. Griffith, Z. D. Gruji ć, P. G. Harris, W. Heil, V. Hélaine, M. Kasprzak, Y. Kermaidic, K. Kirch, P. Knowles, H.-C. Koch, S. Komposch, A. Kozela, J. Krempel, B. Lauss, T. Lefort, Y. Lemière, A. Mtchedlishvili, M. Musgrave, O. Naviliat-Cuncic, J. M. Pendlebury, F. M. Piegsa, G. Pignol, C. Plonka-Spehr, P. N. Prashanth, G. Quéméner, M. Rawlik, D. Rebreyend, D. Ries, S. Roccia, D. Rozpedzik, P. Schmidt-Wellenburg, N. Severijns, J. A. Thorne, A. Weis, E. Wursten, G. Wyszynski, J. Zejma, J. Zenner, and G. Zsigmond, "Observation of gravitationally induced vertical striation of polarized ultracold neutrons by spin-echo spectroscopy", *Phys. Rev. Lett.*, vol. 115, p. 162 502, 16 Oct. 2015. DOI: `10.1103/PhysRevLett.115.162502`.

[58] P. G. Harris, J. M. Pendlebury, and N. E. Devenish, "Gravitationally enhanced depolarization of ultracold neutrons in magnetic-field gradients", *Phys. Rev. D*, vol. 89, p. 016 011, 1 Jan. 2014. DOI: `10.1103/PhysRevD.89.016011`.

[59] S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison, K. Bodek, M. Fertl, B. Franke, P. Geltenbort, K. Green, W. C. Griffith, M. van der Grinten, Z. D. Gruji ć, P. G. Harris, W. Heil, V. Hélaine, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaidic, K. Kirch, H.-C. Koch, S. Komposch, A. Kozela, J. Krempel, B. Lauss, T. Lefort, Y. Lemière, M. Musgrave, O. Naviliat-Cuncic, J. M. Pendlebury, F. M. Piegsa, G. Pignol, C. Plonka-Spehr, P. N. Prashanth, G. Quéméner, M. Rawlik, D. Rebreyend, D. Ries, S. Roccia, D. Rozpedzik, P. Schmidt-Wellenburg, N. Severijns, D. Shiers, J. A. Thorne, A. Weis, E. Wursten, J. Zejma, J. Zenner, and G. Zsigmond, "Gravitational depolarization of ultracold neutrons: Comparison with data", *Phys. Rev. D*, vol. 92, p. 052 008, 5 Sep. 2015. DOI: `10.1103/PhysRevD.92.052008`.

[60] C. Abel, N. J. Ayres, G. Ban, G. Bison, K. Bodek, V. Bondar, E. Chanel, P.-J. Chiu, B. Clement, C. Crawford, M. Daum, S. Emmenegger, P. Flaux, L. Ferraris-Bouchez, W. Griffith, Z. Grujić, P. Harris, W. Heil, N. Hild, K. Kirch, P. Koss, A. Kozela, J. Krempel, B. Lauss, T. Lefort, Y. Lemière, A. Leredde, P. Mohanmurthy, O. Naviliat-Cuncic, D. Pais, F. Piegsa, G. Pignol, M. Rawlik, D. Rebreyend, D. Ries, S. Roccia, K. Ross, D. Rozpedzik, P. Schmidt-Wellenburg, A. Schnabel, N. Severijns, J. Thorne, R. Virot, J. Voigt, A. Weis, E. Wursten, J. Zejma, and G. Zsigmond, "The n2EDM experiment at the Paul Scherrer Institute", *EPJ Web of Conferences*, vol. 219, T. Jenke, S. Degenkolb, P. Geltenbort, M. Jentschel, V. Nesvizhevsky, D. Rebreyend, S. Roccia, T. Soldner, A. Stutz, and O. Zimmer, Eds., p. 02 002, 2019, ISSN: 2100-014X. DOI: 10.1051/epjconf/201921902002.

[61] A. Chulliat, W. Brown, P. Alken, C. Beggan, M. Nair, G. Cox, A. Woods, S. Macmillan, B. Meyer, and M. Paniccia, "The US/UK world magnetic model for 2020-2025 : Technical report", National Centers for Environmental Information (U.S.) and British Geological Survey, Technical Report noaa:24390, 2020. DOI: 10.25923/ytk1-yx35.

