



HAL
open science

HOPE - un piège magnétique pour neutron ultra-froid dédié à la mesure du temps de vie du neutron : conception et premières données expérimentales

Felix Rosenau

► **To cite this version:**

Felix Rosenau. HOPE - un piège magnétique pour neutron ultra-froid dédié à la mesure du temps de vie du neutron : conception et premières données expérimentales. Physique [physics]. Université Grenoble Alpes, 2015. Français. NNT : 2015GREAY036 . tel-01258803

HAL Id: tel-01258803

<https://theses.hal.science/tel-01258803>

Submitted on 19 Jan 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ GRENOBLE ALPES

Spécialité : **Physique subatomique & astroparticules**

Arrêté ministériel : 7 août 2006

Présentée par

Felix ROSENAU

Thèse dirigée par **Oliver ZIMMER**

préparée au sein de l'**Institut Laue Langevin**
dans l'**École Doctorale de physique**

HOPE - un piège magnétique pour neutron ultra-froid dédié à la mesure du temps de vie du neutron: conception et premières données expérimentales

Thèse soutenue publiquement le **10 juillet 2015**,
devant le jury composé de :

M. Dominique REBREYEND

DR CNRS au LPSC Grenoble, Président

M. Gilles BAN

Professeur ENSICAEN au LPC de Caen, Rapporteur

M. Peter FIERLINGER

Professeur à la TU Munich, Rapporteur

M. Oliver ZIMMER

Directeur du groupe NPP à l'ILL Grenoble, Membre



HOPE - a magnetic ultra-cold neutron trap to measure the neutron lifetime: design and first experimental data

Felix Rosenau

Abstract

The lifetime of the free neutron (τ_n) plays an important role in fundamental particle physics as well as cosmological models. Our knowledge of the precise value of τ_n is limited by the systematic inaccuracies of the two commonly used experimental approaches, the so called “beam” and “material bottle” methods. Moreover a systematic deviation of the τ_n -values extracted from both methods has become manifest over the past decades.

The HOPE project is part of a new generation of neutron lifetime experiments that aims to determine τ_n by storing *ultra cold neutrons* (UCN) in a combined magneto-gravitational potential. HOPE generates the necessary gradient magnet fields by a combination of highly potent rare-earth permanent magnets with a maximum B -field strength of about 1.3 T and a set of superconducting coils. In this thesis I give a detailed description of the apparatus, possible systematic effects and how we are planning to investigate and cope with those effects. Subsequently the results from a first beamtime at the PF2 source will be presented and discussed. The results are encouraging as we reached a maximum storage-time constant of $\tau_{stor} = 881(45)_{\text{stat}}$ s, indicating a very small UCN loss rate during storage.

Le temps de vie du neutron libre (τ_n) joue un rôle important dans la physique des particules comme dans des modèles cosmologiques. Notre connaissance de la valeur précise de τ_n est limitée par les incertitudes systématiques des deux méthodes expérimentales couramment utilisées, les méthodes dites de “faisceau” et de “bouteille matérielle”. En outre une déviation systématique des valeurs de τ_n obtenues par les deux méthodes s’est manifestée au cours des dernières décennies.

Le projet HOPE fait partie d’une nouvelle génération d’expériences qui cherchent à mesurer τ_n en stockant des neutrons ultra-froids (UCN) dans un potentiel magnéto-gravitationnel. HOPE génère les champs magnétiques nécessaires par une combinaison d’aimants permanents puissants de terre rare, qui produisent des champs magnétiques maximaux d’environ 1.3 T, et un système de bobines supraconductrices. Dans cette thèse je donnerai une description détaillée de l’appareil, des effets systématiques possibles et comment nous envisageons d’étudier et de maîtriser ces effets. Ensuite les résultats d’un premier temps de faisceau, effectué à la source PF2, seront présentés et discutés. Les résultats sont encourageants puisque nous avons atteint des constantes de temps de stockage maximaux $\tau_{stor} = 881(45)_{\text{stat}}$ s, ce qui indique des pertes d’UCN très faibles pendant la phase du stockage.

Contents

I. HOPE - a magneto-gravitational UCN trap to measure the neutron lifetime	7
1. The neutron and its lifetime in theory and experiment	8
1.1. What the neutron can tell us about the world	8
1.2. State of the art of τ_n -measurements	12
2. Design and systematical aspects of the experiment	19
2.1. Magnet system	21
2.1.1. Permanent magnet system	21
2.1.2. Superconducting coil system	30
2.1.3. UCN potential	33
2.2. Internal parts and experimental methods	37
2.2.1. Overview on systematic effects	37
2.2.2. Internal installations and UCN manipulation	46
2.2.3. Upgrades and variations	50
3. The 2014 beam time on PF2	58
3.1. Measurement scheme and parameters of interest	58
3.2. First data from HOPE	61
3.2.1. Filling rate, background rate and analysis scheme	62
3.2.2. Spectrum cleaning measurements	66
3.2.3. First attempts to determine the storage-time constant	71
3.3. Prospects	77
3.3.1. Measuring at 300 Amps	77
3.3.2. Statistical aspects	77
3.3.3. Systematic aspects - how to deal with the overcritical UCN	81
4. Conclusion	83
II. HOPE - un piège magnétique pour neutron ultra-froid dédié à la mesure du temps de vie du neutron	85
5. Le neutron et son temps de vie en théorie et expérience	86
5.1. Un outil varié pour explorer le monde	86

Contents

5.2. Etat actuel de la mesure de τ_n	88
6. Conception et aspects systématiques de l'expérience	90
6.1. Les aimants permanents	90
6.2. Les bobines supraconductrices	90
6.3. Composants intérieurs et méthodes expérimentales	91
6.3.1. Effets systématiques	91
6.3.2. L'absorbeur	92
6.3.3. Le piston	92
6.3.4. Le switch détecteur	92
7. Le temps de faisceau sur PF2	94
7.1. Procédures opératoires et paramètres d'intérêt	94
7.2. Premières données prises avec HOPE	94
7.2.1. Mesures du nettoyage du spectre	94
7.2.2. Mesures avec des temps de stockage longs	95
7.3. Perspectives	96
7.3.1. Aspects statistiques	96
7.3.2. Aspects systématiques – comment continuer avec les UCN sur- critiques	97
8. Conclusion	98
A. Cryostat	100
B. UCN source SUN2	108
C. Technical details on the 2014 beamtime on the EDM beamline of PF2	112
D. Monte-Carlo Simulations	118
E. Stray field measurements	124

Part I.

HOPE - a magneto-gravitational
UCN trap to measure the neutron
lifetime

1. The neutron and its lifetime in theory and experiment

1.1. What the neutron can tell us about the world

Neutrons have become an indispensable tool in a variety of scientific disciplines, mostly for their unique characteristics when used in scattering experiments. Being electrically neutral, they interact fundamentally different with matter than charged particles or electromagnetic radiation and therefore deliver unique insights when used as material probes. In experimental particle physics the neutron is also considered as a probe but in a more indirect way: as the neutron is one of the least complex non-elementary particles it has become an important test piece for the *standard model of particle physics* (SM) and - even more importantly - *beyond* standard model physics. While the SM is widely considered to be one of the most successful scientific theories that have ever been established, it is equally true that it is a limited theory. E.g. it fails to explain the matter-antimatter asymmetry in the universe and probably most prominently, gravity. Generally beyond standard model theories contain the SM as a low energy limit. Consequently the direct way to search for new physics is to go to higher energies where the effects of some more elementary theory should become directly visible. The *Large Hadron Collider* (LHC) is at the forefront of these efforts, a particle collider with a circumference of 27 km that accelerates protons up to 7 TeV before letting them collide in order to create the particles that are predicted by beyond SM theories.

In comparison, neutron physics experiments might seem rather modest but can absolutely compete in relevance with their big cousins. The measurement of the neutron's *electric dipole moment* (nEDM) is often referred to as the flagship of fundamental neutron physics and an archetype example for the indirect approach to beyond-SM theories. The neutron is not a point-like particle but has a finite spatial extent and an internal charge distribution and therefore can also have an EDM. Indeed, the existence of a permanent EDM violates T and therefore also CP-symmetry, a fundamental principle that states that a system's behavior should be invariant under the simultaneous inversion of its charge and parity. The SM allows for a small amount of CP-violation that would lead to an EDM on the order of 10^{-32} e cm. However, besides allowing the neutron to have an EDM, CP-violation has some more important implications: in the 1960 Russian physicist Andrei Sakharov formulated three basic requirements of any theory aiming to explain the matter-antimatter asymmetry in today's universe, one of them being CP-violation [36]. Without the asymmetry, most of the matter in the universe would have been annihilated shortly after the Big Bang and the universe would be a much more desert place than it is. The amount of CP-violation in the SM is far from being sufficient to create the

1. The neutron and its lifetime in theory and experiment




Figure 1.1.: Feynman diagrams of the neutron β -decay. The left plot shows the process on the quark level. Due to the very short range of the weak interaction the decay can also be described as a point-like 4-body process (right plot).

measured asymmetry. Thus most theories going beyond the SM contain new sources of CP-violation and therefore can be indirectly tested with nEDM searches. Long before the LHC has been turned on, the continuous improvement of the nEDM limit has already been one of the biggest constraints on *supersymmetric* expansions of the SM [21].

Fermi's interaction and the neutron β -decay

Likewise, the β -decay of the neutron is a test piece for fundamental particle physics. A first theory of the β -decay has been given by Enrico Fermi in 1934 [11]. In analogy to the electromagnetic interaction he described it as purely vector-like 4-body process. The decay of the neutron in the Fermi-model is shown on the right-hand side of figure 1.1: a neutron n transforms into a proton p under the emission of an electron e^- and an electron antineutrino $\bar{\nu}_e$.

In 1957, the *Wu experiment* measured the angular distribution of electrons emitted in the *nuclear β^- -decay* of a polarised ^{60}Co probe [45]. The measured distribution was strongly asymmetric as most of the electrons were emitted *against* the probe's polarisation direction. This was the first experimental proof of parity violation in the β -decay. Shortly after the *Goldhaber experiment* confirmed the results and stated that the parity violation was maximal [15]. They showed that all electron neutrinos emitted in the decay of $^{152\text{m}}\text{Eu}$ had left-handed helicity (i.e. their spins' projection is orientated against their direction of motion), right handed neutrinos do not exist in nature. Based on those findings Richard Feynman and Murray Gell-Mann [14], and in an independent effort George Sudarshan and Robert Marshak [42], established a new, parity-violating formulation of the Fermi interaction: the *V-A-theory* - short for "vector-minus-axial-vector-interaction-theory". Following the notation from reference [44], its *transition matrix element* M is given as:

$$M = [G_V \bar{p} \gamma_\mu n - G_A \bar{p} \gamma_5 \gamma_\mu n] [\bar{e} \gamma_\mu (1 - \gamma_5) \nu], \quad (1.1)$$

where \bar{p} , n , \bar{e} and ν are the Dirac spinors of the four particles, γ_μ are the Dirac matrices and G_V and G_A are the coupling constants of the vector- and axial-vector part of the

1. The neutron and its lifetime in theory and experiment

interaction. G_V and G_A are free parameters of the theory and thus have to be determined experimentally. As there is no direct way to measure them they can only be indirectly derived from other observables. One of those is the *lifetime of the free neutron* τ_n :

$$\tau_n^{-1} = \frac{2\pi}{\hbar} |M|^2 \rho, \quad (1.2)$$

where ρ is the phase space density of the final states and M is the matrix element from equation 1.1.

τ_n and the CKM matrix

Determining the precise values of G_V and G_A has implications that go well beyond the realm of the weak interaction. Today we know that the Fermi interaction is only an effective theory that can be used as low-energy limit of the *weak-interaction*. A Feynman diagram of the β -decay of the free neutron as described by the weak interaction is shown on the left-hand side of figure 1.1: The neutron is composed of one up and two down quarks and slightly heavier than the proton that has a *uud* quark composition. Under the emission of a W^- boson one of the neutron's down quarks transforms into an up quark. Subsequently the W^- decays into an electron e^- and an electron-anti-neutrino $\bar{\nu}_e$ yielding the net process:

$$n \longrightarrow p + e^- + \bar{\nu}_e + \Delta mc^2, \quad (1.3)$$

where Δm is the mass difference between the neutron and the end-products and $\Delta mc^2 = 0.782$ MeV [30]. The mass of the W^- boson $M_W = 80.4$ GeV/ c^2 is very large compared to the released energy and it can therefore only exist as a short-lived *virtual particle*. Consequently, the whole process effectively appears as a point-like Fermi interaction.

A characteristic of the weak interaction is that it does not only allow transitions between up and down quarks but also between the different *quark generations*. The SM knows six different quarks that are grouped in three generations:

$$\begin{pmatrix} u \\ d \end{pmatrix}, \quad \begin{pmatrix} c \\ s \end{pmatrix}, \quad \begin{pmatrix} t \\ b \end{pmatrix}. \quad (1.4)$$

Quarks are distinguished by their *flavour quantum number* and their mass, that varies widely between the three generations. Apart from that up, charm and top as well as down, strange and bottom quark all have precisely the same quantum numbers. This identity allows for a particularity: the quarks' flavour (weak) eigenstates $|q'\rangle$ are not identical to their mass eigenstates $|q\rangle$. Their mixing ratio is given by the so called Cabibbo-Kobayashi-Maskawa (CKM) matrix [18]:

$$\begin{pmatrix} |d'\rangle \\ |s'\rangle \\ |b'\rangle \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} |d\rangle \\ |s\rangle \\ |b\rangle \end{pmatrix}. \quad (1.5)$$

1. The neutron and its lifetime in theory and experiment

For the coupling constants of the neutron decay G_V and G_A this means that they are no fundamental but composed quantities:

$$G_V = g_V \cdot |V_{ud}|, \quad G_A = g_A \cdot |V_{ud}|. \quad (1.6)$$

g_V and g_A are the actual coupling constants of the weak interaction while the factor $|V_{ud}|$ accounts for the transition probability $P_{d \rightarrow u} = |V_{ud}|^2$ of a down quark becoming an up quark when undergoing a weak interaction. The formation of charm and top quarks is forbidden by energy conservation since the mass of either of them is large compared to the decay's total energy gain $\Delta E = \Delta mc^2$.

Within the standard model the matrix is assumed to be unitary:

$$\sum_{i=d,s,b} |V_{ui}|^2 = \sum_{i=d,s,b} |V_{ci}|^2 = \sum_{i=d,s,b} |V_{ti}|^2 = 1. \quad (1.7)$$

$\sum_{i=d,s,b} |V_{ui}|^2$ is the combined probability P of a down quark becoming either an up, a strange or a bottom quark when undergoing a weak interaction. If $P < 1$ this means that there is a non-zero probability of the down quark getting either annihilated or being transformed into some yet unknown fourth generation quark. In either case, finding significant deviations from the matrix's unitary would be a direct proof of new physics and a precise determination of its entries is an important test of the SM.

The 2014 *Particle Data Group* (PDG) value for $|V_{ud}|$ is 0.97425(22) [30]. It is determined from superallowed $0^+ \rightarrow 0^+$ nuclear beta decays. The value that can be extracted from neutron lifetime measurements is less precise but consistent with those findings. Including radiative corrections, $|V_{ud}|$ can be expressed in terms of τ_n and the ratio of the decay's coupling constants $\lambda = g_A/g_V$ as follows [27]:

$$|V_{ud}|^2 = \frac{4908.7(1.9) \text{ s}}{\tau_n (1 + 3\lambda^2)}. \quad (1.8)$$

Third party input on λ can be taken e.g. from measurements of the neutron decay angular correlation coefficient $A_0 = -2 \frac{\lambda(\lambda+1)}{1+3\lambda^2}$. Using the 2014 *Particle Data Group* (PDG) world averages of $\tau_n = 880.3 \pm 1.1 \text{ s}$ and $\lambda = -1.2723 \pm 0.0023$ we get $|V_{ud}| = 0.9758(2)_{\Delta rad}(6)_{\Delta \tau}(15)_{\Delta \lambda}$. A test of the results from the superallowed decays by the free neutron decay thus requires both, a more precise determination of λ and τ_n .


τ_n and the primordial nucleosynthesis

The neutron lifetime also plays an important role in the formation of light elements in the early universe, the so called primordial nucleosynthesis. About 150 s after the Big Bang the universe had become cold enough for neutrons and protons to permanently bind and form lighter elements up to ${}^7\text{Li}$. The respective abundances of the then formed nuclei is the result of a complex nuclear chain reaction and the neutron-to-proton ratio n/p at that moment. The relative abundances of the light elements in the universe are basically unchanged since the nucleosynthesis. A comparison of the theoretical predictions and the

1. The neutron and its lifetime in theory and experiment

relative abundances in today's universe thus allows for a consistency check of theory of the early universe. This test requires a precise knowledge of the n/p ratio at the moment of the nucleosynthesis. τ_n as well as the coupling constants of the weak interaction g_V and g_A are important parameters in the calculation of n/p .

Directly after the Big Bang n/p was given by the respective rates of the reactions



and thus followed a Boltzmann distribution:

$$n/p = \exp^{\Delta mc^2/k_B T}, \quad (1.10)$$

where Δm is the neutron-proton-mass difference and T is the temperature of the universe. About 1 second after the Big Bang the universe went through what is called the *nucleon freeze out*: the universe had become too cold to keep the reactions from 1.9 running and their rates dropped drastically. This happened at $n/p \approx 1/6$. However the exact ratio depends on the $n \leftrightarrow p$ reaction rates and thus indirectly on the coupling constants g_V and g_A . After the freeze out and the consequent stop of the $n \leftrightarrow p$ reactions n/p decreased slowly due to the beta-decay of the free neutron before reaching its final value of $\sim 1/7$ that was then frozen out by the nucleosynthesis. The currently limiting factor in the calculation of n/p is the experimental uncertainty of τ_n . A more detailed review on the topic can be found e.g. in [44].

1.2. State of the art of τ_n -measurements

The 2014 world average of the neutron lifetime as given by the particle data group is $\tau_{PDG} = 880.3(1.1)$ s [30]. An overview of the data used in the calculation of τ_{PDG} is shown in figure 1.2. The two predominant experimental methods - and the only ones used for the present τ_{PDG} - are *beam* and *material bottle* experiments. Both techniques have very different systematic effects and can therefore be considered as mutual benchmarks. Beam experiments determine τ_n by measuring the specific decay rate $R = \tau_n^{-1}$ of a neutron beam:

$$R = \frac{\dot{N}}{\rho_n V_D}, \quad (1.11)$$

where \dot{N} is the total decay rate, V_D is the *decay volume* - the section of the beamline where the decays can actually be detected - and ρ_n is the neutron density in V_D . From 1.11 we see that the determination of R implies an *absolute* measurement of \dot{N} and ρ_n as well as an exact determination of the fiducial volume V_D .

In bottle experiments neutrons are stored for long periods of time ($\mathcal{O}(T_{stor}) = \mathcal{O}(\tau_n)$) and τ_n is determined by measuring the *relative* decrease of the neutron number over time. Their primary observable is the trap's storage time τ_{stor} :

$$\tau_{stor}^{-1}(E_{tot}) = \tau_n^{-1} + \sum_i \tau_{loss,i}^{-1}(E_{tot}), \quad (1.12)$$

1. The neutron and its lifetime in theory and experiment




Figure 1.2.: Some experimental τ_n -values from the past 30 years. The circles designate *beam* and the squares *material bottle experiments* (see text for more details). The triangle labels a one of a kind measurement where the neutrons were stored in a “magnetic ring”. The diamond shows the very first competitive result from a new generation of storage experiments using magnetic instead of material bottles. The greyed out experiments are not considered in the calculation of the current (2014) averaged neutron lifetime value as given by the PDG, $\tau_{PDG} = 880.3(1.1)$ s. The shown experiments are: (1) 877(10) Paul 1989, (2) 887.6(3) Mampe 1989, (2') 882.5(2.1) Steyerl 2012 (*new analysis of the data from Mampe 1989*), (3) 893.6(5.3) Byrne 1990, (3') 889.2(4.8) Byrne 1996 (*improvement of the 1990 result*), (4) 882.6(2.7) Mampe 1993, (5) 885.4(1.0) Arzumanov 2000, (5') 881.6(2.1) Arzumanov 2012 (*new analysis of the data 2000*), (6) 886.3(3.4) Nico 2005, (6') 887.7(2.2) Yue 2013 (*improved flux calibration of the Nico 2005 experiment*), (7) 878.5(0.8) Serebrov 2005, (8) 880.7(1.8) Pichlmaier 2010, (9) 878.3(1.9) Ezhov 2014. All information taken from [30] except on the Ezhov 2014 result [10].

1. The neutron and its lifetime in theory and experiment

where the neutron's *total energy* E_{tot} is defined as the neutron's kinetic energy at the trap's potential minimum. We see that τ_n^{-1} is only one addend in 1.12. The second one is a sum over all other possible loss channels $\tau_{loss,i}^{-1}$. As indicated those losses do not only depend on the traps' characteristics but also on E_{tot} and therefore on the spectrum of the stored neutrons. Understanding and limiting those loss channels is the main experimental challenge of the storage method.

Over the last few years the increasing comprehension of those losses has been used to reanalyse past experiments and ease off the controversy about the formerly irreconcilable Arzumanov and Serebrov experiments [2, 3, 38]. Today, all relevant material bottle experiments deliver neutron-lifetime constants within a small interval around τ_{PDG} . In contrast the data from the two most precise beam experiments is conspicuously shifted to higher values. A revision of the 2005 result from Nico et al. has further intensified the problem [29, 46]. New data, especially when drawn from a third, systematically independent technique would be of great help to decide whether the offset is just a statistical fluctuation or due to yet uncontrolled systematic error sources.

HOPE is part of a new generation of neutron lifetime experiments, the *magnetic or magneto-gravitational neutron traps*, that aims to deliver this independent data. Conceptually, magnetic traps are relatively close to material traps but they do not share their major drawback, the inevitable *neutron losses due to wall collisions*. Magnetic neutron storage is a promising new access to τ_n but it still is a comparably recent technique and for a long time the only competitive experimental input came from the *NESTOR* experiment (no. 1 in figure 1.2). *NESTOR* was a magnetic storage ring that stored neutrons by holding them on (relatively) stable circular orbits [32]. In contrast to newer magnetic bottles the neutrons were not permanently stored but could leave the trap through the magnetic walls once they left their trajectories. Only recently a new neutron lifetime value has been published by V. Ezhov et al., $\tau_{Ezhov} = 878.3(1.9)$ s [10], that is in good accordance with the lifetime values from former material bottle experiments.

In the following I will briefly revise the basic experimental concepts of magnetic and material bottle experiments. Since both methods store so called *Ultra Cold Neutrons* (UCN) I will start with a short introduction to UCN. A good overview on the topic can be found e.g. in [16]. For more information on the neutron lifetime I refer to [44].

Ultra Cold Neutrons

When a neutron hits a wall it will primarily interact with the cores of the wall atoms due to strong interaction. The interaction potential is on the order of 10 MeV but also very short ranged (fm-scale). Enrico Fermi proposed to exchange the strong but short ranged force by an effective pseudo-potential with a longer range but a smaller interaction potential [12]. This new potential can now be treated by perturbation theory and in the first Born approximation gets:

$$U_F(\mathbf{r}) = \frac{2\pi\hbar^2 a}{\mu} \delta^{(3)}(\mathbf{r}), \quad (1.13)$$

1. The neutron and its lifetime in theory and experiment

where a is the scattering length of the nuclei and μ the reduced mass $\frac{m_{nuclei}m_n}{m_{nuclei}+m_n}$. Using $U_F(\mathbf{r})$ it is now possible to give an effective *optical potential* $U_{opt}(\mathbf{r})$ for the collective interaction of a neutron with matter:

$$U_{opt}(\mathbf{r}) = V_{opt} + iW, \quad (1.14)$$

with

$$V_{opt} = \frac{2\pi\hbar^2}{m_n} \sum_i n_i(\mathbf{r})b_i \quad W = \frac{\hbar}{2} \sum_i n_i(\mathbf{r})\sigma_{loss,i}(v_n)v_n. \quad (1.15)$$

In 1.15 n_i designates the number density of the isotope i , b_i the bound scattering length of the said isotope and $\sigma_{loss,i}$ its velocity dependent neutron absorption cross section. V_{opt} can be identified as the effective potential barrier of the material¹ - it can be positive *or* negative depending on the scattering lengths of the materials - and W can be used to calculate the neutron loss probability per bounce. A suitable material for neutron storage must combine both a high V_{opt} and a small W . Popular materials for neutron storage are e.g. Beryllium with a potential barrier of 252 neV and a loss factor $f = W/V$ of 0.5×10^{-5} or ^{58}Ni ($V_{opt} = 335$ neV, $f = 8.6 \times 10^{-5}$ [16]). Obviously even a “high” Fermi-potential still is very low in absolute terms. For comparison the mean energy of a cold neutron beam that has been thermalised at 20 K is of 2 meV, a factor 10^4 higher than the typical Fermi-potentials.

Besides the strong interaction and the resulting Fermi-potential there are two other fundamental interactions that equally apply forces in the neV-range to the neutron: the electromagnetic interaction and gravity.

The neutron does not carry a charge but it has a magnetic moment $\mu_n = -1.913\mu_N$, where μ_N is the nuclear magneton, and therefore interacts with external magnetic fields. The resulting magnetic potential is

$$V_{mag} = -\vec{\mu}\vec{B} = \pm 60.3 \text{ neV T}^{-1} \cdot \vec{B}. \quad (1.16)$$

Since the sign of V_{mag} is determined by the neutron’s polarisation - the relative direction of its spin and the magnetic field - UCN in a B -field are commonly divided into *low* and *high field seeking* neutrons (LFS and HFS). LFS neutrons have their spin oriented *antiparallel* to the magnetic field and are repelled by increasing gradient B -fields while HFS neutrons are “attracted” by magnetic gradients. Only LFS neutrons can be permanently trapped in a magnetic bottle. Half of the UCN is lost by default due to their inappropriate polarisation state.

The gravitational potential experienced by a neutron is of

$$V_{grv} = m_n g h = 102 \text{ neV m}^{-1} \cdot h. \quad (1.17)$$

For UCN physics V_{grv} can be both, an asset and a handicap. In combination with V_{opt} and V_{mag} it can be used to construct a hybrid material/magneto-gravitational trap that

¹Throughout this thesis when I talk about the *Fermi-potential* of a material I will generally refer to V_{opt} .

1. The neutron and its lifetime in theory and experiment

is “closed” by gravity on top. Also, the gravitational force can be used to manipulate the UCN spectrum, e.g. by adapting the relative height of UCN source and experiment or by varying the height of the guide system in order to sort out UCN in a certain energy range.

Negative aspects are the reduction of the accessible phase space and thereby of the total UCN number that can be filled into a trap. In section 2.1.3 we will see that in the case of HOPE gravity reduces the final UCN number by about 50 % when passing from a horizontal to a vertical trap orientation. Moreover the strong influence of gravity, in combination with some other disadvantageous characteristics, considerably complicates the theoretical description of UCN experiments: UCN do not behave like a conventional gas and in most cases cannot be described by the means of analytical tools derived from statistical mechanics (see e.g. [16] or section D of this script). In most cases only much more time-consuming particle-tracking simulations deliver satisfying results.

The major drawback of UCN physics (i.e. experimental physics in the energy range of the Fermi-potential) are the generally very low count rates. The most successful UCN source in the world, the PF2 source at the ILL, basically uses the *ultra-low* energy tail of the thermalised neutron spectrum coming from a cold source. It delivers UCN densities of 5 - 50 UCN cm⁻³ [41, 39]. A new generation of *superthermal*² UCN sources promises to increase those numbers by one to two orders of magnitude. In the appendix of this script one of those sources, SUN2, is briefly discussed. More information about UCN sources can be found e.g. in [24, 48].

Fill-and-empty vs. counting-the-dead type measurements

Storage experiments can be divided into two groups, depending on *how* the decay of the trapped UCN is tracked.

Counting-the-dead type experiments detect the neutrons’ decay products and τ_n is determined from an exponential fit to the decaying signal rate. Generally this method allows to determine τ_n from a single *fill*, repeated filling is only necessary in order to increase the statistical accuracy. The method is almost exclusively reserved for magnetic traps like HOPE as the magnetic fields can also be used to guide the decay products to a detector. Without those guiding fields the detection probability drops tremendously as the total surface of the trap’s side walls will generally be much bigger than the effective surface of the used electron or proton detector.

The second method is the much more common *fill-and-empty* measurement (counting the survivors): once the trap has been *filled* the UCN are stored for a time t_i . When the time is up the trap is opened and *empties* itself into an UCN detector. By repeating the procedure with different t_i one gets several measurement points on the neutrons’ decay curve and can once again perform an exponential fit. The measurements at the different t_i have to be repeated (many times) in order to reach a significant statistical accuracy.

²I.e. a technique that is not based on thermalisation but some different physical phenomena that allows for a cooling *below* the temperature of the surrounding material.

1. The neutron and its lifetime in theory and experiment

The fill-and-empty method requires an additional monitoring of the UCN filling rate in order to check for systematic drifts of the initial neutron number.

Apart from the measurement method, storage experiments can also be characterised by their predominant loss channels.

Material traps

The major drawback of material traps are losses occurring during wall collisions. The optical potential already includes wall losses through the imaginary part in equation 1.14. From there it is possible to calculate an energy dependent loss probability per bounce $P_{wall}(E_{tot})$ that can be used to calculate the τ_n from equation 1.12. Unfortunately this easy solution proved to be inappropriate to describe the experimental results as the actual wall-loss probability appears to be considerably higher than predicted by theory (see e.g. [37]). The reason of this discrepancy is not yet conclusively established. Probably it is a combination of several effects namely impurities in the wall material, contamination on the walls' surfaces or cracks and tiny holes in the wall. Since these *anomalous losses* can not be handled theoretically, an empirical approach is necessary:

The wall loss probability per unit time $\tau_{wall}^{-1}(E_{tot})$ can be written as:

$$\tau_{wall}^{-1}(E_{tot}) = f_{wall}(E_{tot}) \cdot P_{wall}(E_{tot}), \quad (1.18)$$

with:

$$f_{wall} = \frac{v}{l} \approx v \frac{A}{4V}. \quad (1.19)$$

A and V are the trap's total surface and volume, l is the UCNs' mean free path and v their velocity. For an ideal gas $l = \frac{4V}{A}$. Due to gravity UCN do not behave like an ideal gas and 1.19 is only an approximation. We see that by altering A/V we can manipulate the wall-collision frequency f_{wall} and thereby τ_{wall}^{-1} . In the ideal case of an infinite storage volume A/V , f_{wall} as well as τ_{wall}^{-1} become 0. Obviously it is not possible to construct an infinite storage volume but it *is possible* to vary A/V , e.g. by displacing the sidewalls of the trapping volume. Thus the general procedure will be to measure τ_{stor}^{-1} as a function of A/V and to *linearly extrapolate* the data to the storage time of a virtually infinite trap *without wall losses*.

However the extrapolation is no trivial operation and, especially for long extrapolation times (e.g. 100 - 120 s in [2] or up to 150 s in [34], entry 5 respectively 8 in figure 1.2), limits the experiments' absolute accuracy. Today the longest storage time in a material trap, and thereby also the shortest extrapolation time, has been achieved by A. Serebrov et al. and their *Gravitrap* experiment ([38], entry 7 in 1.2). By covering their walls with Fomblin grease and cooling the whole trap down to ~ 110 K they reached maximum storage times $\tau_{stor} = 872.2 \pm 0.3$ s, only about 5 s less than their lifetime value $\tau_n = 878.5 \pm 0.8$ s. This still is the most precise measurement of τ_n .

Magnetic traps

Magnetic traps store UCN in a gradient B -field. A subspecies are the *magneto-gravitational* traps as HOPE that use gravity to "close" the trap in vertical direction. In both cases

1. The neutron and its lifetime in theory and experiment

the point is to *exclude* wall collisions and thereby avoid the extrapolation step in the determination of τ_n . Besides the inevitable residual gas collisions magnetic traps have an all new loss channel since the neutrons will only remain trapped as long as they keep their polarisation. However both loss channels should result in rather small systematic corrections $\Delta\tau_{syst} < 0.1$ s (see e.g. section 2.2.1 or [10]) and thus allow for a “systematically clean” determination of τ_n with prospects to overall uncertainties on the order of 0.1 s. Recently V. Ezhov et al. demonstrated the potential of magnetic traps when they published their first results giving a maximum storage time of $\tau_{stor} = 874.6 \pm 1.7$ s and a derived neutron lifetime value of $\tau_{Ezhov} = 878.3 \pm 1.9$ s.

Overcritical UCN

A very critical point in storage experiments is the *spectrum cleaning*, i.e. the removal of UCN with a total energy $E_{tot} > E_{max}$, where E_{max} is the trap’s cut-off energy (e.g. the wall’s Fermi-potential in case of a material trap). Without a thorough cleaning procedure these *overcritical UCN* can cause a significant shift of the mean storage-time constant. The major effect on magnetic traps has been unintendedly demonstrated by an American collaboration that stored their UCN in a 3-dimensional Ioffe-Pritchard-type trap [8]. Their first measurements yielded a lifetime value $\tau_{NIST} = 621^{+18}_{-17}$ s, substantially shorter than τ_{PDG} .

They accredited this discrepancy to so called *marginally trapped UCN*, overcritical neutrons with storage times in the same order as τ_n itself. In a magnetic trap the wall-collision frequency of overcritical UCN will be reduced as, dependent on the angle of incidence, they will be reflected by the gradient magnetic field before touching the wall. Even if they hit the wall the magnetic potential at the point of the collision decreases their impact energy and thereby further reduces their wall loss probability. In combination both effects can explain the longevity of the overcritical UCN and the dramatic shift of the mean storage time. By expanding their measurement protocol by a dedicated spectrum cleaning phase - they lowered the magnetic field in the beginning of each cycle before ramping it up again and starting the actual storage phase - they could reduce the number of overcritical UCN and could extract a lifetime value of $\tau'_{NIST} = 833^{+74}_{-63}$ s.

In their case the overcritical UCN were stored in a *magneto-material hybrid trap* resulting in especially long overcritical storage time constants. As we will see in section 2.2.1 magnetic traps are generally (much) more susceptible to marginally trapped UCN than material traps. Dealing with this problem will definitely be one of the bigger challenges of upcoming experiments.

2. Design and systematical aspects of the experiment

HOPE - the *H*albach-type *O*ctupole *P*ermanent magnet *U*CN trap - is a magneto-gravitational UCN trap designed to measure the lifetime of the free neutron. The project has already been proposed by Oliver Zimmer some 15 years ago [47]. The overall concept of the experiment was worked out by Kent Leung who also started working on HOPE's name-giving octupole permanent magnet system [24, 26]. His thesis contains numerous thoughts on systematic effects and the design of the experiment that I will refer to recurrently throughout this script. I took over his work in 2011, finalised the experiment's design and could perform a first measurement cycle on the PF2 source in 2014. The results from those measurements will be presented in the subsequent chapter. In the following I will focus on the trap's design and the experimental concepts.

Figure 2.1 shows a cut through a SolidWorks model of the apparatus. The “heart” of the experiment, the actual trapping region (1), is approximated by the light blue oval. Its position and form is defined by the permanent magnet system (PMS, 2), the superconducting coil system (SCS, 3) and gravity. While most of the UCN potential will remain static throughout a single measurement cycle the bottom coil of the SCS is designed to allow rapid ramping and will serve as a magnetic valve during the UCN filling and emptying. The UCN will be guided to the experiment via an horizontal stainless steel tube (6) that ends in a T-piece located some 60 cm beneath the potential minimum of the trapping region. The lower side of the T-piece is connected to a slide valve (7) that allows switching between a “filling” and a “counting” position. In the counting position the switch is open towards a He3 UCN detector while, in the filling position, it allows to enter a Teflon plug (5) into the trapping region. The primary idea behind the plug is to speed up the UCN spectrum cleaning by triggering UCN mode mixing via diffuse reflection. During the cleaning process a UCN absorber (4) will be lowered into the trapping region and retracted afterward. Electron-track simulations have been performed in order to assure that a permanently installed absorber on top of the trapping volume is compatible with the future installation of an electron or proton detector on top of the PMS. The whole system is embedded in a helium-bath cryostat (8) manufactured by AS Scientific. In the following chapter I will give a detailed description of all the components mentioned above and demonstrate how the chosen technical solutions respond to the difficulties of a neutron-lifetime measurement.

2. Design and systematical aspects of the experiment




Figure 2.1.: Cut through a SolidWorks[®] model of the HOPE experiment. The numbered parts are: 1: trapping volume 2: permanent magnets 3: superconducting coils (SCS) 4: UCN absorber 5: Teflon plug 6: UCN entrance 7: detector switch (shown in storing/emptying position) 8: LHe bath. The origin of the attached z-axis sits on the mid-plane of the SCS' bottom coil and thus close to the trapping region's lowest point. If not stated otherwise this scale is used throughout the thesis.

2.1. Magnet system

The magnet system has already been extensively discussed in the Kent Leung’s thesis [24]. I will only summarise the essential information and supplement some technical details.

2.1.1. Permanent magnet system

The permanent magnet system is the “heart” of HOPE as it provides the predominant part of the magnetic barrier. The PMS is a 1.2 m long cylinder with an outer diameter of about 20 cm. The whole cylinder is traversed by a concentric bore with an inner radius $\rho_{PMS} = 46.8$ mm that represents the actual trapping region. Apart from regions close to the end planes and the side walls, the B -field strength $|B|$ inside the bore can be approximated by the formula

$$|B(\rho)| = |B_{\perp}(\rho)| = |B_{max}| \left(\frac{\rho}{\rho_{PMS}} \right)^3, \quad (2.1)$$

where $|B_{\perp}|$ is the portion of the magnetic field perpendicular to the bore’s axis, $|B_{max}|$ is the maximal B -field strength and ρ is the radial position. The maximal field strength on the bore walls is about 1.3 T, the z and φ dependency of $|B|$ can be neglected as long as we do not approach the side walls or end planes. In sum, the B -field within the bore points perpendicular to the bore’s axis, reaches its maximum strength on the side walls, has a zero-crossing on its axis and follows a ρ^3 -law in between.

The technical implementation is carried out in the form of a *Halbach-type array* [17]. An illustration of its structure and characteristics is shown on the left of figure 2.2. The pattern combines two favorable properties as the field lines of the single magnets superimpose each other constructively on one side and destructively on the other. While the reinforcing effect on the inside of the bore increases our trap’s storage potential the cancellation on the outside are of practical interest as it would be very restraining to work in the constant presence of a strong magnetic field. The right side of the figure shows a slice of our permanent magnet system and illustrates the arrangement of the individual magnet pieces and how they add up in order to form the desired octupole field on its inside. In total the PMS consists of 12 such modules arranged in a pile, each housed in a stainless steel jacket that acts as counteracting force against their mutual magnetic repulsion.

The optimal pole order is given by the PMS’ dimensions: if we generalise equation 2.1 to the case of an arbitrary number of poles N we get the relation¹

$$|B(\rho)| = |B_{max}| \left(\frac{\rho}{\rho_{max}} \right)^{\frac{N}{2}-1}. \quad (2.2)$$

With 10 poles we would thus have an exponent of 4, with 12 poles an exponent of 5 and so on. The phase-space accessible to the UCN increases with the exponent and we could

¹The derivation can be found e.g in [5].