[62] S. Afach, C. A. Baker, G. Ban, G. Bison, K. Bodek, Z. Chowdhuri, M. Daum, M. Fertl, B. Franke, P. Geltenbort, K. Green, M. G. D. van der Grinten, Z. Grujic, P. G. Harris, W. Heil, V. Hélaine, R. Henneck, M. Horras, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaïdic, K. Kirch, P. Knowles, H.-C. Koch, S. Komposch, A. Kozela, J. Krempel, B. Lauss, T. Lefort, Y. Lemière, A. Mtchedlishvili, O. Naviliat-Cuncic, J. M. Pendlebury, F. M. Piegsa, G. Pignol, P. N. Prashant, G. Quéméner, D. Rebreyend, D. Ries, S. Roccia, P. Schmidt-Wellenburg, N. Severijns, A. Weis, E. Wursten, G. Wyszynski, J. Zejma, J. Zenner, and G. Zsigmond, "Measurement of a false electric dipole moment signal from ^{199}Hg atoms exposed to an inhomogeneous magnetic field", *The European Physical Journal D*, vol. 69, no. 10, p. 225, Oct. 2015, ISSN: 1434-6079. DOI: 10.1140/epjd/e2015-60207-4.

[63] G. Ban, G. Bison, K. Bodek, M. Daum, M. Fertl, B. Franke, Z. Grujić, W. Heil, M. Horras, M. Kasprzak, Y. Kermaidic, K. Kirch, H.-C. Koch, S. Komposch, A. Kozela, J. Krempel, B. Lauss, T. Lefort, A. Mtchedlishvili, G. Pignol, F. Piegsa, P. Prashanth, G. Quéméner, M. Rawlik, D. Rebreyend, D. Ries, S. Roccia, D. Rozpedzik, P. Schmidt-Wellenburg, N. Severijns, A. Weis, G. Wyszynski, J. Zejma, and G. Zsigmond, "Demonstration of sensitivity increase in mercury free-spin-precession magnetometers due to laser-based readout for neutron electric dipole moment searches", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 896, pp. 129–138, 2018, ISSN: 0168-9002. DOI: 10.1016/j.nima.2018.04.025.

[64] K. Green, P. Harris, P. Iaydjiev, D. May, J. Pendlebury, K. Smith, M. van der Grinten, P. Geltenbort, and S. Ivanov, "Performance of an atomic mercury magnetometer in the neutron EDM experiment", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 404, no. 2, pp. 381–393, 1998, ISSN: 0168-9002. DOI: 10.1016/S0168-9002(97)01121-2.

[65] S. Afach, G. Ban, G. Bison, K. Bodek, Z. Chowdhuri, M. Daum, M. Fertl, B. Franke, P. Geltenbort, Z. D. Grujić, L. Hayen, V. Hélaine, R. Henneck, M. Kasprzak, Y. Kermaïdic, K. Kirch, S. Komposch, A. Kozela, J. Krempel, B. Lauss, T. Lefort, Y. Lemière, A. Mtchedlishvili, O. Naviliat-Cuncic, F. M. Piegsa, G. Pignol, P. N. Prashanth, G. Quéméner, M. Rawlik, D. Ries, D. Rebreyend, S. Roccia, D. Rozpedzik, P. Schmidt-Wellenburg, N. Severijns, A. Weis, E. Wursten, G. Wyszynski, J. Zejma, and G. Zsigmond, "A device for simultaneous spin analysis of ultracold neutrons", *The European Physical Journal A*, vol. 51, no. 11, p. 143, 2015. DOI: 10.1140/epja/i2015-15143-7.

[66] G. Ban, J. Chen, P.-J. Chiu, B. Clément, M. Guigue, T. Jenke, P. Larue, T. Lefort, O. Naviliat-Cuncic, B. Perriolat, G. Pignol, S. Roccia, W. Saenz-Arevalo, and P. Schmidt-Wellenburg, $n-n'$ *oscillations: Sensitivity of a first UCN beam experiment*, 2022. arXiv: 2206.08721 [nucl-ex].

[67] W. D. Sáenz Arévalo, "A gaseous detector for ultracold neutrons in the n2EDM project", EMJMD on Experimental Nuclear Physics, Master's thesis, Università degli studi di Padova, 2019.

[68] N. J. Ayres, G. Ban, G. Bison, K. Bodek, V. Bondar, T. Bouillaud, B. Clement, E. Chanel, P.-J. Chiu, C. B. Crawford, M. Daum, C. B. Doorenbos, S. Emmenegger, A. Fratangelo, M. Fertl, W. C. Griffith, Z. D. Grujic, P. G. Harris, K. Kirch, J. Krempel, B. Lauss, T. Lefort, O. Naviliat-Cuncic, D. Pais, F. M. Piegsa, G. Pignol, G. Rauscher, D. Rebreyend, I. Rienäcker, D. Ries, S. Roccia, D. Rozpedzik, W. Saenz-Arevalo, P. Schmidt-Wellenburg, A. Schnabel, N. Severijns, B. Shen, M. Staab, K. Svirina, R. T. Dinani, J. Thorne, N. Yazdandoost, J. Zejma, and G. Zsigmond, "The very large n2EDM magnetically shielded room with an exceptional performance for fundamental physics measurements", *Review of Scientific Instruments*, vol. 93, no. 9, Sep. 2022, ISSN: 1089-7623. DOI: `10.1063/5.0101391`.

[69] C. Abel, N. J. Ayres, G. Ban, G. Bison, K. Bodek, V. Bondar, T. Bouillaud, E. Chanel, J. Chen, W. Chen, P.-J. Chiu, C. B. Crawford, M. Daum, C. B. Doorenbos, S. Emmenegger, L. Ferraris-Bouchez, M. Fertl, A. Fratangelo, W. C. Griffith, Z. D. Grujic, P. Harris, K. Kirch, V. Kletzl, P. A. Koss, J. Krempel, B. Lauss, T. Lefort, P. Mullan, O. Naviliat-Cuncic, D. Pais, F. M. Piegsa, G. Pignol, M. Rawlik, I. Rienäcker, D. Ries, S. Roccia, D. Rozpedzik, W. Saenz-Arevalo, P. Schmidt-Wellenburg, A. Schnabel, E. P. Segarra, N. Severijns, T. Shelton, K. Svirina, R. Tavakoli Dinani, J. Thorne, R. Virot, N. Yazdandoost, J. Zejma, N. Ziehl, and G. Zsigmond, "A large 'active magnetic shield' for a high-precision experiment", *The European Physical Journal C*, vol. 83, no. 11, p. 1061, Nov. 2023. DOI: `10.1140/epjc/s10052-023-12225-z`.

[70] S. F. Emmenegger, "Next generation active magnetic shielding for n2edm and axion-like particle search", Dissertation number: [27846], Doctoral Thesis, ETH Zurich, Zurich, Switzerland, 2021. DOI: `10.3929/ethz-b-000515206`.

[71] J. Elen, W. Franken, I. Horvath, G. Pasotti, M. Ricci, J. Roeterdink, N. Sacchetti, M. Spadoni, C. Spoorenberg, G. Vecsey, and P. Weymuth, "The superconductor test facility SULTAN", *IEEE Transactions on Magnetics*, vol. 17, no. 1, pp. 490–493, 1981. DOI: `10.1109/TMAG.1981.1060978`.

[72] D. Pais, "Development of the caesium magnetometer array for the n2EDM experiment", Dissertation number: [27742], Doctoral Thesis, ETH Zurich, Zurich, Switzerland, 2021. DOI: `10.3929/ethz-b-000511496`.

[73] F. Khan, M. Asgar, and P. Nordblad, "Magnetization and magnetocrystalline anisotropy of NiMo single crystal alloys", *Journal of Magnetism and Magnetic Materials*, vol. 174, no. 1, pp. 121–126, 1997, ISSN: 0304-8853. DOI: `10.1016/S0304-8853(97)00182-0`.

[74] K. Bodek, M. Daum, R. Henneck, S. Heule, M. Kasprzak, K. Kirch, A. Knecht, M. Kuźniak, B. Lauss, M. Meier, G. Petzoldt, M. Schneider, and G. Zsigmond, "Storage of ultracold neutrons in high resistivity, non-magnetic materials with high Fermi potential", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 597, no. 2, pp. 222–226, 2008, ISSN: 0168-9002. DOI: `10.1016/j.nima.2008.09.018`.

[75] N. Yazdandoost, Private Communication, Jul. 2024.

[76] G. Zsigmond, "The MCUCN simulation code for ultracold neutron physics", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 881, pp. 16–26, 2018, ISSN: 0168-9002. DOI: `10.1016/j.nima.2017.10.065`.

[77] W. Henggeler, M. Boehm, P. Allenspach, and A. Furrer, "The phase space transformer instrument", *Journal of Neutron Research*, vol. 13, pp. 251–260, Dec. 2005. DOI: `10.1080/10238160512331329965`.

[78] B. Clement, *STARucn - Monte Carlo simulations for the optimization and data analysis of particle physics experiments*, `https://sourceforge.net/projects/starucn/`, Accessed: 2024-04-26.