2. Design and systematical aspects of the experiment



Figure 2.2.: The left hand side shows a schematic illustration of a Halbach-type array. The array consists of single permanent magnet pieces with different magnetisation directions. The particular pattern is chosen in such a way that the field lines of the single pieces will reinforce each other on one side of the magnets while they are canceling each other out on the other side. The picture on the right hand side shows one of our permanent magnet modules. It consists of a total of 32 single magnets. The magnetisation direction of the single magnets is indicated by the white arrows while the gray lines denote course of the magnetic field lines as they have been obtained from some FEMM simulations. Compared to the schematic the pattern contains additional transition pieces as the magnetisation direction from one piece to another only changes by 45° . Therefore the 32 pieces add up to form 8 virtual magnetic poles resulting in the eponymus octupole field inside the bore.

2. Design and systematical aspects of the experiment

store more neutrons if our trap had a higher pole order. However, if we further approach magnet slices with opposing magnetisation directions, mutual field-cancellation effects will increase, diminish $|B_{max}|$ and outweigh the positive effects of the higher pole order.

A supplementary problem is the mutual demagnetisation of the individual magnet pieces. The magnets are so called *Vacodym* magnets purchased from *Vacuumschmelze GmbH* [43]. They are rare-earth alloys that are based on the chemical structure $\text{Nd}_2\text{Fe}_{14}\text{B}$ and offer a wide range of variations in their characteristic values such as their remanence B -field or coercivity. For our magnet system two different alloys (Vacodym 745HR and 633HR) have been used in order to maximise the system's overall performance. While the 745HR-type magnets have a slightly higher remanence field, the 633HR-alloy has a higher coercivity and was chosen for the *off-pole* pieces (the pieces with a magnetisation tangential to the bore) where the magnet's autodemagnetisation gets maximal. Still FEMM simulations predict that there should be some demagnetisation on the inner corners of the off-pole magnet pieces [28, 24]. Those effects should be negligible and until now all our experience confirms that they are: Neither has Kent seen major weak points in the PMS field when he measured it using a *LakeShore* Hall probe. Nor have we measured an otherwise unexplainable low storage-time constant during the 2014 measurements (see section 3.2.3). However the simulations show that we are reaching the materials limits.

Cooling the magnets

We also have to consider the demagnetising effect of the superconducting coils on the permanent magnet system: without further precaution, we would definitely enter a regime of irreversible deformation of the PMS' magnetic field when we ramp the shutter coil up to its maximum value [24]. Obviously we can not shield the magnets from the coils' field as this would also weaken the UCN potential. Fortunately there is an alternative possibility since the permanent magnets' coercivity is inversely proportional to their temperature. Additionally the magnets' remanence field has a maximum somewhere between 100 and 150 K depending on the chosen alloy [22]. At the temperature of its field maximum the magnet's coercivity should be sufficient to prevent any demagnetising effects from the superconducting coils. Thus we opted for an adjustable cooling mechanism that would allow us to tune the final temperature to an optimal working point.

Figure 2.1 shows that the PMS is surrounded by a liquid Helium bath and it is tempting to use the bath to cool the PMS. However there are numerous practical considerations that veto such a solution. The most prominent is that, in the case of a coupled system, any maintenance work on the PMS would require the warm-up of the entire cryostat. Furthermore the cooling by the cryostat would be unregulated and we would still need a tool to fine-tune the PMS' final temperature.

Therefore we decided to use an independent cooling system for the PMS: a LN2 circuit made from 6 mm diameter copper tubes that we wrapped around the PMS' outer aluminum jacket. In order to finalise the procedure I performed a couple of test cool-downs, together with my trainee Yoann Kermaidic, in April/May 2012. Within a few iterations we managed to sufficiently stabilise the system's temperature by using a simple *P-loop mechanism*: We regulated the nitrogen flow via a standard on/off LN2 valve that was op-

2. Design and systematical aspects of the experiment

erated by a *LakeShore* 218S temperature monitor [19]. The 218S is equipped with eight relays that are rated up to 30 VDC² and 5 A and can be controlled via any of the eight temperature inputs. We operated the relays in low-alarm mode. The alarm is turned on once the temperature T falls below some set temperature T_{set} and turned off again once T rises above T_{set} plus some predefined *deadband* value dT_{db} . A non-zero deadband avoids very short switching cycles that would only shorten the valves lifetime without leading to a considerably higher temperature stability. Figure 2.3 shows some photos of the cooling system right before the PMS' first installation in the cryostat. Compared to the test runs I only implemented a few notable changes:

- In 2012 the cryostat did not yet exist and we used a 1.5 to 2 m long DN250 ISO-K tube with an inner diameter of 261 mm as a substitute vacuum chamber.
- We use two copper spirals, one in the upper and one lower third of the PMS. During the test runs both spirals were connected to the same valve and the LN2 flow in both was always synchronous. I later decided to use two separate valves and allow for an independent usage of both spirals in order to minimise temperature gradients along the PMS' axis.
- During the test runs we used simple flex lines for the LN2 transfer from the dewar to the vacuum chamber. Those were replaced by an isolated He-transfer line from *AS Scientific* in order to limit the heating of the line in between the cooling cycles and reduce the LN2 consumption.
- Aside from the three PT-100 sensors shown on photo we had three additional sensors on the PMS' inner side wall during the test cool-downs.

The plots shown in figure 2.4 summarise the outcome of our cooling tests. In the upper plot we can see the initial cool down. During the test runs in 2012 the top and bottom PT-100 sensors were not mounted as shown in figure 2.3 but directly next to the top respectively the bottom spiral. Consequently the top_{AI} (dark blue curve) and $bottom_{AI}$ (green curve) sensors react immediately once the LN2 flow is started while the $middle_{AI}$ sensor (purple) and the three sensors inside the middle bore follow with varying delays. In this phase we get maximum temperature gradients of about 70 K measured between the top_{AI} and the $middle_{Mag}$ sensor. At $t \approx 5000$ s the top_{AI} curve starts falling faster than the $bottom_{AI}$ curve. The different cooling power of the two spirals is due to an asymmetry in the set-up. As the common LN2 entrance was on top of the vacuum chamber, the travel path to the bottom spiral was longer than the path the top spiral. Therefore the LN2 has already lost part of its cooling power before reaching the bottom spiral and/or the LN2 flux in the lower spiral is lower due to the higher flow resistance.

We had chosen to control our P-loop via the $bottom_{AI}$ sensor which reaches T_{set} about 7.5 hours after the cooling starts. The last sensor, Mid_{Mag} , reaches the set temperature about 4 hours later. In the meantime the top_{AI} sensor showed an underswing followed

²The electrovalves we use are 24 VAC but the relays work just as fine with alternate as with direct current.

2. Design and systematical aspects of the experiment



Figure 2.3.: The PMS before its installation in the cryostat. We equipped the PMS with two independent copper spirals for the LN₂ cooling. We have only three feedthroughs for the two spirals as their exit is short-cut (right hand side, top). The red circles show the positions of the PT100 temperature sensors we use as feedback for our P-loop temperature control. The spirals were held in place by the two copper sheets we wrapped around the PMS and fixed by several collars (bottom right). If we had mounted the collars directly on the spirals their fixing screws would have stuck out too far and touched the cryostat's inner bore.

2. Design and systematical aspects of the experiment

by saturation. About 12 hours after the cooling was started the system reaches stable working conditions.

Once the system is stabilised we profit from its high thermal inertia as can be seen in the two plots in the middle row of figure 2.4. The two plots show phases with stabilised temperatures with slightly different T_{set} and dT_{db} values. The different dT_{db} -values result in significantly different valve switching intervals. In the case of $dT_{db} = 0.5$ K the interval gets long enough for the LN2 in the transfer line to vaporise and cause the “heating boost” in the beginning of each cooling cycle.

Directly on the magnets the temperature fluctuations are much smaller (< 0.1 K). At ~ 140 K the temperature dependence of the magnet’s remanence field is $\lesssim 10^{-3}$ T K⁻¹ [22]³. Using equation 1.16 - $V_{mag} = 60.3$ neV T⁻¹ - this yields maximum fluctuations of the UCN potential on the order of 10^{-2} neV.

The bottom plot in figure 2.4 shows the PMS’ warm-up once the LN2 flow is stopped. We did not include any heating device and the warm-up thus is solely due to radiation and convection. In contrast to the cool-down, no significant temperature gradients build up, only the middle_{Mag} sensor reacts with a little delay as it is much less affected by radiative heating. However the overall warm-up time is considerably longer than the cool-down as its driving force, the temperature gradient between the surroundings and the PMS, naturally decreases with rising temperature. Starting at 140 K it takes about 40 hours to reach 20°C. As the permanent magnets are very corrosive it is strongly counter indicated to break the vacuum at lower temperatures in order to avoid dew formation.

We were very satisfied with the results as we had found a simple and relatively efficient way to cool the PMS and stabilise its temperature in the desired range. Moreover the cooling should get even more efficient inside the cryostat since the total heat input will be reduced.

Ironically the aforementioned modifications accompanied by the changed thermal conditions created an altogether less favorable situation. Figure 2.5 shows some temperature data from the 2014 beamtime. It indicates that both, the static temperature gradient as well as the temperature fluctuations, were bigger during the beamtime than during the 2012 test runs. The big static gradient between the top and the two other sensors suggests an imbalance between the thermal heat load on the top and bottom side of the PMS. As before the temperature fluctuations measured with the bottom sensor are caused by the warm-up of the LN2 transfer line. Due to the reduced heat load inside the cryostat the cooling intervals expanded up to almost 30 minutes. In the meantime even our custom made He-transfer line unfreezes and must be re-cooled in the beginning of each cooling cycle.

The situation is not dramatic but some small modifications could be implemented prior to future data taking. During the beamtime dT_{db} was set to 0.3 K. By reducing it - or even setting it to 0 - the temperature fluctuations can be easily reduced. The static gradient could be decreased by moving the top spiral further towards the PMS’ top end. Finally it could help to reposition the PT100 sensors. The arrangement shown in figure 2.3 -

³We are planning to determine the exact amplitude during our prospected magnetic-field measurements (see section 2.2.3).

2. Design and systematical aspects of the experiment




Figure 2.4.: The four plots show the cooling of the permanent magnets, two phase of temperature stabilisation with different settings and the warm-up once the LN2 cooling stopped. The ending *Al* and *Mag* in legend stand for aluminum and magnet and distinguish between the temperature sensors that were mounted on the outer aluminum shell and the ones that directly stuck to the magnets. T_{set} and dT_{db} are defined in the text. The spikes in the top magnets warming curve are due to a bad electrical contact.

2. Design and systematical aspects of the experiment




Figure 2.5.: PMS temperatures during the 2014 beam time. The curves correspond to the three PT100 temperature probes shown in figure 2.3. Despite the virtually improved system - independently operated spirals, lower heat input, etc. - it's overall performance is slightly deteriorated in terms of total temperature gradient and fluctuation amplitude as compared to the 2012 cooling tests.

with all three sensors sitting rather far away from the spirals - was rather unfortunate. Due to the system's inertia the sensors' reactions were delayed and prevented a smooth operation of the feedback loop.

Mounting the magnets

Mounting the PMS' within the cryostat was one of the most challenging tasks. The design had to be adapted to experimental as well as technical constraints:

- In order to limit UCN losses due to residual-gas collisions we must reach final gas pressures $< 10^{-6}$ mbar (see section 2.2.1). It therefore seemed useful to separate the PMS with its glued magnets from the UCN storage volume by a *bore tube*.
- The importance of the inner diameter of the *bore tube* is a subject of discussion. The trap's cut-off energy is defined by the magnetic field strength on the tube's inner wall which has a very steep radial dependency ($|B(\rho)| \propto \rho^3$). Thus it might seem appropriate to maximise the tube's outer diameter. However, in his thesis, Kent proposes to deliberately reduce the tube's diameter in order to add a thermal isolation and active heating equipment in between the tube and the PMS [24]. With a warm bore tube the experimental vacuum would be accessible at all times, *without* having to warm up the PMS first. I decided to accept the disadvantages of a cold bore tube and maximise its diameter instead. My main concern was the technical challenge of maintaining the temperature gradient between the PMS and the bore tube (~ 180 K mm^{-1}) over the entire inner surface of the PMS.

2. Design and systematical aspects of the experiment

- In order to minimise UCN back reflections on their way in and/or out the trap we paid closest attention to avoid steps in the lower part of the experimental vacuum chamber.
- The permanent magnet system as it is shown in figure 2.3 - including the bore tube and the extension of the experimental vacuum on top - measures some 2.5 m. This installation has to be connected to the cryostat via vacuum-tight seals on its two extremities in order to create the isolation vacuum for the PMS. As high precision is needed at the vacuum seals the suspension of the PMS has to be flexible in order to compensate for possible angular and axial misalignments.
- The PMS cooling will also affect the bore tube. The resulting thermal contractions are compensated by the bellows shown in figure 2.6. The bellow's flange is equipped with an Indium seal in order to allow for the low temperatures at this point.

In the following I will briefly review the most interesting and tricky parts of the final design.

The bore tube is a pleated steel sheet with 0.5 mm wall-thickness and a longitudinal welding seam. It is made from 316L stainless steel which combines the three favourable characteristics of having a very low magnetic susceptibility, a low thermal conductivity and a rather high Fermi potential (185 neV). Due to the tube's anticipated ovality and the bore's irregularities - rather than a circle it is a polyhedron with regular steps at the junctions of the individual magnet pieces - it was difficult to estimate the tube's maximal outer diameter. I fixed it to $92.9^{0}_{-0.5}$ mm, leaving a $0.7^{0}_{-0.5}$ mm clearance compared to the bore's inner diameter. The final tube "just fits" into the bore, meaning that it sits tight but can still be moved without (brute) force.

Figure 2.6 allows to discuss some technical aspects of the PMS fixation. It shows an overview of the PMS set-up including close-ups of the regions of interest. The PMS' only supporting connection are the 12 M6 rods that link them to the (dark-green) upper tube (see detail 2 and upper picture). The rods shown in detail 3 have no supporting function. The circle they are holding enters the groove in the flange without touching its ground and centralises the PMS in radial direction. Detail 2 also shows that the bore tube is not directly connected to the PMS but only to the (turquoise) bellows. Thus the PMS' vertical position is not predefined but can be tuned via the M6 rods.

Detail 1 shows the suspension mechanism of the (dark-green) tube on top of the cryostat. The upmost flange (A) belongs to the tube and supports the entire weight of the magnet system. The flange sits on six M10×1.0 thread rods (turquoise) with rounded heads placed in six cavities in flange C. This system allows us to compensate for angular misalignments. The flange C establishes a vacuum-tight connection to the cryostat while flange B is necessary to seal the vacuum between the tube and flange C.

On its top-end side the bore tube is equipped with a welded flange that connects the tube to the bellow (detail 2). A second preinstalled flange on its bottom end was impractical - at least if we did not want to reduce the tube's inner diameter - as we still had to push it through the PMS. Welding on the tube once it had been installed was no alternative either and we finally decided to go for a glued connection using *Stycast*

2. Design and systematical aspects of the experiment




Figure 2.6.: Composition of 3D drawings and pictures showing technical details of how the PMS is suspended and centered inside the cryostat. Details on the individual drawings can be found in the text.

2850FT as adhesive and *Catalyst 23 LV* as curing agent. The right picture in figure 2.6 shows the tubular adapter piece already glued to the bore tube. The two rings that sit in between the adapter piece and the PMS have been added to avoid possible deformations in the case of unidirectional atmospheric pressure. On detail 3 the adapter piece is shown in dark-blue. As on the cryostat's top side the vacuum connection to the major flange is done via a radially compressed O-ring.

2.1.2. Superconducting coil system

In the following I will describe the coil system and its characteristics. A more complete discussion can be found in [24].

Figure 2.7 shows a 3D model of the superconducting coil system. It consists of three coaxial coils that are placed side by side on a common support tube with 300 mm inner

2. Design and systematical aspects of the experiment




Figure 2.7.: 3D model of the superconducting coil system. In vertical orientation the left-end side would be at the bottom. The three light-green bars hold the superconducting switches (brown) of the three coils.

diameter. Going through the picture from left to right we see:

- The *bottom* or *shutter coil* sits on the system’s bottom end and serves as a magnetic UCN shutter. According to the manufacturer *Oswald Elektromotoren* it provides a 1.73 T field at $I_{bot} = 300$ A and can be ramped from 0 to 300 A within 10 seconds using a 15 V power supply [31]. However during our 2014 beam time we measured a ramping voltage of $U = 5.5$ V at a ramping speed of $\dot{I} = 5.7$ A s⁻¹, yielding an inductance of $L = \frac{U}{\dot{I}} = 965$ mH. Those numbers are in good agreement with the theoretical predictions for a coil with corresponding dimensions and current density. In order to reach a ramping speed of 30 A s⁻¹, a 30 V supply voltage is thus necessary.
- The *bias coil* allows to generate a uniform bias field throughout the entire trapping region. Like the bottom coil it can be ramped up to 300 A where it would produce a 1.31 T field.
- The *top coil* is by far the strongest one and reaches 5.04 T at 300 A. However Kent’s FEMM simulations showed that such high fields would irreversibly demagnetise the permanent magnets and proposed a maximum field strength < 4 T [24].
- The (defective) *superconducting switches* are used to “switch” between *open* and *persistent* mode. When the switches are heated above their critical temperature the superconducting circuits are interrupted and the coils can be energised by an external current source. Once the heating is turned off the switches return to their superconducting phase: the coils’ circuits are closed and the current is maintained until the next ramping cycle. Unfortunately the *self-heating* of the switches prevented a proper operation during the 2014 beamtime. Above ~ 200 A

2. Design and systematical aspects of the experiment

the switches stayed permanently normal conducting and we could not put the coils into persistent mode. Consequently we could only measure with a shutter-coil field of 200 A during the beamtime, resulting in a considerably decreased UCN filling rate (see section 3.3.2). The switches will be exchanged prior to future data taking.

Depending on the experimental procedure the bias and the top coil serve for different purpose:

In a pure *fill-and-empty* measurement the top coil becomes obsolete as the trapping region's extent in vertical direction is limited by gravity and thus no additional magnetic barrier is needed at the trap's top end⁴. The bias coil can be used for systematic investigation of spin-flip induced UCN losses. Theoretical considerations predict negligible spin-flip rates for minimal B -field strengths in the order of 10 mT (see section 2.2.1 and [24]). The left plot in figure 2.8 indicates that the bottom coil should satisfy this requirement. It shows the z -component of the magnetic field on the system's middle axis as a function of the vertical position for $I_{bot} = 300$ A and different bias-field strengths. At $z = 1000$ mm - approximately the highest point a trapped UCN can reach - $B_z(I_{bias,0A})$ still is in the range of 10 mT and the spin-flip probability should thus be largely suppressed throughout the entire trapping region. However due to the inductive coupling of the coils it is not possible to simply leave the bias coil untouched. The left plot in figure 2.8 shows the B -field generated by the bias coil as we measured it during the 2014 beam time on PF2. The baseline gives the field strength for a bias coil current of 30 A and no bottom coil field ($B_{0A,30A} = 101.5$ mT). With the bottom coil at 50 A and the bias coil at 30 A the overall field gets $B_{50A,30A} = 111.5$ mT. The negative spike shows the effect of the inductive coupling as we ramped the bottom coil to 50 A with the bias coil being in persistent mode ($B_{50A,?A} = 95.8$ mT). From those figures we can now calculate the inductive coupling $C_{bot-bias}$ between the bottom and the bias coil and get:

$$C_{bot-bias} = \frac{B_{50A,30A} - B_{50A,?A}}{B_{0A,30A}} \frac{I_{bias}}{I_{bot}} = \frac{111.5 \text{ mT} - 95.8 \text{ mT}}{101.5 \text{ mT}} \frac{30 \text{ A}}{50 \text{ A}} = 93 \text{ m I}_{bias} [\text{A}] \text{ I}_{bot} [\text{A}]^{-1}, \quad (2.3)$$

meaning that the bias coil is ramped down by 93 mA for every Amp the bottom coil is ramped up. Ramping the bottom coil to 300 A this adds up to 27.9 A respectively a B -field reduction of $\Delta B_{bias} = 122$ mT. Referring again to the calculated B -field values from the left plot, we see that if we started from an uncharged bias coil we would end up with a 0-field crossing right in the middle of the storage region. In order to avoid that we will either have to put the bias coil into open-loop mode every time we ramp the bottom coil or *preramp* it to a "compensation" current. A rigorous approach would probably call for a similar treatment of $C_{bias-top}$ in order to make sure that there are no overlooked back-door effects.