[79] W. Schreyer, T. Kikawa, M. Losekamm, S. Paul, and R. Picker, "PENTrack—a simulation tool for ultracold neutrons, protons, and electrons in complex electromagnetic fields and geometries", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 858, pp. 123–129, 2017, ISSN: 0168-9002. DOI: `https://doi.org/10.1016/j.nima.2017.03.036`.

[80]　F. Atchison, T. Bryś, M. Daum, P. Fierlinger, A. Fomin, R. Henneck, K. Kirch, M. Kuźniak, and A. Pichlmaier, "The simulation of ultracold neutron experiments using GEANT4", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 552, pp. 513–521, Nov. 2005. DOI: `10.1016/j.nima.2005.06.065`.

[81]　I. H. Lambert, *Photometria sive de mensura et gradibus luminis, colorum et umbrae*, Photometry, or: on the measure of light intensity, color and shade. Augustae Vindelicorum [Augsburg]: Sumptibus viduae Eberhardi Klett, 1760, ETH-Bibliothek Zürich, Rar 1355, Public Domain Mark. DOI: `10.3931/e-rara-9488`.

[82]　R. L. Cook and K. E. Torrance, "A reflectance model for computer graphics", *ACM Trans. Graph.*, vol. 1, no. 1, pp. 7–24, Jan. 1982, ISSN: 0730-0301. DOI: `10.1145/357290.357293`.

[83]　S. Imajo, H. Akatsuka, K. Hatanaka, T. Higuchi, G. Ichikawa, S. Kawasaki, M. Kitaguchi, R. Mammei, R. Matsumiya, K. Mishima, R. Picker, W. Schreyer, and H. M. Shimizu, "Diffuse scattering model of ultracold neutrons on wavy surfaces", *Phys. Rev. C*, vol. 108, p. 034 605, 3 Sep. 2023. DOI: `10.1103/PhysRevC.108.034605`.

[84]　N. Hild, "Studies of the deuterium used in the PSI UCN source", Dissertation number: [26412], Doctoral Thesis, ETH Zurich, Zurich, Switzerland, 2019. DOI: `10.3929/ethz-b-000390393`.

[85]　V. F. Sears, "Neutron scattering lengths and cross-sections", *Neutron News*, vol. 3, no. 3, pp. 26–37, 1992. DOI: `10.1080/10448639208218770`.

[86]　U.S. Nuclear Regulatory Commission, *Physical and chemical properties of tritium*, 2020. [Online]. Available: `https://www.nrc.gov/docs/ML2034/ML20343A210.pdf`.

[87]　C.-Y. Liu, A. R. Young, and S. K. Lamoreaux, "Ultracold neutron upscattering rates in a molecular deuterium crystal", *Phys. Rev. B*, vol. 62, R3581–R3583, 6 Sep. 2000. DOI: `10.1103/PhysRevB.62.R3581`.

[88]　T. Brys, "Extraction of ultracold neutrons from a solid deuterium source", Dissertation number: [17350], Doctoral Thesis, ETH Zurich, Zurich, Switzerland, 2007. DOI: `10.3929/ethz-a-005540305`.

[89]　M. Kasprzak, "Ultracold neutron converters", Matrikel-Nummer: 0448516, Doctroal Thesis, University of Vienna, Vienna, Austria, 2024.

[90]　C. L. Morris, J. M. Anaya, T. J. Bowles, B. W. Filippone, P. Geltenbort, R. E. Hill, M. Hino, S. Hoedl, G. E. Hogan, T. M. Ito, T. Kawai, K. Kirch, S. K. Lamoreaux, C.-Y. Liu, M. Makela, L. J. Marek, J. W. Martin, R. N. Mortensen, A. Pichlmaier, A. Saunders, S. J. Seestrom, D. Smith, W. Teasdale, B. Tipton, M. Utsuro, A. R. Young, and J. Yuan, "Measurements of ultracold-neutron lifetimes in solid deuterium", *Phys. Rev. Lett.*, vol. 89, p. 272 501, 27 Dec. 2002. DOI: `10.1103/PhysRevLett.89.272501`.

[91]　C. V. Raman and K. S. Krishnan, "A new type of secondary radiation", *Nature*, vol. 121, no. 3048, pp. 501–502, Mar. 1928, ISSN: 1476-4687. DOI: `10.1038/121501c0`.

[92]　R. L. McCreery, *Handbook of Vibrational Spectroscopy*. American Cancer Society, 2006, ISBN: 9780470027325. DOI: `10.1002/0470027320`.

[93]　D. Groom, "Cosmic rays and other nonsense in astronomical CCD imagers", *Experimental Astronomy*, vol. 14, no. 1, pp. 45–55, Aug. 2002, ISSN: 1572-9508. DOI: `10.1023/A:1026196806990`.

[94]　Kaiser Optical Systems, Inc., *Analyzer control 4.3 SDK developer's guide*, 2011009 R2, Kaiser Optical Systems, Inc., 371 Parkland Plaza, Ann Arbor, MI 48103, 2011.

[95]　F.-M. Raoult, "Recherches expérimentales sur les tensions de vapeur des dissolutions", Experimental research on the vapor pressures of solutions, *J. Phys. Theor. Appl.*, vol. 8, no. 1, pp. 5–20, 1889. DOI: `10.1051/jphystap:0188900800500`.

[96] N. J. Ayres, Z. Berezhiani, R. Biondi, G. Bison, K. Bodek, V. Bondar, P.-J. Chiu, M. Daum, R. T. Dinani, C. B. Doorenbos, S. Emmenegger, K. Kirch, V. Kletzl, J. Krempel, B. Lauss, D. Pais, I. Rienäcker, D. Ries, N. Rossi, D. Rozpedzik, P. Schmidt-Wellenburg, K. S. Tanaka, J. Zejma, N. Ziehl, and G. Zsigmond, "Improved search for neutron to mirror-neutron oscillations in the presence of mirror magnetic fields with a dedicated apparatus at the PSI UCN source", *Symmetry*, vol. 14, no. 3, 2022, ISSN: 2073-8994. DOI: 10.3390/sym14030503.

[97] K. Motizuki, "Theory of the para-ortho conversion in solid deuterium", *Journal of the Physical Society of Japan*, vol. 12, no. 2, pp. 163–170, 1957. DOI: 10.1143/JPSJ.12.163.

[98] G. Girardi and L. Gasser, *Angebot nr. 2190006224, Deuterium*, Letter, Lenzburg: Messer Schweiz AG, Sep. 2019.

[99] Sigma-Aldrich, *Deuterium hydride*, Certificate of Analysis, Product Number: 488690, Batch Number: MBBD3154, CAS Number: 13983-20-5, MDL Number: MFCD04126430, Formula: HD, Miamisburg, Ohio, US: Sigma-Aldrich, Dec. 2021.

[100] A. Anghel, T. L. Bailey, G. Bison, B. Blau, L. J. Broussard, S. M. Clayton, C. Cude-Woods, M. Daum, A. Hawari, N. Hild, P. Huffman, T. M. Ito, K. Kirch, E. Korobkina, B. Lauss, K. Leung, E. M. Lutz, M. Makela, G. Medlin, C. L. Morris, R. W. Pattie, D. Ries, A. Saunders, P. Schmidt-Wellenburg, V. Talanov, A. R. Young, B. Wehring, C. White, M. Wohlmuther, and G. Zsigmond, "Solid deuterium surface degradation at ultracold neutron sources", *The European Physical Journal A*, vol. 54, no. 9, p. 148, Sep. 2018, ISSN: 1434-601X. DOI: 10.1140/epja/i2018-12594-2.

[101] G. Bison, M. Daum, K. Kirch, B. Lauss, D. Ries, P. Schmidt-Wellenburg, and G. Zsigmond, "Ultracold neutron storage and transport at the PSI UCN source", *The European Physical Journal A*, vol. 58, no. 6, p. 103, Jun. 6, 2022. DOI: 10.1140/epja/s10050-022-00747-1.

[102] G. Bison, W. Chen, P.-J. Chiu, M. Daum, C. B. Doorenbos, K. Kirch, V. Kletzl, B. Lauss, D. Pais, I. Rienäcker, P. Schmidt-Wellenburg, and G. Zsigmond, "Time-of-flight spectroscopy of ultracold neutrons at the PSI UCN source", *The European Physical Journal A*, vol. 59, no. 9, p. 215, 2023. DOI: 10.1140/epja/s10050-023-01104-6.

[103] Z. Berezhiani, "Through the looking-glass: Alice's adventures in Mirror World", in *From Fields to Strings: Circumnavigating Theoretical Physics*, World Scientific, Feb. 2005, pp. 2147–2195. DOI: 10.1142/9789812775344_0055.

[104] Z. Berezhiani, "Mirror World and its cosmological consequences", *International Journal of Modern Physics A*, vol. 19, no. 23, pp. 3775–3806, Sep. 2004. DOI: 10.1142/s0217751x04020075.