In a *counting-the-dead* type measurement the above discussion of the spin-holding field gets secondary as the default bias field will exceed the necessary holding field by about two orders of magnitude. The bias field would then serve as a guiding field and

⁴Of course this only holds for a vertically oriented trap. In a horizontal set-up there is no gravitational barrier and we would have needed the top coil to close the trap.

2. Design and systematical aspects of the experiment




Figure 2.8.: The plot on the left side shows B_z on the cryostat’s middle axis for different bias fields and the bottom coil running at 300 A. The magnetic field values have been calculated using the *Mathematica* plug-in *Radia* (see next section). The plot on the right shows the z -component of the B -field as measured with a *LakeShore* Hall probe. For the measurements the probe was mounted on a level with the bias coil ($z \approx 500$ mm) and roughly aligned with the cryostat’s mid-axis. The B -field rise occurs when the bottom coil is ramped and the bias coil is in *open-loop mode*, the decrease when the bottom coil is ramped with the bias coil in *persistent mode* (see text).

help transport the decay products to a particle detector placed above the UCN storage region. If we chose to detect the emitted protons, the top coil can be used to focus them onto a small sized detector (e.g. the silicon drift detectors that have also been used in the aSPECT experiment, see [40]). In contrast it is not possible to use it as an electron lens as the acceleration voltage needed to lift the electrons over the magnetic barrier scales with their kinetic energy which is several orders of magnitude higher than the protons’ energy. The top coil would then be used to prolongate the bias coil’s guiding field. Some more information on charged particle detection within HOPE, including a detailed discussion of the necessary bias fields will be given in section 2.2.3.

2.1.3. UCN potential

Taking equation 1.16 and 1.17 from section 1.2 the UCN potential $E_{pot}(\mathbf{r})$ throughout our trapping region gets:

$$E_{pot}(\mathbf{r}) = \mathbf{B}(\mathbf{r}) \cdot 60.3 \text{ neV T}^{-1} + \mathbf{e}_z \times \mathbf{r} \cdot 102 \text{ neV m}^{-1}, \quad (2.4)$$

where $\mathbf{B}(\mathbf{r})$ is the magnetic field at \mathbf{r} and \mathbf{e}_z is the unit vector in positive z -direction. $\mathbf{B}(\mathbf{r})$ can only be approximated by analytical formulas. Therefore I chose a numerical approach for my calculations: I used the *Mathematica* plug-in *Radia*, an in-house development of our neighboring institute, the *European Synchrotron Radiation Facility* (ESRF) that is “dedicated to 3D magnetostatics computation. It is optimized for the design of undulators and wigglers made with permanent magnets, coils and linear/nonlinear soft magnetic materials [9].” *Radia* has everything we need as it allows for 3D modeling, permanent-

2. Design and systematical aspects of the experiment

and electromagnetic field sources and even considers the magnetic susceptibility of the materials used. It is based on the *boundary element method* and solves a given problem by subdividing the field sources into small sized *objects* and iteratively calculating the final magnetisation state of everyone these objects.

In figure 2.9 we can see a B -field map of the storage region in the ρ - z -plane calculated with *Radia*. The high-field region on the bottom is dominated by the shutter coil while the upper regions show the ρ^3 -dependency of the permanent magnet system. Together the two magnet systems form a magnetic well for neutrons. By including gravity we get the 3-D trapping potential shown in the other three plots. Besides the optimised fill-and-empty setting - where *optimised* only means ramping up the bottom coil to its maximum - figure 2.9 also presents potential maps for the counting-the-dead mode - which adds a guiding field to the former one - and the coil setting as it was during the 2014 beam time.

However, in order to quantitatively compare different settings the sole knowledge of the cut-off energy or the spatial dimension of the UCN potential is not sufficient. We also have to take into account the UCN spectrum. This is problematic as the spectrum of the stored UCN is not identical to the spectrum of the source that the trap is connected to. There is no analytic way to determine the final spectrum and an ideal treatment would involve a particle-tracking simulation of the filling process in order to get a feeling for the final UCN spectrum. In the present case we will content ourselves with an approximate solution that should be sufficiently precise to draw some comparative conclusions.

The idea is to consider the filling as an equilibration process that results in a uniform *phase-space density* (i.e. a state where a particular UCN is equally likely to be found in any of it's accessible microstates). A detailed discussion of the topic can be found e.g. in [16]. Once the equilibrium is reached the spatial probability density p depends on the kinetic energy E_{kin} as follows:

$$p(E_{kin}) \propto \sqrt{E_{kin}}. \quad (2.5)$$

Using the local dependency of the kinetic energy

$$E_{kin}(\mathbf{r}) = E_{tot} - E_{pot}(\mathbf{r}), \quad (2.6)$$

where E_{tot} is the UCN's *total energy* (i.e. its kinetic energy at trap's potential minimum), we can express p in dependency of \mathbf{r} :

$$p(\mathbf{r}, E_{tot}) \stackrel{2.6}{\propto} \sqrt{E_{tot} - E_{pot}(\mathbf{r})}. \quad (2.7)$$

Equation 2.7 shows that the UCN density in the trap does not only depend on \mathbf{r} but also on E_{tot} . We end up with the following double integral for the total UCN number N_{tot} :

$$N_{tot} = \int_V \left[\int_{E(\mathbf{r})} n(E_{tot}) \sqrt{E_{tot} - E_{pot}(\mathbf{r})} dE_{tot} \right] d\mathbf{r}, \quad (2.8)$$

where the integration boundaries $E(\mathbf{r})$ of the inner integral are defined by $E_{pot,max} \geq E(\mathbf{r}) \geq E_{pot}(\mathbf{r})$, V is a connected volume that satisfies the relation $E_{pot,max} - E_{pot}(V) \geq 0$ and $n(E)p(\mathbf{r}, E)dE$ is the UCN density a point \mathbf{r} in the energy range $(E, E + dE)$.

2. Design and systematical aspects of the experiment




Figure 2.9.: The plot on the far left shows a map of the B -field in the trapping region with the coils set to standard fill-and-empty mode while the neighboring plot shows the associated UCN trapping potential as given by equation 2.4. The two plots on the right show trapping potentials for alternative coil settings. The contour lines indicate the total B -field strength in T respectively the UCN potential in neV, the thick black lines indicate the cut-off energy of the respective setting. All three UCN potential plots have been assigned individual offsets in order to satisfy $E_{pot,min} = 0$ neV.

2. Design and systematical aspects of the experiment

Constant $B_{bot} = 1.73$ T						
$B_{bias/top}$ [T]	0	0.1	0.2	0.3	0.4	0.5
$dN \propto E^{1/2}dE$	1	0.86	0.71	0.58	0.46	0.35
$dN \propto EdE$	1	0.83	0.66	0.51	0.39	0.28
Constant $B_{bias/top} = 0$ T						
B_{bot} [T]	1.73	1.61	1.50	1.38	1.27	1.15
$dN \propto E^{1/2}dE$	1	0.72	0.50	0.34	0.21	0.13
$dN \propto EdE$	1	0.68	0.45	0.28	0.17	0.09
Horizontal HOPE						
$B_{bot} = B_{top} = 1.73$ T, $B_{bias} = 0$ T						
$dN \propto E^{1/2}dE$	2.00					
$dN \propto EdE$	1.90					

Table 2.1.: Filling rates calculated using equation 2.8 for two different UCN spectra and various different coil settings. As 2.8 is no actual equation but only a relation the given numbers have only a comparative interest, no statement is made on the total filling rates.

We have no precise knowledge of $n(E)$ and can only make an educated guess on the final spectrum. In section B we will see that the theoretical spectrum of a superfluid helium source such as SUN2 is given by the relation:

$$dN/dE \propto \sqrt{E}. \quad (2.9)$$

However this distribution will get further modulated through energy-dependent UCN losses in the source itself and during the transport to the experiment. In his thesis Kent mentions a second energy-distribution function stating that it might more appropriately describe the final UCN spectrum:

$$dN/dE \propto E. \quad (2.10)$$

For my evaluation of equation 2.8 I have used both proposed spectra. The actual integration was done numerically using the potential maps introduced in figure 2.9 and a step size of 1 mm. Table 2.1 summarises the results for some relevant settings.

If we want to compare the results for the two different spectra it is important to first clarify *how* they should be compared. Let's consider the values for $B_{bot} = 1.15$ T. In the case of $dN \propto E^{1/2}dE$ we end up with 13 % of the neutrons we would have at $B_{bot} = 1.73$ T while it would be only 9 % for the alternate spectrum. On first glance one could interpret this as "At $B_{bot} = 1.15$ T I will have 50 % more neutrons with spectrum 2.9 than with spectrum 2.10." However this interpretation contains a logic error as it assumes that $N_{tot}(B_{bot} = 1.73$ T) is equal for both spectra. In reality we do not know the relative offset of the two spectra and the only conclusion we can draw is "By decreasing B_{bot} to 1.15 T I loss 4 % more UCN with spectrum 2.10 than with spectrum 2.9." If

2. Design and systematical aspects of the experiment

we now look at the table we find that the results for the two spectra show only slight differences $\leq 7\%$. This suggests that 2.8 does not depend too strongly on $n(E_{tot})$ and that our further conclusions should still have some validity even if neither 2.9 nor 2.10 give an accurate description of our final UCN spectrum.

Obviously N_{tot} strongly depends on B_{bias} : at $B_{bias} = 0.4$ T we have already reduced N_{tot} by more than 50 %. This relation will become important for our future counting-the-dead measurements as the electron/proton detection efficiency rises with increasing guiding field and we will have to find a balance between the two counteracting effects. Some preliminary simulations on the matter will be presented in section 2.2.3.

The dependence on B_{bot} is even more important than the one on B_{bias} . This is particularly interesting regarding our 2014 beamtime as it allows us to extrapolate to the total UCN numbers we could achieve in future beamtimes, once the superconducting switches have been repaired and we can use the bottom coil at its nominal value. The calculations suggest that this improvement could yield a factor ~ 8 or even more. In section 3.3.1 I will compare those predictions with some experimental data from the beamtime.

Finally there are the numbers for HOPE in horizontal position. Compared to the vertical version its total volume is considerably higher as it is much less influenced by gravity. However, aside the statistical gain, the horizontal trap design is very disadvantageous. First there are technical issues such as the support of the PMS inside the cryostat's bore but also experimental drawbacks like the extended emptying times (the UCN can not just "fall out" of the trap) or the exclusion of counting-the-dead measurements with electrons due to the magnetic mirror effect. We therefore opted for the vertical solution in spite of the lower neutron numbers. A detailed discussion of the pros and cons of both concepts can be found in [25].

2.2. Internal parts and experimental methods

In order to provide a reliable and precise measure of the neutron lifetime it will be crucial that we understand the systematics of our experiment. I will start this section with an overview of the most important systematic problems we will have to deal with. Afterward I will present the tools we developed to examine and cope with those effects.

2.2.1. Overview on systematic effects

UCN depolarisation

The possibility of losing UCN by depolarisation is unique to magnetic traps. Nevertheless the underlying physics is well known and it is possible to give relatively good theoretical predictions on the spin-flip rate R_{flip} . A handwavy criterion for the suppression of spin-flips is the compliance with the *adiabatic condition*

$$\omega_L = \frac{2\vec{\mu}_n \mathbf{B}}{\hbar} \gg \frac{1}{|\mathbf{B}|} \left| \frac{d\mathbf{B}}{dt} \right|, \quad (2.11)$$

where ω_L is the Larmor frequency (the precession frequency of the UCN's spin around the magnetic field \mathbf{B}), $\vec{\mu}_n$ is the neutron's magnetic moment and \hbar the Planck constant.

2. Design and systematical aspects of the experiment

As long as 2.11 is fulfilled the neutron spin stays constantly aligned with the magnetic field and R_{flip} is negligible. Obviously the condition is maximally violated for $\mathbf{B} \rightarrow 0$. In most applications it is sufficient to avoid zero-field crossings in order to satisfy 2.11 and thus get acceptably low spin-flip rates.

In Kent's thesis we can find a more solid discussion of the spin-flip probability in HOPE [24]. He measured the storage time constant of the PMS in dependency of the bias field and found that it saturated at about 1.5 mT, suggesting that spin-flips were efficiently suppressed from thereon. However his set-up differed greatly from the final experiment. Back then the SCS was not yet delivered and he had to close the trap with a Teflon plug on the bottom. The bias field was generated by a normal conducting copper coil. Also the statistical accuracy of his measurements was far from the level we are ultimately aiming for. Nevertheless his measurements give a good first estimate and predict acceptable spin-flip rates for minimum field strengths in the order of a few mT.

He also performed some Monte-Carlo based spin-tracking simulations in order to give a theoretical prediction on R_{flip} . The UCN he simulated were not free to propagate through the trap but were only going back and forth on a straight line passing at maximum speed right through the system's middle axis. From a spin-holding point of view this trajectory is a worst case scenario and his results thus provide upper limits. In order to extract R_{flip} he calculated the quantum mechanical transition probability for every single step and averaged over the entire integration time

$$R_{flip} = \frac{\sum_n |\langle \chi_n | \chi_- \rangle|^2}{N}. \quad (2.12)$$

Here $\langle \chi_n |$ is the UCN's spin state at iteration step n , $\langle \chi_- |$ the high-field seeking spin-state and N the total number of steps. His simulations included only an idealised octupole and a uniform bias field. For 10 mT he obtained $R_{flip} = 1 \times 10^{-7} \text{ s}^{-1}$. If we assume a uniform spin-flip probability for all trapped UCN, a neutron lifetime of 880 s and do not consider any other loss channels we then end up with total storage times of 879.9 s, i.e. an upper limit of 0.1 s for the systematic shift. In section 2.1.2 we have seen that we will easily reach field strengths in this order. Even without an additional bias field the total field strength does not fall beyond ~ 8 mT throughout the entire trapping region.

In the immediate vicinity of the permanent magnets possible magnetic defects could lead to much higher field gradients than predicted by the ideal octupole field. However an estimation shows that a considerable contribution to R_{flip} is improbable. By using $\left| \frac{d\mathbf{B}}{dt} \right| = \left| \frac{d\mathbf{B}}{dr} \frac{dr}{dt} \right|$, where $\frac{dr}{dt} = v$ is the velocity of the neutron we can reformulate 2.11:

$$\omega_L |\mathbf{B}| v \gg \left| \frac{d\mathbf{B}}{dr} \right|. \quad (2.13)$$

Close to the magnets the B-field strength is high ($|\mathbf{B}| \approx 1 \text{ T}$) and the neutron velocity low ($v = \sqrt{2(E_{tot} - V_{mag})/m_n} \lesssim 1 \text{ m/s}$). Using $\omega_{L,neutron}/2\pi = 29.17 \text{ MHz T}^{-1}$ we get $30 \times 10^6 \text{ T/m} \gg \left| \frac{d\mathbf{B}}{dr} \right|$. I.e. in order to violate 2.11 the B-field gradients close to the wall must be in the order of $\sim 1 \text{ T } \mu\text{m}$. If there were gradients in that order they would most probably get evened out by the autodemagnetisation of the permanent magnets.

2. Design and systematical aspects of the experiment




Figure 2.10.: The low- and the high-field seeker potential on the middle axis of the magnet system ($B_{bot} = 1.73$ T, $B_{bias} = 0$ T). Both potentials are shown with (the same) arbitrary offset. The light-blue and light-red areas indicate the effective storage regions for both cases.

An empirical determination of R_{flip} can be done using the same tools that will be discussed in broad detail in the context of overcritical UCN contamination: either via the UCN leakage during the storage period or by measuring the trap’s storage-time constant several times with different coil settings which implies different spin-flip probabilities.

Trapped high-field seekers

HOPE works without preliminary spin selection meaning that both spin states are free to enter the trapping region during the filling phase. Once the bottom coil has been ramped up, the shutter field acts as a potential barrier for the low-field seeking UCN and prevents them from leaving the trap on the bottom. For the high-field seekers the same field looks like a potential well that draws them out of the trapping region. Due to gravity the well is tilted towards its lower end and most of the high-field seeking neutrons will finish by falling out of the well. However during the ramping process they lose some of their potential energy - the inverse process of the spectrum heating of the low-field seekers - and it is possible that a low energetic UCN gets “sucked” into the potential well during its formation and eventually gets trapped in the high-field region. In contrast to the low-field seekers the *trapped high-field seekers* are not stored in a purely magneto-gravitational potential but in a *magneto-material* trap and are thus exposed to wall collisions and the associated loss channel. Nevertheless storage times in the order of 10^2 seconds could be possible. Even though the effect might seem exotic it has already been observed at the ILL some years ago [7] and we will have to address this issue. In the following section I will argue that in our current set-up, using the Teflon piston, there should be no trapped high-field seekers left after the cleaning process.

2. Design and systematical aspects of the experiment

Residual gas collisions

As the UCN energy regime is negligible compared to the energy regime of the residual gas particles, any interaction between an UCN and a gas particle will inevitably raise the UCN's energy several orders of magnitude above the trap's cut-off energy. The loss rate due to gas collisions thus is equal to the UCN's scattering rate τ_{scat}^{-1} :

$$\tau_{scat}^{-1} \propto \text{number of scatterers} \stackrel{T=const}{\propto} p, \quad (2.14)$$

where p is the pressure in the trapping region. If we omit all other possible loss channels we got the following formula for the total storage-time constant $\tau_{stor}(p)$:

$$\tau_{stor}^{-1}(p) = \tau_n^{-1} + \tau_{scat}^{-1}(p) \xrightarrow{\tau_{scat}(p)=c/p} \tau_{stor}(p) = \frac{c\tau_n}{c + p\tau_n}, \quad (2.15)$$

where τ_n is the neutron lifetime and c is the proportionality constant between τ_{scat}^{-1} and p . Unfortunately we can not precisely determine the constant c without knowing the composition of the residual gas. In order to get an estimate for $\tau_{stor}(p)$ I fixed c indirectly by referring to pressure dependent storage times found in the *Gravitrap* experiment [38], which they determined by variation of the vacuum conditions. For a final pressure of 5×10^{-6} mbar they give a deviation $\Delta\tau_{scat}(p) = \tau_n - \tau_{stor}(p) = 0.4$ s. From that we can calculate c :

$$c = \frac{p\tau_n\tau_{stor}(p)}{\tau_{shift}(p)} = 9.7 \text{ mbar s}. \quad (2.16)$$

Figure 2.11 shows $\Delta\tau_{scat}(p)$ and its associated error $\delta\Delta\tau_{scat}$ using c from equation 2.16. $\delta\Delta\tau_{shift} = \frac{d\Delta\tau_{scat}(p)}{dp}\delta p$ is calculated assuming that the uncertainty of the pressure measurement δp is dominated by the inherent inaccuracy of the pressure gauges. During the 2014 beamtime we monitored the gas pressure in the main vacuum chamber using two *PKR251 FullRange Gauges* from *Pfeiffer*. In the region of interest they have an accuracy of $\sim \pm 30\%$ and a reproducibility of $\sim \pm 5\%$ [33]. The plot suggests that at pressures $< 10^{-6}$ mbar the storage time shift should be negligible. However in order to confirm that, we will have to find a more reliable way to determine c . A first step could be to measure the residual gas composition with a mass spectrometer and calculate c from the neutron cross-sections of the found particles.

On the experimental side probably the most difficult part will be to precisely determine the pressure in the storage region. During the past beamtime we had a constant pressure gradient of about an order of magnitude between the two gauges that were connected to the main vacuum chamber. One of them was separated from the storage region by narrow passages in the vacuum system and the other one was sitting right next to a turbo pump. Thus neither of them was suited to give an exact measure of the pressure in the region of interest.