[105] A. A. Neath and J. E. Cavanaugh, "The Bayesian information criterion: Background, derivation, and applications", *WIREs Computational Statistics*, vol. 4, no. 2, pp. 199–203, 2012. DOI: https://doi.org/10.1002/wics.199.

[106] Y. A. Freiman and Y. Crespo, "Isotopic and spin-nuclear effects in solid hydrogens", *Low Temperature Physics*, vol. 43, no. 12, pp. 1345–1361, 2017. DOI: 10.1063/1.5012785.

[107] G. Bison, B. Blau, W. Chen, P.-J. Chiu, M. Daum, C. B. Doorenbos, N. Hild, K. Kirch, V. Kletzl, B. Lauss, D. Pais, I. Rienäcker, D. Ries, P. Schmidt-Wellenburg, V. Talanov, and G. Zsigmond, "Characterization of ultracold neutron production in thin solid deuterium films at the PSI ultracold neutron source", *Phys. Rev. C*, vol. 107, p. 035 501, 3 Mar. 2023. DOI: 10.1103/PhysRevC.107.035501.

[108] G. Bison, M. Daum, K. Kirch, B. Lauss, D. Ries, P. Schmidt-Wellenburg, G. Zsigmond, T. Brenner, P. Geltenbort, T. Jenke, O. Zimmer, M. Beck, W. Heil, J. Kahlenberg, J. Karch, K. Ross, K. Eberhardt, C. Geppert, S. Karpuk, T. Reich, C. Siemensen, Y. Sobolev, and N. Trautmann, "Comparison of ultracold neutron sources for fundamental physics measurements", *Phys. Rev. C*, vol. 95, p. 045 503, 4 Apr. 2017. DOI: 10.1103/PhysRevC.95.045503.

[109] E. Shabagin, "Thermische Effusion von Deuterium zwischen Moderator und Membranvakuummmeter in der UCN Anlage", Paul Scherrer Institute, Technical report, 2016.

[110] A. Knecht, "Towards a new measurement of the Neutron Electric Dipole Moment", Advisor: Prof. Dr. Ulrich Straumann, Doctoral Thesis, University of Zurich, Zurich, 2009.

[111] C. D. Technologies, "CASCADE-U: A novel ionization chamber for thermal neutron beam monitoring", CASCADE Detector Technologies, Tech. Rep., 2019, Accessed: 2024-03-30. [Online]. Available: https://n-cdt.com/wp-content/uploads/2019/09/CASCADE-U.pdf.

[112] B. Blau, M. Daum, M. Fertl, P. Geltenbort, L. Göltl, R. Henneck, K. Kirch, A. Knecht, B. Lauss, P. Schmidt-Wellenburg, and G. Zsigmond, "A prestorage method to measure neutron transmission of ultracold neutron guides", *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 807, pp. 30–40, Jan. 2016, ISSN: 0168-9002. DOI: 10.1016/j.nima.2015.10.075.

[113] H. Hernandez, *Probability distribution and bias of the sample standard deviation*, Technical report, Jan. 2023. DOI: 10.13140/RG.2.2.22144.51205.

[114] J. Gurland and R. C. Tripathi, "A simple approximation for unbiased estimation of the standard deviation", *The American Statistician*, vol. 25, no. 4, pp. 30–32, 1971. DOI: 10.1080/00031305.1971.10477279.

[115] G. Schwarz, "Estimating the dimension of a model", *The Annals of Statistics*, vol. 6, no. 2, pp. 461–464, 1978, ISSN: 00905364. [Online]. Available: http://www.jstor.org/stable/2958889 (visited on 05/02/2024).

[116] Paul Scherrer Institute, *NARZISS: Polarised neutron reflectometer*, https://www.psi.ch/en/sinq/narziss, Accessed: 2024-06-07.

[117] V. D. Bondar, S. Chesnevskaya, M. Daum, B. Franke, P. Geltenbort, L. Göltl, E. Gutsmiedl, J. Karch, M. Kasprzak, G. Kessler, K. Kirch, H. C. Koch, A. Kraft, T. Lauer, B. Lauss, E. Pierre, G. Pignol, D. Reggiani, P. Schmidt-Wellenburg, Y. Sobolev, T. Zechlau, and G. Zsigmond, "Losses and depolarization of ultracold neutrons on neutron guide and storage materials", *Physical Review C*, vol. 96, p. 035 205, 2017. DOI: 10.1103/PHYSREVC.96.035205.