Overcritical UCN

Overcritical UCN - neutrons that have a total energy above the trap's threshold - are not an exclusive feature of magnetic UCN traps. However the special conditions in magnetic

2. Design and systematical aspects of the experiment




Figure 2.11.: Shift of the mean UCN storage time due to residual gas collisions in dependency of the pressure in the trap. The red line gives the systematic error of the correction. E.g. at a final gas pressure of 10^{-5} mbar we get a correction of ~ 0.8 s and an error $\delta\Delta\tau_{\text{scat}} = \pm 0.4$ s.

2. Design and systematical aspects of the experiment

traps make that the mean storage time $\bar{\tau}_{stor}$ of overcritical neutrons is substantially longer than it would be in a material trap (see e.g. [8]). Especially neutrons with energies that are only slightly higher than the cut-off can reach storage times of several hundreds of seconds. There is no way to distinguish such a long-lived overcritical neutron from an actually storable one once it has found its way into the final data and it will thus shift the measured storage time to lower values.

There are several reasons for this disadvantageous behavior of the overcritical UCN portion:

- The *effective cleaning*⁵ *area* is much smaller than in a material trap. Let's consider a UCN that has only a very small energy excess compared to the cut-off. In the case of a box-shaped material trap its effective cleaning area is essentially given by the entire bottom plate of the trap as it is only there that the UCN reaches its maximum kinetic energy and can exceed the wall's Fermi potential. Referring to the UCN potentials shown in figure 2.9 we see that in our case the cleaning area is limited to only a small spot in the spacial center of the bottom coil.
- In a material trap the effective cleaning area coincides with the UCN potential minimum. As the only force acting on the UCN is gravity, they are constantly accelerated to the bottom plate and their collision frequency will be accordingly high. Figure 2.9 shows that in HOPE the potential minimum is situated somewhere in the trap's center and hence not congruent with the "exit point". Moreover the high symmetry of magneto-gravitational potentials favours closed (and hence stable) trajectories of overcritical neutrons. Once stuck on such a stable trajectory even "high-energetic" UCN could become *marginally trapped* and stay within the trap for many hundreds of seconds. In his thesis Kent demonstrated that, theoretically, such trajectories are possible up to ~ 160 neV!
- Overcritical UCN with energies that are not "very close" to the cut-off energy will also be able to collide with the side walls. However their loss probability per bounce will be very small. In contrast to the slightly overcritical UCN in a material trap they will not hit the wall with a kinetic energy that actually exceeds the wall material's Fermi potential but only with a small kinetic energy given by $E_{kin}(\mathbf{r}) = E_{tot} - V(\mathbf{r})$. Consequently their loss probability will be given by contamination and wall defects and probably again be in the order of⁶ 10^2 s. We could have enhanced that probability by choosing an absorbing material for our bore tube. However we preferred the stainless steel tube as we want to use the *UCN leakage rate* as tool to investigate our systematics. By adding an additional loss channel the interpretation of the leakage rate would have become much more imprecise.

⁵In this context I mean the passive *auto-cleaning* of the trap, in contrast to the actively accelerated spectrum cleaning we perform at the beginning of each measurement cycle.

⁶This is the *magneto-material hybrid trapping* effect from [8] that has already been presented in section 1.2

2. Design and systematical aspects of the experiment

- The well-shaped form of the magneto-gravitational potential also hinders the spectrum cleaning. In a material-trap the cleaning can be easily achieved using a flat absorber entered from the top. The cut-off energy is defined by the absorber's height and the overcritical UCN are directed towards the absorber by diffuse reflections on the trap's side walls. Our equipotential surfaces are not flat but of oval shape. Also there is no counterpart to the mode mixing effect of the diffuse wall reflections and we must cut deep into the trapping region if we want to be sure to interfere with all possible overcritical trajectories. On the other hand cutting deeper into the storage region means losing a larger amount of storable neutrons.

Accounting for the overcritical neutron contamination, the mean storage time $\bar{\tau}_{stor}$ can be given as

$$\bar{\tau}_{stor} = \frac{N_{storable}\tau_n + \int_{E_{max}}^{\infty} n(E)\tau_{crit}(E)dE}{N_{storable} + \int_{E_{max}}^{\infty} n(E)dE}, \quad (2.17)$$

where $N_{storable}$ is the total number of storable UCN, τ_n is the neutron lifetime (other loss channels are neglected), E_{max} is the trap's cut-off energy, $n(E)dE$ is the energy-dependent amount of neutrons with energy in the range $(E, E + dE)$ and $\tau_{crit}(E)$ is the storage-time constant of an overcritical UCN with energy E . As we know neither $n(E)$ nor $\tau_{crit}(E)$ I will use a (very) simplified model throughout this thesis and describe the entire overcritical UCN share by only two parameters, their total number N_{crit} and their mean storage-time constant τ_{crit} . Equation 2.17 then becomes:

$$\bar{\tau}_{stor} = \frac{N_{storable}\tau_n + N_{crit}\tau_{crit}}{N_{storable} + N_{crit}}. \quad (2.18)$$

Figure 2.12 shows the dependency of $\bar{\tau}_{stor}$ on τ_{crit} and the overcritical neutron share $n_{crit} = \frac{N_{crit}}{N_{storable} + N_{crit}}$ as given by equation 2.18. If we require the systematic shift $\Delta\tau_{sys,crit} = \tau_n - \bar{\tau}_{stor}$ to be < 1 s and exclude scenarios with extremely long-lived UCN ($\tau_{crit} \gtrsim 700$ s) we see that we will have to decrease n_{crit} to a few times 10^{-3} .

In summary, it will be a complex task to get $\Delta\tau_{sys,crit}$ under control. We will have to find a way to clean the initial spectrum without severely reducing the number of storable neutrons and figure out a way to determine n_{crit} with the required sensitivity. Furthermore it would be desirable to monitor the overcritical UCN spectrum in order to refine equation 2.18.

Spectrum heating during storage

Another process to be avoided is a gain of small amounts of energy of UCN during storage, whereby they may become overcritical. The permanent magnets' temperature drifts presented in section 2.1.1 are a possible source of spectrum heating. Those fluctuations create a synchronous fluctuation of the magnets' remanence field and thus of the UCN potential. Like during the ramping of the bottom coil the stored UCN will be affected

2. Design and systematical aspects of the experiment




Figure 2.12.: Dependence of the UCN ensemble's mean storage lifetime $\bar{\tau}_{stor}$ on the overcritical storage-time constant τ_{crit} and the overcritical neutron share n_{crit} . $\bar{\tau}_{stor}$ is indicated by the colour coding, given on the right, and by a few selected dashed lines. The upper plot shows the extreme case of very small n_{crit} -values while τ_{crit} is varying widely. The lower plot shows the inverse scenario with only very large τ - values and varying n_{crit} -values.

2. Design and systematical aspects of the experiment

by those changes and their total energy E_{tot} will evolve as

$$E_{tot}(t) = E_{tot}(t_0) + \int_{t_0}^t \frac{dV(\mathbf{r}(t'), t')}{dt'} dt', \quad (2.19)$$

where $E_{tot}(t_0)$ is the total energy of a particular UCN at the starting time t_0 of the storage period, $\mathbf{r}(t)$ is its position at time t and $V(\mathbf{r}(t), t)$ is the UCN potential at point \mathbf{r} and time t . As equation 2.19 depends on the trajectory $\mathbf{r}(t)$ of the every single neutron, a precise evaluation of 2.19 is complicated. However the maximal energy gain can be easily given as $\Delta E_{tot,max} = n \times V_{peak-to-peak}$, where n is the maximum number of temperature oscillations during one storage period and $V_{peak-to-peak}$ is the peak-to-peak amplitude of the corresponding potential oscillation. With $V_{peak-to-peak} \propto 0.01$ neV and $n = \frac{\text{maximum storage time}}{\text{oscillation period}} \approx \frac{3000}{300} = 10$ we get $\Delta E_{tot,max} \lesssim 0.1$ neV.

Another energy modulating process is Doppler heating due to mechanical vibrations of the magnet system. Our cryostat experiences regular tangible vibrations coming from the cold head that is used to cool the coils' current leads (see section A). In order to get an idea of the possible energy gain one could measure the amplitude and frequency of the vibrations and make a similar estimate as for the temperature drifts. However the most reliable approach will be to directly monitor the spectrum of the stored UCN and its time evolution. In section 3.2.3 I present a possible way to investigate the UCN spectrum that gets along without additional installations.

Filling rate variations

The filling of the trap is a random process and consequently the initial number of neutrons will change from cycle to cycle. As long as those fluctuations are purely statistical they will not influence our measurements. However if there are some systematic drifts we have to include an individual weighting factor to every single measurement. Such systematic drifts can be long-term effects e.g. due to the steady burn-up of the ILL's fuel element. Or they can be arbitrary variations between any two measurement cycles. On a LHe UCN source like SUN2 such fluctuations could e.g. be triggered by temperature changes of the converter material.

Probably the most elegant solution to this problem is to monitor the neutrons' decay products instead of just counting the number of surviving UCN. By measuring the decay products, every single cycle individually records the entire decay curve. This way we would no longer have to rely on the relative comparability of the particular measurements. Another possibility would be to constantly monitor the incoming UCN flux in order to determine individual weighting factors.

Determination of the storage time

The emptying process of the trap is not a point-like event but has a finite time span that is given by the down-ramping speed of the bottom coil. The slow decrease of the shutter field does not only blur the end point of the storage phase but also adds an energy

2. Design and systematical aspects of the experiment

dependence to the total duration of the storage time. Thus, a diligent treatment of the time variance requires knowledge on the final UCN spectrum of every single measurement point in order to assign them their individual time uncertainties δt . In reverse we can also use the emptying peak in order to obtain information on the UCN spectrum (see section 3.2.3). Eventually this information could even be used to determine δt itself.

2.2.2. Internal installations and UCN manipulation

The design of HOPE's main vacuum chamber and its mechanical parts was mainly influenced by five - partly congruent, partly conflicting - criteria:

- Providing all necessary tools to manipulate the neutron ensemble with a main focus on the spectrum cleaning phase.
- Allowing for a continuous monitoring of the UCN leakage rate during the storage phase. We deem this a major asset for investigation of systematic effects.
- Optimising the “statistical characteristics” of the trap. This involves maximising the trap's effective volume and minimising the UCN losses during the filling and emptying phase.
- Limiting the constraints on a future charged-particle detector.
- Providing the necessary vacuum conditions.

In the following I will go through the main components of HOPE's *experimental region* - the main vacuum chamber, the detector switch, the absorber and the Teflon piston - present their most interesting features and explain how they address above requirements.

Main vacuum chamber

Figure 2.14 shows the system's top flange. It contains a vacuum feedthrough for the absorber actuation and a *DN 63 ISO-F* port that serves as the main pumping port. The lower part of the system contains several additional pumping bypasses. They have much smaller cross-sections (*DN 25 ISO-KF* or smaller) and for the most part are separated from the trapping region by narrow slits as they are connected to potential dead volumes (e.g. one can see a pumping bypass just beneath the Teflon plug on detail 1 of figure 2.13). Unfortunately it is not possible to further increase the pumping power by adding ports to the lower parts of the vacuum chamber since they would lead to important UCN transport losses.

Below the top flange comes a 700 mm long pipe with 214 mm inner diameter. In the future it will house our charged-particle detector and its electronics. In order to keep maximum flexibility, we maximised the available space. The pipe's lower end is connected to the *bore tube* via a short bellows (see figure 2.6 from section 2.1.1).

Detail 3 of figure 2.13 shows the bore tube's bottom part. It is directly flanged to the T-piece (grey) that connects the cryostat, the horizontally incoming UCN guide and the detector switch. We deemed it preferable to minimise the vertical distance, and thereby

2. Design and systematical aspects of the experiment

the gravitational potential difference, between the cryostat and the UCN guide. The figure shows that the final clearance between the bottom flange (faint-red) and the T-piece is just sufficient to still be able to tighten all the surrounding screws. To minimise UCN losses all shown parts have been custom made or specially adapted in order to guarantee high-precision junctions. Generally the critical tolerances are < 0.1 mm. All surfaces that are exposed to UCN are either made from stainless steel or coated with NiMo.

Detector switch

The detector switch allows to change between the *counting position* (direct connection between storage region and the UCN detector) and the *piston position* (storage region is accessible to the Teflon piston). It is a $517 \times 152 \times 190$ mm aluminum case (orange in figure 2.13) that contains a $365 \times 110 \times 160$ mm Teflon slide (bright-yellow) and is actuated via a pneumatic piston (detail 1).

The Teflon slide has two lead throughs, a straight one for the Teflon piston and one that is inclined by 52.5° and leads to the UCN detector. The plug lead through is a simple hole with 91 mm diameter. The inclined hole for the UCN contains a stainless steel insert in order to enhance its guiding properties. On its lower end the slope continues in an extension piece that ends in a vertical cone. The cone serves as adapter piece to the standard 80 mm OD stainless steel guides that are used on PF2. The total horizontal off-set from the trapping region's axis is of ~ 250 mm.

The switch is actuated by a custom made pneumatic piston that fits our high vacuum requirements. Most commercially available pneumatic drives use a dynamic O-ring system as vacuum feedthrough. In order to avoid the dynamic seals we use a *COMVAT* bellows as motion feedthrough. It is equipped with *CF*-type flanges on both sides since simple *Viton* O-rings are not rated for pressure differences < 1 bar. Its total stroke is limited to 110 mm by the stop positions in the switcher case. At 0.5 to 1 bar overpressure it moves gently taking some 2 to 3 seconds from one end position to another.

Teflon piston

The Teflon piston is our remedy against *stable overcritical UCN trajectories*. During the spectrum cleaning phase we will lift it up into the trapping region where it will serve as a “mode mixer”: when a UCN hits the piston it will be arbitrarily deflected and eventually be directed straight onto the absorber. In section 2.1.3 we have seen that our UCN potential is a 3-D well with its minimum E_{min} sitting somewhere between $z = 200$ and 300 mm, depending on the coils' setting. All UCN trajectories, closed orbits as well as arbitrary flight paths, must circle around this minimum. Therefore we can exclude stable UCN trajectories during the cleaning phase by lifting the piston above E_{min} . Furthermore it resolves the problem of the trapped high field seekers as it will wipe out all of during its two passages through the high B -field region.

Using a vertical “paddle” instead of the piston would reduce the Doppler heating during the lifting process. We decided to test the piston first due to a systematic downside

2. Design and systematical aspects of the experiment

of the paddle variation: Since the paddle's movement is parallel to its surface diffuse UCN reflections during the motion can result in both, Doppler heating *and* cooling, independently of the paddle's direction of motion. In contrast the Teflon plug will heat the spectrum *only* on its way into the trap and cool it on its way out. Some more details on the effect can be found in [24].

We chose to make the piston out of Teflon as it combines several advantageous characteristics: it has a relatively high Fermi potential (~ 100 neV), low friction, is lightweight and non-magnetic. Nevertheless we are planning to test alternatives. E.g. its surface could be covered with Fomblin grease or a thin metal sheet could be added in order to further increase the UCN potential.

The plug and its actuation system can be seen in figure 2.13. Detail 2 shows the actual drive, a stepper motor from *Nanotec* (motor type *ST8918S*) that is coupled to a M20 \times 4 trapezoidal thread rod. The rotary movement of the thread rod is translated into the vertical displacement of slide *A* via the screw *B*. The slide is guided by two high-precision shafts (*C*). The motion feedthrough is done using a *DN 25 COMVAT* bellows (*D*) with a maximum stroke of 900 mm which at once determines the piston's maximum traveling distance. The bellows is traversed by a 1.3 m long stainless steel rod that holds the Teflon plug. At the positions *e* and *f*, end switches limit the motion in the case of step losses, software malfunction, etc. A third end switch at position *g* serves as feedback for the detector switch: only when the piston is completely retracted from the system switch *g* is closed. When *g* is opened the wire connection of the detector switch is mechanically disconnected and the switch will automatically move into its piston position.

At a step-size of 1.8° we could run the system at 2500 steps s^{-1} , resulting in a translation velocity of 50 mm s^{-1} . Higher velocities should be feasible but have not been tested. In order to avoid step losses and the resulting misalignments the motor can also be switched to a closed-loop mode. The necessary decoder unit for the feedback-loop is already installed on the stepper's rear side.

Absorber

The absorber (bright-red in the sketch on the left-hand side of figure 2.14) enters the trapping region from the top. It is a 120 mm long polyethylene cylinder with an outer diameter slightly smaller than the bore's inner diameter in order to allow for a smooth displacement. On its bottom side the cylinder has a conical 100 mm deep cavity. Its inner surface is tapered and adapted to the trap's conical equipotential surfaces (see figure 2.9). The absorber is *work in progress* and its current design only its first iteration.

For the counting-the-dead measurements the absorber scheme will have to be adapted as, in its current design, the absorber would completely blocks the decay products' access to a charged-particle detector. Electron-track simulations showed us that the electron trajectories are avoiding certain regions on their way up to the detector. A projection of the trajectories onto the x-y-plane showed a cross-like pattern rather than a uniformly filled circle. Consequently we only have to limit the absorber's body to the "holes" in the electron pattern in order to grant free electron passage to the detector. A more detailed discussion of the tracking simulations together with an illustration (figure 2.18) of the

2. Design and systematical aspects of the experiment

electron pattern will be given in the next section.

In contrast to the piston mechanics the absorber automation is not completely custom made but recycles a fully functional linear table from *ISEL Germany AG* (see figure 2.14). The table uses an *ISEL MS160W* stepper motor and has a maximum travel distance of ~ 300 mm. The motion feedthrough is once again done using a *COMVAT* bellows with 300 mm axial stroke. The absorber is hold by a 1.8 m long aluminum rod that enters the vacuum system through a O-ring crimp terminal (detail 1 of figure 2.14). The O-ring feedthrough allows to freely vary the absolute absorber position with the total length of the aluminum rod as sole limitation. The ample adjustability is necessary as the high-point of the trapping potential varies extensively with coil currents. Furthermore we only have an approximate idea of the absorber's ideal working point. Thus the overall procedure will be to roughly adjust the absorber's absolute position by hand before starting an automated fine-adjustment using the linear table.

2.2.3. Upgrades and variations

B-field mapper

Before they were stacked together the magnetic fields of all 12 permanent magnet modules were measured separately without finding any grave magnetic defects, a summary of those measurements can be found in [24]. A mapping of the entire B -field, including the SCS, is pending and can still provide a great deal of new information:

- The measurements will be an empirical cross-check of the *Radia* calculations presented in section 2.1.3.
- It is a control measurement to make sure that the magnets' performance has not been corrupted e.g. by mechanical shocks, corrosion or strong external magnetic fields.
- We will measure the system's temperature dependence and determine the ideal magnets' temperature for our measurements.
- We will enhance the spatial resolution compared to the measurements presented in [24]. Those measurements were done using a 3-axis Hall probe from *LakeShore* [19]. The active areas of its three axis B_x , B_y and B_z are pairwise separated by about 2 - 3 mm. Due to the high angular, radial - and close to the shutter coil also vertical - dependence of the magnetic field, this spatial spread leads to non-negligible deviations in the calculation of $|B| = \sqrt{B_x^2 + B_y^2 + B_z^2}$. Furthermore the probe is not rated for temperatures below 10°C. Consequently we plan to use a set of new 3-axis probes from the Slovak manufacturer *Arepro s.r.o.* [1]. The probes can be used from 1.5 to 350 K and all three axis are measured within a $250 \times 250 \times 250$ μm cube.

I already designed the mapping robot for these measurements but finally could not conduct them myself due to the ongoing problems with the switches of the superconducting

2. Design and systematical aspects of the experiment




Figure 2.14.: The left-hand side shows a *SolidWorks* sketch of the absorber mechanics. Right next to the absorber bellows sits the main turbo pump. The intermediate piece between the turbo pump and the flange contains three additional ports that can be used for vacuum gauges, a mass spectrometer, etc. The picture on the right-hand side shows the top flange with all its installations during the 2014 beamtime on PF2.

2. Design and systematical aspects of the experiment

coils. A *SolidWorks* sketch of the system is shown in figure 2.15. Only the parts that are depicted in a bright yellow are new, all other components are “recycled” from the piston mechanics.

The mapper has two automated axis, φ and z , the ρ -axis must be adjusted manually. The translation in z -direction is done using the piston mechanics that has been presented in figure 2.13. For the angular actuation a second stepper motor (an *Orientalmotor PK266M*) is mounted on the rear side of the moving platform (detail 3 of figure 2.15). The bellows from the piston set-up is exchanged by a simple O-ring feedthrough that allows for both, linear *and* rotary movements (detail 2). The detail also shows that the piston’s holding rod is replaced by a 12×2 mm stainless steel tube from *Swagelok*. The Hall probes’ wires are led through the tube and exit the vacuum system through a hole right above the moving platform. Once the wires are installed the hole is vacuum sealed with *Stycast*. The rather complicated wiring guarantees that the cables do not perform any movements relative to the mapping robot *inside* the vacuum chamber. On the outside we have continuous access to the cables and can rapidly intervene if the movements cause any damage.

The most elaborated piece of the mapper, the (Teflon) holding table for the Hall probes, is shown in detail 1. The table contains nine different mounting positions for the probes (red) in total, eight on its top face and one on the side. The eight positions on the top face correspond to eight different radii ($\rho = 0, 10, 15, 20, 25, 30, 35$ and 40 mm) and compensate for the missing ρ -actuation. In order to fix their angular orientation once and for all the probes are not directly mounted on the table but sit in small “orientation units” where their position is fixed by a stud screw. In contrast to the probes the orientation units are not round but have a prolate form and can be quickly permuted without losing their angular alignment.

The horizontal position is special as its radial position is not predetermined but voluntarily floating. The probe sits in a cylindrical Teflon piece with a “dome” on one of its end planes. The Teflon piece is held by four horizontal aluminum pins that define its φ - and z -position while it is still free to move back and forth in ρ -direction. Small springs on the aluminum pins apply a steady force to the Teflon piece and *gently* press its dome face against the bore tube. Ideally the dome and the bore tube only touch in a singular spot at a radial coordinate $\rho_{bore}(\varphi, z)$ that is given by the bore tube’s undulations. Consequently, the Hall probe does not scan the B -field on a given radius but directly on the bore tube’s surface. Thereby the scan yields an “effective maximum B -field” instead of the field on an arbitrarily defined maximum radius. It is important that the springs’ force is not too powerful as we do not want to deform the bore tube but only guarantee a stable contact between the tube and the Teflon piece.

Electron track simulations

My final design had to be compatible with the future installation of a charged-particle detector. Therefore I needed some basic idea of the decay particles’ behaviour.

Kent had already started investigating the electron and proton trajectories in the trap. He had simulated their movement in an ideal octupole field of infinite length. The

2. Design and systematical aspects of the experiment




Figure 2.15.: Sketch of the B -field mapper with some details of its main parts.

2. Design and systematical aspects of the experiment

particles' initial energy was drawn from the electron/proton spectra of a free neutron decaying at rest (the kinetic energy on the UCN scale can be neglected) and the direction of emission was uniformly distributed over the entire solid angle. His calculations were done in *Matlab* using a build-in adaptive-step size *ordinary differential equation* (ODE) solver. A more detailed description as well as an overview on his results can be found in [24].

I refined his code, conducting several modifications. First of all I expanded his classical ansatz by the corresponding relativistic equations:

$$\ddot{\mathbf{r}} = \frac{\pm q}{m_{rel}} \{ \mathbf{E} + \dot{\mathbf{r}} \times \mathbf{B} \}, \quad (2.20)$$

where \mathbf{E} and \mathbf{B} are the electric and magnetic field, q is the elementary charge, \mathbf{r} is the particle's position and m_{rel} its relativistic mass

$$m_{rel} = \frac{m_{classical}}{\left(1 - \frac{v^2}{c^2}\right)^{1/2}}. \quad (2.21)$$

A more important revision was to exchange the infinite octupole field by the B -field maps presented in section 2.1.3. Thereby my simulations do not only take into account the PMS' finite length and deviations from the idealised octupole field close to the walls, but also contain the SCS' field components. Furthermore, I drew the starting points of the decay products from the UCN distribution in the trap under the assumption of a uniform phase-space density and a "SUN2-like" UCN spectrum (i.e. $dN \propto E^{1/2}dE$, see section B). As I calculated the B -field maps and UCN potentials using the *Mathematica* plug-in *Radia* I could not simply import the routines in *Matlab*. In order to avoid gigantic field maps and/or interpolation errors I implemented a communication between the two and used *Mathematica* as an aid program for the *Matlab* solver.

The \mathbf{E} -field in equation 2.20 is only interesting in combination with proton simulations. In order to considerably influence the particles' trajectories the electric-field strength must be comparable to the particles' kinetic energy. The protons that are emitted in a neutron beta decay typically have energies below 1 keV. In contrast, beta electrons have much higher kinetic energies ($\mathcal{O}(E_{kin}) \sim 10^2$ keV) making the field strengths that would be necessary in order to achieve major deflections illusionary. As my main focus was on electrons I only considered magnetic fields and set the electrical field to 0 throughout my simulations.

The most interesting results of the simulations are illustrated in the figures 2.16 through 2.18. As the pictures are quite explicit I decided to let them tell their own tale.

2. Design and systematical aspects of the experiment




Figure 2.16.: Three projections of a typical electron trajectory. The electron is created at around $z = -200$ mm and initially drifts in negative z -direction. Once it reaches the strong field regions around the bottom coil it gets reflected back and starts moving up again. Just like the bottom coil field the PMS' field rejects the electron, keeps it from touching the bore wall and forces it onto a stable trajectory in the x - y -plane. On this trajectory the electron slowly drifts in positive z -direction and eventually reaches the detector plane located at the top end of the permanent magnet system.

2. Design and systematical aspects of the experiment




Figure 2.17.: The green line shows the electrons' probability of reaching the detector ϵ in dependence of the bias field (in the legend ϵ is referred to as *detection probability* by convenience while the actual *d.p.* will also depend on the impact angle, the internal efficiency of the detector, etc.). By increasing B_{bias} from 100 to 300 mT ϵ is tripled while for higher bias fields the effect becomes less important. Contrarily to ϵ , the initial number of UCN per fill N_0 decreases with increasing B_{bias} due to the downsizing of the potential well (see section 2.1.3). The red and the blue curve show the effect for the two spectra $N_{0,1}$ ($dN \propto \sqrt{E}dE$) and $N_{0,2}$ ($dN \propto E dE$). The optimum bias-coil setting maximises the product of ϵ and $N_{0,i}$ and simultaneously the total number of electrons that reach the detector. For both spectra the optimum setting can be expected somewhere between 200 and 300 mT.




Figure 2.18.: Crossing points of 20000 electron tracks with a horizontal plane at $z = 1125$ mm. The insert shows the magnetic field lines in the horizontal plane as calculated with *Radia*. Since the electrons are “confined” on stable trajectories that are defined by the magnetic field shape, they do not uniformly fill the entire trapping region. Their trajectories form a cross-shaped pattern that leaves four spots completely untouched. For the future combination of a UCN absorber with an electron detector this is a major asset since it will allow for the design of an absorber that can permanently reside within the central bore tube without disturbing the electron tracks.

3. The 2014 beam time on PF2

The 2014 beamtime was the very first time HOPE was fully assembled and our primary goal was the successful commissioning of the apparatus. The rest of the time was scheduled for systematic tests and the optimisation of the measurement parameters. An overview on the commissioning phase and some of the technical problems we encountered can be found in section C. In the following chapter I will present and discuss the data we could collect in the second phase of the beamtime, once the apparatus was steadily running. An outlook on the future of HOPE will be given at the end of the chapter.

3.1. Measurement scheme and parameters of interest

All measurements we performed were fill-and-empty measurements and followed the scheme described in table 3.1. Besides the timing information given in the table there are several other measurement parameters that had to be defined, namely the *filling parameters* (the filling time, ramping speed, bias field, etc.) and *cleaning parameters* (absorber and piston position as well as the cleaning time). It seems preferable to optimise the filling parameters first in order to perform all other measurements with the best possible statistics. However it is not the *initial* number of neutrons N_0 that must be maximised but the number of *storable* neutrons N_{stor} . Therefore the optimal filling procedure can not be determined without a reliable cleaning process and I decided to focus the measurements on the cleaning scheme, all other parameters remained constant during the beamtime.

The movement of the inner parts of our set-up is illustrated in figure 3.1. We tested an alternative approach by inverting the chronology of bottom coil ramping and piston movement. We hoped that by pushing the UCN out of the region of maximal B -field strength the spectrum heating effect during the coil ramping would be reduced. Regrettably the result was as clear as disappointing as we did not store any UCN at all when using this approach. As the central bore tube is not precisely round we had to chose the piston's outer diameter somewhat smaller than the tube's inner diameter. It seems that nearly all UCN leave the central bore tube through this small gap before the piston reaches the PMS where the UCN can not approach the side walls anymore due to the radial magnetic field. We could not think about any other promising alternative to the cleaning procedure as proposed in figure 3.1. It would be possible to facilitate the cleaning by holding the bottom-coil field at a lower value during the cleaning. Also the UCN's energy could be already limited before they enter the trap by means of a gravitational UCN spectrometer. However in both cases the spectrum would still have to be post-cleaned as the UCN will always gain energy during the ramping.

3. The 2014 beam time on PF2

Phase	Time [s]	Action	Aim
Filling	95	switch is in detector position after 20 s the switch is closed and the plug is moved to its filling position	determine the filling rate fill the trap
Closing	60	bottom coil is ramped to its maximum value	close the trap
Cleaning	120 or 240	UCN shutter is closed, plug and absorber are moved to their cleaning positions	remove overcritical UCN
Storing	variable: 30 to 2500	plug and absorber are removed, detector switch is opened	observe leakage of overcritical UCN the variable storage time is the main parameter of the fill-and-empty method
Counting	80	bottom coil is ramped to 0	count remaining UCN
Background	100	-	determine background rate cool-down magnet system

Table 3.1.: Overview of the six phases of a typical measurement cycle. An illustration of the movement of the inner parts (piston, absorber and detector switch) is shown in figure 3.1.

3. The 2014 beam time on PF2




Figure 3.1.: Movement of plug and absorber with respect to the UCN potential. The UCN potential is shown for the bottom coil running at 200 A (0 A during the filling) and a bias field of 20 mT. The positions of the absorber and the piston during the cleaning phase varied throughout the beamtime.




Figure 3.2.: Typical measurement cycle as described in table 3.1. Shown is the data from the ^3He detector summed over 8 runs with identical settings for the absorber and the plug movement. The peak in the beginning of the cleaning phase occurs when the plug moves and the remaining UCN in the guide system leak into the detector. The constantly high count rate during the storing is due to overcritical UCN falling out of the trapping region.

Concerning spin-flips we decided not to perform any new measurements. Instead we just set the system to some state that should decrease the spin-flip probability beneath our detection efficiency. Following the discussion in section 2.2.1 this goal could be easily reached. We only had to prevent the formation of an “artificial” zero-field region due to the inductive unloading of the bias coil during the bottom-coil ramping. Therefore we ramped the bias coil to 25 A and left it in persistent mode throughout the measurements. Taking the inductive-coupling constant $C_{bot-bias} = 93 \text{ m I}_{bias}[\text{A}] \text{ I}_{bot}[\text{A}]^{-1}$ from equation 2.3 and a maximal shutter-coil current of 200 A we get a minimal bias-coil current $I_{bias} > 6 \text{ A}$ after the shutter-coil ramping. Given the estimate from section 2.2.1 - $\Delta\tau_{sys,spinflip}(B_{bias} = 10 \text{ mT}) < 0.1 \text{ s}$ - the corresponding bias field $B_{bias} > 25 \text{ mT}$ should be sufficiently high.

There are numerous other parameters that would have been worth investigating such as the ramping speed or the filling time. However due to our rather tight schedule we eventually decided to save those for future beamtimes and fully concentrate our efforts on the cleaning scheme.

3.2. First data from HOPE

Figure 3.2 shows the UCN counts per s of the ^3He detector for a particular set of measurement parameters. While the main interest of the measurements is to determine the number of surviving UCN after a certain time, the emptying peak is far from being the




Figure 3.3.: Spectra of the monitor and main detector. As monitor detector we used NaI scintillator with a ^{10}B layer as neutron converter. The α and the ^7Li -peak from the $^{10}\text{B}(n,\alpha)^7\text{Li}$ reaction are nicely separated in the spectrum. The small peak on the right comes from the reaction without γ -emission. The red curve shows the spectrum of our main detector, a so called Strelkov UCN detector based on the $^3\text{He}(n,p)^3\text{H}$ -reaction. Both, the proton and the Triton deposit their entire energy in the gas chamber and only a single detection peak appears in the spectrum. The discriminator levels used in the recording of the shown spectra are not the same as for the actual data taking.

only exploitable source of information. The peak in the beginning of each cycle yields information on the filling rate while the count rate during the last 100 s is used to determine the background rate. The slowly decreasing count rate during the storage time is due to overcritical UCN falling out of the trapping region. I will start the presentation of the measurement results with the filling and background rates and an explanation of my analysis scheme before giving a detailed review on our spectrum cleaning efforts. In the end I will present results from measurements with long storage times up to 2500 s.

3.2.1. Filling rate, background rate and analysis scheme

Figure 3.3 shows the spectra of our monitor and main detector. The monitor detector was connected to our guide system via a T-piece, sitting relatively close to the UCN turbine (a scheme of the set-up is shown in the lower right corner of figure 3.4). Its purpose was to measure the flux Φ of incoming UCN during the filling phase and thereby deliver a weighting factor for the main detector's data. I later decided to additionally monitor Φ using the main detector and therefore added the short monitoring phase in the beginning of each cycle - the first 20 seconds of the "filling phase" in table 3.1. My main concern was that the data from the monitor detector could be falsified by the unpredictable behavior of the UCN shutter (red in the scheme). During the beamtime the shutter repeatedly

3. The 2014 beam time on PF2

lost its opening position or did just refuse to work at all (more details are given in section C).

Indeed figure 3.4 shows some inconsistency in between the rates extracted from the main and the monitor data. In the data from the main detector occurs a step between the 12th and the 13th of October that does not seem to be present in the monitor data. In return the data from the monitor detector is separated into two distinct bands. The correlation of the count rates from the two detectors clearly depicts the inconsistency of the two data sets (lower left plot in figure 3.4) . Especially the two bands from the monitor detector are completely uncorrelated to the data from the main detector.

The step in the main detector's count rate could be explained by a sudden change of the shutter's opening position. Pushing further one could think of two slightly different opening positions that result in different back-scattering rates and explain the two bands in the monitor count rate. However both scenarios only explain why one of the detectors changes its behavior but can not explain why the other one should be unaffected. Also it is not obvious where the overall decline of the count rate - 1 % over the nine days period for the monitor and 4 % for the main detector - comes from. I cross checked it with the reactor power but it was only reduced by ~ 0.3 % in the same period of time.

Due to the continuous trouble we encountered using the UCN shutter and the big inconsistencies in between the two data sets I ignored the data from the monitor detector and based the weighting of the individual cycles solely on the main detector. For the aim of simplicity I will simply refer to it as *the* detector from now on as there is no more use for any distinction.

Due to my distrust in the questionable behavior of the UCN shutter I decided to constantly monitor the evolution of the background rate at the end of each cycle rather than performing dedicated background-measurement runs. Before we could resolve the problem with the shutter the background rate was massively fluctuating and sometimes rose up to several Hz. The fluctuations were obviously caused by an increase of the UCN leakage rate through the shutter as we only had to readjust its closing position in order to restore the normal (low) background rate. By constantly monitoring the background rate one could have seen immediately if the shutter started losing its position again or at least easily sort out the corrupted measurements during the data analysis. Fortunately the background rate remained stable to a passable degree as can be seen in the left plot of figure 3.5 that shows the day-by-day evolution of the background rate during the last nine days of the beam time. The mean background rate during this period was of $0.0318(8)$ counts s^{-1} .

The overall background rate is the average of two distinct background rates as can be seen in the right plot of figure 3.5. The plot shows a histogram of the number of counts during the background observation period. The red bars show the distribution for the entire data set (517 runs in total) while the green and the blue bars separate the data set depending on whether the *turbine switcher* moved or not. On PF2 there are three beam lines - the EDM-, the MAMBO- and the UCN-position - that are all fed by the same primary guide. In the beginning of each measurement cycle the respective beam line is connected to the primary guide via the turbine switcher. Once the filling period is over the switcher will remain in its current position until a request from any of the other

3. The 2014 beam time on PF2




Figure 3.4.: The top row shows the evolution of the detector and the monitor count rate during the last 9 days of the beam time (9th to 17th of October). The y-axis gives the count rate per second averaged over the observation period (between 10 and 20 s after the start of each cycle for the detector and during the first 150 s for the monitor count rate). The color code gives the number of events per bin were an event means a single measurement. The bottom plot on the left shows the correlation between the two rates, with the monitor rate given on the x-axis. The sketch in the lower right corner has been added for illustrative reasons. It shows the positions of the two detectors during the beamtime.

3. The 2014 beam time on PF2




Figure 3.5.: Evolution of the background rate during the measurements. The left plot shows the day-by-day evolution of the background rate with the dotted line depicting the error boundaries of the background mean value. The right plot shows the distribution of the number of background counts over the different measurement cycles. The three bar graphs show the slightly different distributions for different positions of the turbine’s switcher (see text).

positions comes in. Once the switcher moved away the beam line is not hermetically sealed but the number of UCN entering the guide will be reduced considerably. As our monitor detector was placed in between the turbine and our UCN shutter it constantly kept track of the incoming flux. With the turbine switcher showing in our direction we got monitor rates $R_{monitor,1} > 1200 \text{ counts s}^{-1}$, once the switcher moved away the rate dropped by a factor of 50 to $R_{monitor,2} \approx 25 \text{ counts s}^{-1}$. As our UCN shutter is not perfectly UCN tight this big discrepancy also reflects in the slightly changed background rates of figure 3.5. Since $\frac{R_{monitor,2}}{R_{monitor,1}} \ll \frac{R_{background,2}}{R_{background,1}}$ we can conclude that most of the background can not be caused by UCN coming from the turbine. Probably more rigid discriminator levels could have further reduce the background rate.

In my data analysis I distinguished the two scenarios by means of the monitor counting rate in the beginning of the counting phase and used the according background rate in the data treatment. This approach could still be refined but in the following sections we will see that our overall error is dominated by the very low statistics due to low count rates and few measurement cycles.

Having defined the filling and the background rate I extracted the mean number of UCN per measurement point N_p and its associated error δN_p using the following formulas:

$$N_p = \frac{\sum_{i=1}^{n_p} (N_{i,p} * w_{i,p})}{n_p} - Background_p \quad (3.1)$$

3. The 2014 beam time on PF2

$$\begin{aligned} \delta N_p &= \sqrt{\left(\frac{s_p}{\sqrt{n_p}}\right)^2 + (\delta Background_p)^2} = \\ &= \sqrt{\left(\frac{\sqrt{\frac{1}{n_p-1} \sum_{i=1}^{n_p} (N_{i,p} * w_{i,p} - \overline{N_{i,p}} * w_{i,p})^2}}{\sqrt{n_p}}\right)^2 + (\delta Background_p)^2} \end{aligned} \quad (3.2)$$

The designation ‘‘measurement point’’ stands for the totality of all measurements that have been done using the same parameter set p . The parameters in question are the absorber and the piston position, the cleaning time, the storage time and the bottom-coil current during storage. The number of cycles per parameter set is given by n_p . The weighting $w_{i,p}$ is determined for every single measurement and calculated as $w_{i,p} = \frac{w}{fillingrate_{i,p}}$, where the filling rate is the mean of the detector count rate in between the 10th and the 20th second of each cycle and w is the mean filling rate of all cycles. The number of stored neutrons per cycle $N_{i,p}$ is the sum over a well defined region within the counting phase. The limits of the region of interest are shown in the left plot of figure 3.6. $Background_p$ is individually determined for every single measurement as it depends on whether or not the turbine switcher has been moved.