[118] V. Kletzl, "Mitigating systematic effects in the n2EDM experiment using optically pumped cesium magnetometers", Dissertation number: TBA, Doctoral Thesis, ETH Zurich, Zurich, Switzerland, 2025.

[119] A. Ayres, G. Ban, G. Bison, V. Bondar, T. Bouillaud, D. C. Bowles, G. L. Caratsch, E. Chanel, W. Chen, P.-J. Chiu, C. Crawford, M. Daum, C. B. Doorenbos, M. Fertl, A. Fratangelo, C. Griffith, K. Kirch, V. Kletzl, J. Krempel, M. Meier, K. Michielsen, J. Micko, P. Mullan, P. Navon, O. Naviliat-Cuncic, B. Lauss, R. Li, T. Lefort, A. Lejuez, D. Pais, F. M. Piegsa, G. Pignol, D. Rebreyend, I. Rienaecker, D. Ries, S. Roccia, W. Saenz-Arevalo, P. Schmidt-Wellenburg, A. Schnabel, E. P. Segarra, N. Severijns, K. Svirina, K. S. Tanaka, R. Tavakoli, J. Thorne, S. Touati, J. Vankeirsbilck, J. Voigt, A. Weis, N. Yazdandoost, N. Ziehl, and G. Zsigmond, "A magnetic-field scanner to detect minuscule magnetic contamination in the n2EDM experiment", *European Physical Journal C*, 2024, In preparation.

[120] C. Abel, S. Afach, N. J. Ayres, G. Ban, G. Bison, K. Bodek, V. Bondar, E. Chanel, P.-J. Chiu, C. B. Crawford, Z. Chowdhuri, M. Daum, S. Emmenegger, L. Ferraris-Bouchez, M. Fertl, B. Franke, W. C. Griffith, Z. D. Gruji ć, L. Hayen, V. Hélaine, N. Hild, M. Kasprzak, Y. Kermaidic, K. Kirch, P. Knowles, H.-C. Koch, S. Komposch, P. A. Koss, A. Kozela, J. Krempel, B. Lauss, T. Lefort, Y. Lemière, A. Leredde, A. Mtchedlishvili, P. Mohanmurthy, M. Musgrave, O. Naviliat-Cuncic, D. Pais, A. Pazgalev, F. M. Piegsa, E. Pierre, G. Pignol, P. N. Prashanth, G. Quéméner, M. Rawlik, D. Rebreyend, D. Ries, S. Roccia, D. Rozpedzik, P. Schmidt-Wellenburg, A. Schnabel, N. Severijns, R. T. Dinani, J. Thorne, A. Weis, E. Wursten, G. Wyszynski, J. Zejma, and G. Zsigmond, "Optically pumped Cs magnetometers enabling a high-sensitivity search for the neutron electric dipole moment", *Phys. Rev. A*, vol. 101, p. 053 419, 5 May 2020. DOI: 10.1103/PhysRevA.101.053419.

[121] D. J. R. May, "A high precision comparison of the gyromagnetic ratios of the ^{199}Hg atom and the neutron", Doctoral Thesis, University of Sussex, Sep. 1998.

[122] P. Ehrenfest, "Bemerkung über die angenäherte Gültigkeit der klassischen Mechanik innerhalb der Quantenmechanik", *Zeitschrift für Physik*, vol. 45, no. 7, pp. 455–457, Jul. 1927, ISSN: 0044-3328. DOI: 10.1007/BF01329203.

[123] P. Schmidt-Wellenburg, Private Communication, Oct. 2020.

[124] S. Afach, C. Baker, G. Ban, G. Bison, K. Bodek, M. Burghoff, Z. Chowdhuri, M. Daum, M. Fertl, B. Franke, P. Geltenbort, K. Green, M. van der Grinten, Z. Grujic, P. Harris, W. Heil, V. Hélaine, R. Henneck, M. Horras, P. Iaydjiev, S. Ivanov, M. Kasprzak, Y. Kermaïdic, K. Kirch, A. Knecht, H.-C. Koch, J. Krempel, M. Kuźniak, B. Lauss, T. Lefort, Y. Lemière, A. Mtchedlishvili, O. Naviliat-Cuncic, J. Pendlebury, M. Perkowski, E. Pierre, F. Piegsa, G. Pignol, P. Prashanth, G. Quéméner, D. Rebreyend, D. Ries, S. Roccia, P. Schmidt-Wellenburg, A. Schnabel, N. Severijns, D. Shiers, K. Smith, J. Voigt, A. Weis, G. Wyszynski, J. Zejma, J. Zenner, and G. Zsigmond, "A measurement of the neutron to 199Hg magnetic moment ratio", *Physics Letters B*, vol. 739, pp. 128–132, 2014, ISSN: 0370-2693. DOI: 10.1016/j.physletb.2014.10.046.