The individual error for every measurement point is calculated as the *standard error of the mean* and thus given as the *sample standard deviation* s_p divided by the square root of the respective number of cycles n_p . The error due to the background subtraction is calculated separately and added to give the final error δN_p .

The right plot of figure 3.6 serves as a cross check for the reliability of the data. It shows the deviation of the single measurements $N_{i,p}$ from their parameter set’s average value N_p normalised by the sample standard deviation: $\frac{N_{i,p} - N_p}{s_p}$. I chose this representation as the number of repetitions per parameter set was very small - in most cases less than 10 - and therefore the distributions of the single sets would not have been significant. Taking the totality of the measurements we can see that the data behaves like one would presume for a random variable and forms a curve of nearly Gaussian shape.

3.2.2. Spectrum cleaning measurements

Optimising the spectrum cleaning is a very challenging task. Our ultimate goal will be to maximise the number of storable UCN after the cleaning N_{stored} while reaching a ratio of storable to overcritical UCN $\frac{N_{stored}}{N_{crit}}$ on the order of 10^{-3} (see 2.2.1). The three parameters at our disposal in order to reach that goal are the respective cleaning positions of the UCN absorber and the piston as well as the cleaning time. In the same time we have to find a reliable way to determine $\frac{N_{stored}}{N_{crit}}$ to the needed precision.

We should first have a look at the last point as it determines how the measurements have to be performed. I came up with the idea to roughly determine the storage-time constant τ_p for each parameter set p using only two measurement points at $t_{short} = 30$

3. The 2014 beam time on PF2




Figure 3.6.: The left plot shows the emptying peaks summed up over all measurement cycles and separated into runs with bottom coil running at 200 respectively 300 A. The curves look different as the bottom coil’s maximum value influences the emptying time. Consequently the region of interest has been extended for the higher closing field. The right plot shows the distribution of the stored neutron numbers around their individual mean normalised with their individual standard deviation σ_p .

s and $t_{long} = 480$ s. An improved cleaning procedure should lead to a longer τ_p and ultimately we should end up at the actual neutron lifetime. By definition this approach can not enable us to measure $\frac{N_{stored}}{N_{crit}}$ with the needed precision as it uses τ_n to estimate the *absolute* cleaning efficiency and would therefore lead to a circular reasoning. However it is perfectly well suited for a first *comparative* approach.

Unfortunately the measurements are quite time consuming. Making ten runs at both storage times for one particular parameter set takes around four hours while $\delta\tau_{p,stat}$ is still on the order of several hundreds of seconds. In the time available a measurement of $\frac{N_{stored}}{N_{crit}}$ with a meaningful precision for a bigger set of cleaning parameters was not achievable. Nevertheless the data we produced delivered significant insights.

Figure 3.7 shows the number of UCN left after 30 s of storage N_{30} depending on the chosen parameter set. In order to facilitate the readability the data is subdivided into four main groups that are distinguished by their respective piston position and cleaning time. The absorber position on the x -axis is defined as in figure 3.1, the lines connecting the different data points are only included to guide the eyes. There is a clear hierarchy between the four main groups. The upmost curve shows the cleaning results without use of the piston and 120 s of cleaning time. For absorber positions above ~ 650 mm the cleaning has hardly any effect. This height roughly corresponds to the upper limit of the storage region as shown in figure 3.1. Once we start cutting into the storage region the number of surviving UCN decreases almost linearly with the absorber position.

The comparison of measurements with different cleaning times allows us to conclude on the cleaning-time constant τ_{clean} . If τ_{clean} is considerably smaller than 120 s the

3. The 2014 beam time on PF2




Figure 3.7.: The plot shows the number of UCN detected after 30 s of storage for various settings of the spectrum-cleaning parameters.

absorption probability of the UCN that are still alive after 120 s should be small and the ratio $\frac{N_{30,120s}}{N_{30,240s}}$ should be dominated by the neutron decay: $\frac{N_{30,120s}}{N_{30,240s}} \approx \exp^{-120s/\tau_n} \approx 0.85$. As the ratio is significantly smaller we can conclude that the absorber is still removing UCN after 120 s and that τ_{clean} must be relatively long. By entering the piston into the storage region we hoped speeding up the cleaning process by mixing the UCN trajectories via diffuse reflections. Indeed the piston drastically increased the cleaning efficiency but not in the desired manner. τ_{clean} is still relatively long but the number of surviving UCN is dramatically reduced. Even with a fully retracted absorber the piston removes about 50 % of the initial UCN within 120 s.

The UCN leakage rate during the storage phase is another rich source of information. Figure 3.8 shows the time evolution of the leakage rate for different cleaning settings. The background rate has been subtracted and the data has been piecewise averaged over 60 s bins. The averaging has no bigger effect on the results but renders the plots more readable as it reduces the number of data points and the size of the error bars. The uncleaned UCN spectrum starts with a very high leakage rate of about 250 mHz and decays rapidly. The very thoroughly cleaned spectrum has a leakage rate that is at least an order of magnitude lower and gets compliant with 0 after 200 to 300 s. The right plot shows the leakage rate for the uncleaned spectrum over a longer period of time. The data seem to be in good agreement with the fitted exponential decay curve which suggests a storage-time constant τ_{crit} of about 250 s. As one would expect the leakage rate drops to 0 after some 4-5 τ_{crit} . It is not self-evident that the leakage rate can be described by a single-exponential decay as τ_{crit} should depend on the total energy of every individual UCN.

The leakage rate can also be used to determine the total number of overcritical UCN N_{crit} . If we only consider UCN losses due to leakage and free neutron decay, τ_{crit} can be

3. The 2014 beam time on PF2




Figure 3.8.: The two graphs show the trap’s leakage of overcritical UCN per second during the storage phase. The left plot shows the evolution of the leakage for three different cleaning scenarios. Each point is an average over a 60 s period. The right plot shows the leakage in the case of an uncleaned spectrum. The red line is an exponential fit to the data, the fitting parameters are shown in the plot. In both plots the background has been subtracted and the first 30 s of the storage phase have been omitted as the detector entrance remains blocked until the plug is fully retracted and the detector switch opens.

calculated using

$$\tau_{crit}^{-1}(E_{tot}) = \tau_{leak}^{-1}(E_{tot}) + \tau_n^{-1}. \quad (3.3)$$

where τ_n is the neutron lifetime and $\tau_{leak}(E_{tot})$ the energy dependent leakage rate. For the aim of simplicity I will neglect the energy dependence for now. 3.3 then reduces to

$$\tau_{crit}^{-1} = \tau_{leak}^{-1} + \tau_n^{-1}, \quad (3.4)$$

If we want to calculate the leakage rate at time t , $\dot{N}_{leak}(t)$, as a function of the overcritical-decay rate $\dot{N}_{crit}(t)$, we have to take into account the branching ratio $\frac{1/\tau_{leak}}{1/\tau_{crit}}$:

$$\dot{N}_{leak}(t) = \frac{N_{crit,t=0}}{\tau_{crit}} \exp^{-t/\tau_{crit}} \frac{1/\tau_{leak}}{1/\tau_{crit}}. \quad (3.5)$$

By solving 3.5 for N_{crit} and setting $t = 0$ we get

$$N_{crit,t=0} = \dot{N}_{leak}(t = 0) \tau_{leak}. \quad (3.6)$$

If we want to use the fitting parameters from figure 3.8 in equation 3.6 we still have to solve equation 3.4 for τ_{leak} as the storage-time constant given by the fit is the overall decay constant τ_{crit} . Setting τ_n to 880 s and using the fitting parameters from figure 3.8 ($\tau_{crit} = 257 \pm 17 \text{ s} \rightarrow \tau_{leak} = 363 \pm 33 \text{ s}$, $\dot{N}_{leak,t=0} = 0.28 \pm 0.02 \text{ s}^{-1}$) we get an initial

3. The 2014 beam time on PF2




Figure 3.9.: On the left: number of overcritical UCN as function of the chosen cleaning parameters. The N_{crit} values have been calculated using equation 3.7. On the right: an estimate of the amount of well-storable UCN “killed” during the cleaning process. In both plots the data points are grouped in the same manner as in figure 3.7. As the error bars are partially overlapping the x-axis position of the individual data points has been slightly spread around the actual position in order to facilitate the readability.

number of overcritical UCN of $N_{crit,t=0} = 102 \pm 12$ for the uncleaned spectrum¹. While this approach may be well suited for the high leakage rates of none or poorly cleaned spectra it does not work for low leakage rates where we can hardly perform a useful fit. In this case we can resort to a slightly more involved scheme that calculates $N_{crit,t=0}$ from the total number of UCN having left the trap in between $t' = 0$ and $t' = t$:

$$\Delta N_{leak}(t) = \int_{t'=0}^{t'=t} N_{crit,t=0} \exp^{-t'/\tau_{crit}} \frac{1/\tau_{leak}}{1/\tau_{crit}} dt' \longrightarrow$$

$$\longrightarrow N_{crit,t=0} = \Delta N_{leak}(t) \frac{1}{1 - \exp^{-t/\tau_{crit}}} \frac{\tau_{leak}}{\tau_{crit}}. \quad (3.7)$$

The derivation of 3.7 follows the same scheme as before. For its evaluation I again set τ_n to 880 s. For τ_{crit} I made a (hopefully decent) guess setting it to 200 s. $\Delta N_{leak}(t)$ is the total number of counts during the storage phase after background subtraction.

The left plot of figure 3.9 shows the initial number of overcritical UCN in dependency of the cleaning parameters as calculated using equation 3.7. For the uncleaned spectrum it yields $N_{crit} = 87 \pm 5$ which is in reasonably good agreement with the N_{crit} value we got from equation 3.6. The overall appearance of the plot is similar to figure 3.7: the

¹As explained in the caption of figure 3.8, for this and the following calculations, t_0 is shifted by 30 s compared the starting of the cleaning phase as defined in table 3.1.

3. The 2014 beam time on PF2

four groups follow the same hierarchy and their overall tendency is decreasing for a more severe cleaning phase. Still there is some new information to extract.

For the cleaning procedure corresponding to the blue curve we have reached the zero-contamination level within the (large) uncertainties. Indeed for a very low absorber position of 460 mm all groups show good results. Especially interesting in this context is the behavior of the black curve which is obtained without the piston. Linear extrapolation of this curve suggests that it should reach the zero-contamination level for an absorber position of 420 mm. Plot 3.7 showed that the number of surviving UCN after the cleaning process was much higher when we did not use the piston for cleaning. Consequently, cleaning without piston might be most advantageous, remembering that the number to be optimised is the ratio of overcritical to well-stored UCN. However the high statistical uncertainties do not allow final conclusions.

If we combine all information from the past section we can also gain a better understanding of how the cleaning with and without piston differ. The right-hand plot of figure 3.9 shows an estimate of the number of well-storable UCN that get lost during the cleaning phase. It is calculated as

$$N_{lost} = N_{storable} - (N_{30} - N_{crit}), \quad (3.8)$$

where $N_{storable}$ is an estimate for the number of well-stored UCN we would have in the trap without cleaning losses and N_{crit} is the number of overcritical UCN left in the trap after cleaning as calculated by equation 3.7. I estimated $N_{storable}$ from the data for the uncleaned spectrum giving a theoretical number of $N_{storable} = 85 \pm 9$. The error bars shown in the plot are no absolute errors but only relative to the other data points as they do not include the (common) offset $\delta N_{storable} = \pm 9$.

The plot suggests that the usage of the piston during the cleaning phase adds an offset to the loss of well-storable UCN. On average the loss of potentially well-storable UCN is by about 36 UCN higher with piston than without piston, for otherwise same settings. For now the piston has been rather disappointing. It considerably reduces the final number of storable UCN and does not seem to significantly speed up the the removal of the overcritical ones. However there still is room for improvement. It is possible that the high UCN losses are mainly due to Doppler heating and that a change in the piston's transition speed has some corrective effect. In order to improve the efficiency of the cleaning process one could rise the piston further which should result in a higher collision rate and thus an acceleration of the UCN's mode mixing.

3.2.3. First attempts to determine the storage-time constant

In order to get a first estimate of the trap's storage-time constant τ_{stor} I also performed a couple of measurements with considerably longer storage intervals ranging up to almost 2500 s. In total I took three different data sets that differ in their respective cleaning processes. In the following I will refer to the three sets as the *highly cleaned data* (absorber at 500 mm, piston at 160 mm and $T_{clean} = 120$ s), the *medium cleaned data* (as before but without piston) and the *uncleaned data* (120 seconds of waiting without absorber or piston).

3. The 2014 beam time on PF2




Figure 3.10.: Storage measurements with five, respectively six storage times for different cleaning parameter sets. The data sets with less well cleaned spectra have been fitted once with a double-exponential-decay law and once with a single-exponential neglecting the data point with the shortest storage time. The plot in the upper right corner gives an overview of all data presented here.

3. The 2014 beam time on PF2

The left column of figure 3.10 shows plots of the highly, medium and uncleaned data sets. I could fit the highly cleaned data with a single-exponential decay law but the medium and the uncleaned data have a much higher overcritical UCN ratio and I had to use a double-exponential fit in order to get a reasonable match. I included an alternative analysis of the medium and uncleaned data that is shown in the right column of figure 3.10. For those plots I did not consider the data point with the shortest storage time ($T_{stor} = 30$ s). The residual data points could then be fitted with a single-exponential law since the overcritical UCN ratio gets automatically smaller for longer storage times.

Let's first have a look at the double-exponential fits. The fit to the uncleaned data yields $N_{storable}$, N_{crit} and τ_{crit} values that are in relatively good accordance with the calculations from the last chapter. In contrast, the double-exponential fit to the medium cleaned data is incompatible with the former calculations. As to be expected for fits with more free parameters, the fitting errors, especially of the medium cleaned data, are considerably higher than the errors of the single-exponential fits. It gets evident that the double-exponential fits will not allow us to precisely separate the overcritical from the storable UCN but are only a very hand wavy tool that can be used for rough estimations.

If we only consider the single-exponential fits we see that the τ_{stor} values of the medium and highly cleaned data are both consistent with the approved world average of the neutron lifetime $\tau_{PDG} = 880.3(1.1)$ s [30]. From

$$\tau_{stor}^{-1} = \tau_n^{-1} + \tau_{other}^{-1}, \quad (3.9)$$

where τ_{other} is the sum over all UCN loss channels except beta decay, we can calculate an average ‘‘collateral loss rate’’ $\tau_{other}^{-1} = -0.4 \pm 4.4 \times 10^{-5} \text{ s}^{-1}$, a very promising first result! In contrast the uncleaned data comes with a considerably shorter storage-time constant that is separated from τ_n by about 5σ . Apparently even after 480 seconds it still contains a significant amount of overcritical UCN that drag down the average storage time of the ensemble.

An overview of the data is given in table 3.2. Additionally to the data from the fits it also contains some alternative values that have been calculated from the leakage rate: the shown N_{crit} values are given by equation 3.7 and $N_{storable}$ has been calculated as

$$N_{storable} = (N_{30} - N_{crit}) \times \exp^{30 \text{ s} / \tau_n}. \quad (3.10)$$

As in the last section N_{30} is the average number of UCN left in the trap after 30 seconds of storage. The τ_{crit} values are obtained by fitting a single-exponential to the leakage rate as demonstrated in figure 3.8 (for the highly cleaned spectrum that approach did not yield an exploitable result). The $\frac{N_{crit}}{N_{storable}}$ ratios given in the table are all calculated from the N_{crit} and $N_{storable}$ values taken from the leakage-rate measurements. For the single-exponential fits of the medium and uncleaned data the ratios had to be corrected:

$$\frac{N_{crit,480s}}{N_{storable,480s}} = \frac{N_{crit,0s} \exp^{-480 \text{ s} / \tau_{crit}}}{N_{storable,0s} \exp^{-480 \text{ s} / \tau_n}}. \quad (3.11)$$

where, as in the previous section, I set $\tau_{crit} = 200$ s.

3. The 2014 beam time on PF2

When we compare the data from the fits and from the leakage rate analysis we will find several inconsistencies. The $\frac{N_{crit}}{N_{storable}}$ ratios of the three single-exponential fits are a good example. In the precedent paragraph we have seen that the uncleaned data has a considerably shorter τ_{stor} than the other two and we concluded that this must be due to its higher $\frac{N_{crit}}{N_{storable}}$ ratio. However following table 3.2 the uncleaned and the highly cleaned data have rather similar $\frac{N_{crit}}{N_{storable}}$ ratios and only the medium cleaned data has a significantly lower one.

When we look at the double-exponential fits we also find contradictions. While the leakage-rate analysis and the fit to the uncleaned data are in good agreement, the two analysis schemes give vastly different results for the medium cleaned data.

Given the (very) high statistical uncertainties these discrepancies are not yet dramatic. However they can be seen as a first proof for the insufficiency of the “one storage-time” model I use to describe the overcritical UCN portion. As I have already mentioned in section 3.2.2 τ_{crit} is not a constant but a function of the UCN energy. It is not surprising that different cleaning procedures will lead to different energy spectra of the overcritical UCN and we can explain the observed discrepancies as follows:

From the examination of the leakage rate it is evident that the uncleaned spectrum contains a large amount of high energetic UCN with a relatively short τ_{crit} on the order of 200 s. As their storage time is very different from the storage time of the *undercritical* UCN the double-exponential fit is able to clearly discern the two ensembles and thus gives τ_{crit} and N_{crit} values that are in good accordance with the values we calculated from the leakage data. When we take away the data point with the lowest storage time we reduce the overcritical UCN portion to a point where the double-exponential fit can no longer distinguish the long- from the short-lived UCN and we achieve better results using a single-exponential fit. However the amount of short-lived overcritical UCN remains high and shifts the ensembles mean storage-time constant to lower values.

In the case of the medium cleaned we start with a reduced number of short-lived overcritical UCN as compared to the uncleaned data. Apparently this leads to a situation where the double-exponential fit can no longer discern two different storage-time constants. However the overcritical UCN rate must still be too high to allow for a decent single-exponential fit. Consequently the single as well as the double-exponential fit both yield unsatisfying results. By omitting the first data point we drastically reduce the amount of short-lived overcritical UCN and end up with a mean storage-time constant that is close to τ_n .

In order to concretise those thoughts it would be of great help if we could directly measure the UCN spectrum. Fortunately our measurement protocol already contains sort of a spectrum measurement: as the trap can not be instantly opened but the shutter coil is slowly ramped down, faster UCN leave the trap earlier and the shape of the trap’s emptying peak should contain some information on the UCN spectrum. Emptying peaks of the three data sets of figure 3.10 are shown in figure 3.11. Indeed they are quite different and seems to support my former reasoning. The emptying curve of the highly cleaned data has its main peak at the far right meaning that it mainly contains low energetic UCN that only leave the trap once the coil is almost entirely switched off. In contrast

3. The 2014 beam time on PF2

	no cleaning	only absorber	plug & absorber
single-exponential fits			
τ_{fit} [s]	731 ± 30	886 ± 51	881 ± 45
$N_{storable}$	128 ± 8	73 ± 4	50 ± 2
$\frac{N_{crit}}{N_{storable}}$ [%]	16 ± 3	5 ± 1	14 ± 7
double-exponential fit			
τ_{fit} [s]	763 ± 57	1061 ± 197	-
$N_{storable}$	117 ± 18	56 ± 15	-
τ_{crit} [s]	202 ± 99	233 ± 83	-
N_{crit}	69 ± 16	63 ± 15	-
$\frac{N_{crit}}{N_{storable}}$ [%]	101 ± 16	35 ± 8	-
calculated from leakage rate (see last section)			
$N_{storable}$	85 ± 9	81 ± 9	43 ± 5
τ_{crit} [s]	257 ± 17	381 ± 51	-
N_{crit}	87 ± 5	28 ± 3	6 ± 2

Table 3.2.: Overview on the number of storable and overcritical UCN as extracted from the exponential fits and from the calculations presented in the last section. The exact calculation methods of the shown data are given in the text.

the emptying curve of the uncleaned data is much more uniform as it contains also a large amount of high energetic UCN that leave the trap already at the very beginning of the ramping process.