[125] K. Prabhu, *Window Functions and Their Applications in Signal Processing*. CRC Press, 2018, ISBN: 9781315216386. DOI: 10.1201/9781315216386.

[126] C. Abel, S. Afach, N. J. Ayres, C. A. Baker, G. Ban, G. Bison, K. Bodek, V. Bondar, M. Burghoff, E. Chanel, Z. Chowdhuri, P.-J. Chiu, B. Clement, C. B. Crawford, M. Daum, S. Emmenegger, L. Ferraris-Bouchez, M. Fertl, P. Flaux, B. Franke, A. Fratangelo, P. Geltenbort, K. Green, W. C. Griffith, M. van der Grinten, Z. D. Gruji ć, P. G. Harris, L. Hayen, W. Heil, R. Henneck, V. Hélaine, N. Hild, Z. Hodge, M. Horras, P. Iaydjiev, S. N. Ivanov, M. Kasprzak, Y. Kermaidic, K. Kirch, A. Knecht, P. Knowles, H.-C. Koch, P. A. Koss, S. Komposch, A. Kozela, A. Kraft, J. Krempel, M. Ku źniak, B. Lauss, T. Lefort, Y. Lemière, A. Leredde, P. Mohanmurthy, A. Mtchedlishvili, M. Musgrave, O. Naviliat-Cuncic, D. Pais, F. M. Piegsa, E. Pierre, G. Pignol, C. Plonka-Spehr, P. N. Prashanth, G. Quéméner, M. Rawlik, D. Rebreyend, I. Rienäcker, D. Ries, S. Roccia, G. Rogel, D. Rozpedzik, A. Schnabel, P. Schmidt-Wellenburg, N. Severijns, D. Shiers, R. Tavakoli Dinani, J. A. Thorne, R. Virot, J. Voigt, A. Weis, E. Wursten, G. Wyszynski, J. Zejma, J. Zenner, and G. Zsigmond, "Measurement of the permanent electric dipole moment of the neutron", *Phys. Rev. Lett.*, vol. 124, p. 081 803, 8 Feb. 2020. DOI: 10.1103/PhysRevLett.124.081803.

[127] J. Dormand and P. Prince, "A family of embedded Runge-Kutta formulae", *Journal of Computational and Applied Mathematics*, vol. 6, no. 1, pp. 19–26, 1980, ISSN: 0377-0427. DOI: 10.1016/0771-050X(80)90013-3.

[128] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, "SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python", *Nature Methods*, vol. 17, pp. 261–272, 2020. DOI: 10.1038/s41592-019-0686-2.

[129] C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, "Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization", *ACM Trans. Math. Softw.*, vol. 23, no. 4, pp. 550–560, Dec. 1997, ISSN: 0098-3500. DOI: 10.1145/279232.279236.

[130] C. G. Broyden, "The Convergence of a Class of Double-rank Minimization Algorithms 1. General Considerations", *IMA Journal of Applied Mathematics*, vol. 6, no. 1, pp. 76–90, Mar. 1970, ISSN: 0272-4960. DOI: 10.1093/imamat/6.1.76.

[131] R. Fletcher, "A new approach to variable metric algorithms", *The Computer Journal*, vol. 13, no. 3, pp. 317–322, Jan. 1970, ISSN: 0010-4620. DOI: 10.1093/comjnl/13.3.317.

[132] D. Goldfarb, "A family of variable-metric methods derived by variational means", *Mathematics of Computation*, vol. 24, pp. 23–26, 1970. DOI: 10.1090/S0025-5718-1970-0258249-6.

[133] D. F. Shanno, "Conditioning of Quasi-Newton methods for function minimization", *Mathematics of Computation*, vol. 24, pp. 647–656, 1970. DOI: 10.1090/S0025-5718-1970-0274029-X.

[134] G. Pignol, *Systematics on neutron frequency extraction, effects associated with 8 hz and 30 hz pulses*, Internal nEDM communication, Jan. 2010.

[135] B. Blau, A. Anghel, M. Dubs, and J. Welte, "Beschreibung des kryosystems der kalten quelle von UCN", Description of the cryosystem of the cold source of UCN, Paul Scherrer Institut, Internal Technical Report TI-118300/BB34-01.2, Feb. 2010.