The right column illustrates the time evolution of the spectrum. It shows the ratio $\frac{N_{t=480s}}{N_{t=30s}}$ as function of the emptying time. The dashed line gives the theoretical ratio in the case were the only loss channel was the free neutron decay. The values of the uncleaned emptying curve stay beneath the theoretical ratio during the entire emptying period and show a strong disequilibrium in favor of later emptying times. Again this is a clear indication for a high energetic UCN portion with a short storage-time constant.

In case of the medium and the highly cleaned data the situation is different as there are several points that lie *over* the dashed line. The error bars of the affected points are huge and they are all lying within a $1\text{-}\sigma$ region of the theoretical line and it could all just be a statistical fluctuation. If it was not the only possible explanation is a systematic heating of the UCN spectrum. A possible source for this could be Doppler heating due to the mechanical vibrations of the cryostat (see section 2.2.1). A huge heating effect seems to be unlikely as it would be contradictory to the long storage-time constants I have just presented. However the plots emphasize how important it will be too thoroughly investigate the UCN spectrum and its time evolution.

3. The 2014 beam time on PF2




Figure 3.11.: The plots in the left column show the count rate per second during the emptying phase. As the shutter coil current is slowly ramped down the faster UCN can leave the trap earlier than the slower ones and thus the shape of the emptying curve should contain information on the UCN spectrum. The right column shows the time evolution of the ratio $\frac{N(480s)}{N(30s)}$ for the same time interval. Comparably to the plots on the left this representation should show the spectrum's time evolution. The horizontal dashed-line shows ratio expected for neutron decay alone, i.e. $\frac{N(480s)}{N(30s)} = \frac{\exp^{-480/880}}{\exp^{-30/880}} = 0.6$.

3.3. Prospects

3.3.1. Measuring at 300 Amps

Our experiment was handicapped by the dysfunctional switches of the superconducting coils and the resulting limitation to 200 amps of the shutter coil. Leaving the shutter coil's power supply permanently running was no option due to the very high LHe consumption as for the continuous rise of the HTS' temperatures (see section A). However it seemed perfectly arguable to at least run at 300 A for some hours in order to get an impression of the future gain we can expect.

The measurement scheme I chose for those measurements was the same as for the former measurements. I set the cleaning time to 120 s but as the dimensions of the storage region will change with the bottom-coil current my former measurements did not allow me to draw direct conclusions on how to position the piston and the absorber at the new coil current. In order to produce comparable results I decided not to use either of them. In total I did 12 runs, six with 30 s and another six with 480 s of storage time.

Using the same analysis scheme as in section 3.2.2 the measurements gave an initial number of storable UCN of 333 ± 15 . Compared to the uncleaned spectrum at 200 A ($N_0 = 85 \pm 9$) this is a gain of 291 ± 45 %. For the overcritical portion we got 72 ± 5 UCN with a storage-time constant of 233 ± 37 s (see figure 3.12) which is almost identical to results we got from the 200 A measurements.

There is a rather big discrepancy between the measured gain factor and the calculations I presented in section 2.1.3. Those calculations had predicted a gain of a bout 1000 % when passing from 200 to 300 A. This discrepancy could be due to a strong deformation of the UCN spectrum during the 2014 beamtime. On PF2 the trap's potential minimum was situated about 1.3 m above the primary neutron guides resulting in a gravitational potential difference of about 130 neV. With the shutter coil running at 300 amps we can store UCN in an energy range between 0 and 60 neV relative to the trap's potential minimum. In a first approximation - that does not take into account the UCN heating during the coil ramping - those UCN would thus have kinetic energies between 130 and 190 neV in the primary guides. As those guides are made from stainless steel ($E_{pot} \approx 185$ neV) the high energy part of the storable spectrum lies above the guides' Fermi potential and we must expect high transport losses in that energy range. When the bottom coil is running at 200 A this effect should almost vanish as the trap's cut-off energy will be decreased by ~ 30 neV.

This interpretation surely is speculative and the discrepancy between the measured values and the predictions could just be a result of the oversimplified simulation. However there are strong indications that the actual gain factor will be > 3 and it must be a predominant concern to exchange the failing switches before starting any more data taking.

3.3.2. Statistical aspects

Compared to material-UCN traps HOPE has a rather small effective volume and consequently one of our biggest concerns are statistics. During the last beam time the situation

3. The 2014 beam time on PF2




Figure 3.12.: UCN leakage with the bottom coil running at 300 A and uncleaned spectrum. The representation is the same as in figure 3.8.

has been aggravated by the very tight time schedule but with the current set-up even 50 days of data taking will not get us close to our ± 1 s accuracy goal.

In section 3.2.3 I presented data from long storage-time measurements and their respective fitting parameters. The data set with the lowest error on the storage-time constant was the *highly cleaned data* with a fitting error of $\delta\tau_{stor} = \pm 45$ s. Under consideration of the dead time between two measurement cycles the total data-taking time for the set was of approximately 22.3 h². If we presume that the error is purely statistical and thus goes with $1/\sqrt{N_{tot}}$, where N_{tot} is the number of trapped neutrons during the entire data-taking period, we have to trap approximately 2100 times as many neutrons as we did in order to get down to $\delta\tau_{stor} = \pm 1$ s. In terms of neutrons per fill in case of an entire cycle of data taking, i.e. 50 days, this would require an improvement by a factor of 38 or about 3800 UCN before respectively 1900 after the cleaning process (presuming that we lose about 50 % of the storable UCN during cleaning). If we refer to the *medium cleaned data* without the data point with the shortest storage time, the numbers increase accordingly. With an initial fitting error of ± 51 s after 24.6 h of data taking we would require an improvement factor of 57 respectively some 5700 UCN per fill before cleaning.

As a cross check I also performed some Monte-Carlo simulations using the code presented in section D. The results of those simulations can be seen in figure 3.13. It shows the fitting error $\delta\tau$ after 50 days of data taking in dependency of the initial number of

²The precise time span is somewhat arbitrary as it depends on how long we have to wait for the beam in the beginning of each cycle. I set it to 200 s for the calculation of the data-taking time.

3. The 2014 beam time on PF2




Figure 3.13.: Simulated dependence of the fitting error $\delta\tau$ on the number of stored UCN per fill. Each data point is an average over 1000 simulated beamtimes (50 days of data taking per beamtime). The shown lines are $\frac{c}{\sqrt{N}}$ -fits to the data points. The fitting parameter c of the individual fits is included in the plots legend. The bar chart on the right shows the measurement parameters of the three sets.

UCN per fill N_0 . The three shown curves correspond to different simulation-parameter sets. The differences between the sets are illustrated in the bar chart in figure 3.13. The chart shows the number of measurement cycles per storage time for all three sets. For the red and the green curve the distribution are chosen as during the 2014 beamtime. For the alternative set, the blue curve, I unified the number of measurements per storage time and decreased the dead time - the time lost in between two storage intervals - from 650 to 450 seconds. The lines that link the individual data points are $c/\sqrt{N_0}$ -fits to the data. From the fitting parameter c the necessary number of UCN per fill at the break-even point can be easily calculated as $N_0(\delta\tau = 1s) = c^2$, which yields 4500 UCN for the medium cleaned data, 2200 for the highly cleaned and 1450 for the alternative data set.

Of course the comparably high number I get for the medium cleaned data is due to the very disadvantageous choice of the measurement parameters, namely the fact of not considering storage times below 480 seconds. However my Monte-Carlo simulations show that the measurement parameters I had chosen for the highly cleaned data can still be considerably improved which led to a decrease of $N_0(\delta\tau = 1s)$ by almost 35 % in the given case³. The average over the measured and the simulated $N_0(\delta\tau = 1s)$ values for the highly cleaned data is of 2050 UCN. Taking into account a (conservative) estimate of 20 % for the additional improvement due to the optimisation of the measurement parameters we get our final estimate for the necessary UCN number after cleaning: $N_0(\delta\tau = 1s) \sim 1650$ UCN. This is still some 33 times more than we had in the past beamtime.

³It is worth noting that further examinations showed that this improvement is dominated by the redistribution of the measurement points and not by the reduction of the dead time.

3. The 2014 beam time on PF2

Once the superconducting switches have been repaired we will directly gain a factor > 3 but this is not the only possible enhancement:

- For future beam times we will optimise the filling parameters (the ramping speed and the filling time).
- The bias field can be further reduced and the middle coil's behavior during the ramping process can be improved (i.e. maintaining the bias field strength during the bottom coil ramping, see figure 2.8).
- As the system has been optimised for usage on SUN2 the T-piece that connects the UCN guides, the detector switch and the bore tube is designed for 50 mm diameter UCN guides. Compared to the PF2 standard (80 mm diameter tubes) the cross section thus gets reduced by a factor of ~ 2.5 . Putting in an adapted T-piece before returning on PF2 should increase our filling rate even though the enhancement factor would most probably not reach 2.5.
- On PF2 we have to deal with a relatively long distance between the UCN turbine and the cryostat (4 to 5 m). In consequence we have been susceptible to losses at tube connections. We will have to pay some more attention to the guide system during future beam times (use only the guides with the best surface conditions, thoroughly clean the used guides). It would also be a possibility to go for a custom-made guide system. This would have the additional advantage that we could drastically reduce the number of transitions and thus the number of slits in the system.
- Indeed if we can reduce the UCN losses in the system this does not only increase transmission probability but also the UCN storage-time constant in the guide system and thereby the trap's filling rate. In a very simplified model we can describe the filling procedure of the trap by the following differential equation:

$$\dot{N}_{sys}(t) = \dot{N}_{incoming} - \frac{1}{\tau_{sys}} N_{sys}(t), \quad (3.12)$$

where $N_{sys}(t)$ is the total number of UCN in the system at time t , $\dot{N}_{incoming}$ is the number of UCN entering the system from the turbine per second and τ_{sys} is the systems⁴ overall storage-time constant. Equation 3.12 describes the equilibration process of the UCN supply via the turbine and the constant UCN losses in the system. It is solved by

$$N_{sys}(t) = N_{incoming} \times \tau_{sys} (1 - \exp^{-t/\tau_{sys}}). \quad (3.13)$$

Equation 3.13 shows that the trap's filling rate is directly proportional to τ_{sys} .

⁴Here system refers to the entire guiding system plus the trapping region before the shutter coil rampage.

3. The 2014 beam time on PF2

- We must not forget that we are talking about the number of storable UCN that remain in the trap *after* the cleaning process. In the case of the highly cleaned data set about 50 % of the potentially storable UCN are lost in the cleaning process and there are still overcritical UCN remaining in the bunch. Conversely we have seen that alternative cleaning schemes (e.g. without plug usage but a lower absorber position) could erase the overcritical UCN just as good while being less harmful to the storable portion. Whether or not we will be able to provide a sufficiently clean spectrum without losing more than half of our storable UCN must remain unanswered for the moment even though it seems rather unlikely.

It is rather unlikely that those improvements, even combined, will add up to the necessary gain factor and most probably we will not be able to reach our accuracy goal with the UCN densities that PF2 can provide. However we should bare in mind that HOPE has never been conceived for being used on PF2. Our long term prospective has always been to install HOPE on SUN2 for our final measurements. SUN2 has not yet accommodated any experiment and I can not give solid estimates on the UCN densities it will provide. However up to now the results have been encouraging and it seems that there is still room for improvement. A short overview on the source including some recent numbers is given in section B.

3.3.3. Systematic aspects - how to deal with the overcritical UCN

A successful treatment of the overcritical UCN contamination consists of three crucial points:

1. Defining an efficient way to clean the neutron spectrum in the beginning of each measurement cycle.
2. Finding a reliable measure of the ratio of overcritical UCN left in the trap.
3. Get at least an approximate knowledge of the storage-time constants of the overcritical UCN.

In section 3.2.2 I presented promising results concerning the first point and with the leakage rate during the storage phase and the time dependence of the emptying peak I could identify two tools that allow us to examine the overcritical portion of the UCN spectrum. The equations I derived in section 3.2.2 in order to calculate the overcritical UCN number from the leakage rate are obviously very limited but with some more input on the neutron spectrum (e.g. extracted from the emptying peak) it should be possible to refine those equations and gain valuable input helping us to define an efficient cleaning protocol. However if we want to reliably reduce the statistical error caused by overcritical UCN to below 1 second we will very probably have to look out for a more precise measurement scheme as will show the following estimation:

I start from equation 3.6 that established a relation between the neutron leakage rate at the beginning of the storage phase $\dot{N}_{leak}(t = 0)$ and the total number of overcritical

3. The 2014 beam time on PF2

UCN N_{crit} . By dividing equation 3.6 by $N_0 = N_{crit} + N_{storable}$ we get the following relation for the overcritical UCN ratio n_{crit} :

$$n_{crit} = \frac{\dot{N}_{leak}(t=0)\tau_{leak}}{N_0}, \quad (3.14)$$

where τ_{leak} is related to τ_{crit} by equation 3.4. By using $n_{storable} + n_{crit} = 1$ we can transform equation 2.18 from section 2.2.1 into an alternative formulation of n_{crit} :

$$\bar{\tau}_{stor} \stackrel{2.18}{=} n_{storable}\tau_n + n_{crit}\tau_{crit} \rightarrow n_{crit} = \frac{\tau_n - \bar{\tau}_{stor}}{\tau_n - \tau_{crit}}. \quad (3.15)$$

By combining 3.14 and 3.15, replacing τ_{leak} by $\frac{\tau_{crit}\tau_n}{\tau_n - \tau_{crit}}$ and solving it for N_0 we get

$$N_0 = \dot{N}_{leak}(t=0) \frac{\tau_{crit}\tau_n}{\tau_n - \bar{\tau}_{stor}}. \quad (3.16)$$

We interpret $\tau_n - \bar{\tau}_{stor}$ as the systematic error introduced by the overcritical UCN and set it to 1 s. For $\dot{N}_{leak}(t=0)$ we reclaim a signal-to-noise ratio of 10 and thus get 0.32 counts s^{-1} at the measured background of 0.032 counts s^{-1} . For the minimal filling rate N_0 this yields

$$N_0 = \tau_{crit} \times 282 s^{-1}. \quad (3.17)$$

At an overcritical storage time of 200 seconds we would thus already need 56000 (!) UCN per fill in order to still get a decent leakage rate signal at the contamination rate we are aiming for. Even if we reduce our requirement on the signal-to-noise ratio and manage to further reduce the background rate these filling rates will presumably still be out of our reach. Furthermore the techniques sensitivity decreases linearly with the storage time of the overcritical UCN which runs contrary to our needs as we must suppose that the cleaning efficiency also decreases and consequently the amount of overcritical UCN increases with the rising storage time.

Alternatively one could try to determine the overcritical UCN ratio via the spectrum measurements. This will be tricky as we did not provide a tool to perform in-situ spectrum measurements and any ex-situ process will have to correct for the spectrum shifts that occur during bottom-coil ramping. Of course this holds also for the spectrum determination using the emptying peak.

A possibility to avoid those problems would be a ‘‘brute force approach’’, measuring the storage time with two different cleaning schemes and thus two different UCN spectra. If both yield the same result this would be a clear proof that none of them contains a (significant) amount of overcritical UCN as already very small deviations in the mean-storage time of the two spectra should become visible. Even though this procedure would in any case be an important cross-check it is more than doubtful whether it will be practicable as it would further constrict the statistical bottleneck we are already in.

4. Conclusion

In the past, the two predominant experimental techniques to measure the lifetime of the free neutron have been UCN material traps and beam experiments. Both techniques come with their own set of systematic uncertainties. While their respective accuracy has been continuously improved their results disagree. Magnetic UCN traps are an independent way to measure the neutron lifetime and could play a deciding role. Their most prominent systematic error sources are UCN losses due to depolarisation and gas collisions as well as overcritical UCN. In this thesis I presented the magneto-gravitational trap HOPE, demonstrated how it addresses the systematic issues of magnetic UCN storage and discussed first experimental data from the trap's commissioning on the PF2 UCN source.

From the technical side the beam time has been a success in almost any respect. For the very first time we assembled the entire apparatus, commissioned all the individual mechanical pieces and managed to put the entity into a remote-controlled working mode. Towards the end of the cycle we even reached a steady state allowing for an almost continuous operation over a 9 days period. The few remaining problems mostly concerned the hardware communication. Once they will be resolved the experiment could even be run by a single person as the only permanently required user interaction will be the regular exchange of the cryogenic fluids and the occasional change of the measurement parameters.

From a merely experimental point of view there is still more room for improvement. Our major handicap during the beamtime were the very low counting rates. Depending on the chosen cleaning scheme we ended up with $\sim 20 - 100$ UCN per fill and a resulting statistical accuracy $\delta\tau_{stor} \gtrsim \pm 45$ s after 24 hours of data taking. A main reason for the low storage rates were the dysfunctional superconducting switches that limited our storage potential to only 2/3 of its nominal value. The exchange of the switches in combination with various smaller adaptations should increase the initial neutron numbers by a factor $\sim 5 - 10$, possibly even a bit more in a best case scenario. However the decisive push should come from a relocation of HOPE to the new superthermal UCN source SUN2. In its latest test runs SUN2 delivered over 600000 UCN per fill. If we manage to store only 1 % of the available UCN we could already reach our provisional accuracy goal $\delta\tau_{stat} < 1$ s within 50 days of beamtime.

We could extract a first storage-time constant $\tau_{stor} = 881(45)_{stat}$ s that is consistent with the approved world average of the neutron lifetime $\tau_{PDG} = 880.3(1.1)$ s [30]. Its associated loss rate $\tau_{loss}^{-1} = -0.4 \pm 4.4 \times 10^{-5} \text{ s}^{-1}$ is accordingly small. Due to the huge statistical errors it was not possible to experimentally decompose the loss rate into its different channels but I gave some estimates on the expected storage-time shifts $\Delta\tau_{sys,spin,flip}$ and $\Delta\tau_{sys,scat}(p)$. At the given bias field the depolarisation rate should be

4. Conclusion

negligible and result in a shift $\Delta\tau_{sys,spinflip} < 0.1$ s. The UCN losses due to residual gas collisions are less well defined, mainly due to our insufficient knowledge of the gas pressure in the storage region. We aim to further reduce the final gas pressure, well below 10^{-6} mbar, during future beamtimes.

I devoted most of the beamtime to the investigation of the cleaning process and have been able to considerably reduce the amount of overcritical UCN. Unfortunately I found that high cleaning efficiencies were intimately bound to a strong reduction of the number of storable neutrons. In future beamtimes we will have to focus on creating a more favorable balance between the cleaning efficiency and the overall neutron loss. The leakage rate during the storage phase proved being a good tool to investigate the overcritical UCN rate. However estimations suggest that it will not be sufficiently precise for our final accuracy goal ($\frac{N_{crit}}{N_{storable}} \leq 10^{-3}$).

In summary the results of our first measurements are quite promising. We will have to experimentally confirm the estimations on the storage-time shifts $\Delta\tau_{sys,spinflip}$ and $\Delta\tau_{sys,scat}(p)$. By now it seems that we will be able to limit both of them to a few 100 ms or less. In view of the overcritical UCN it will be necessary to further refine our spectrum cleaning and monitoring techniques. The main challenge will be to reliably determine $\frac{N_{crit}}{N_{storable}}$ with the required precision. A possible way could be to explore the energy dependence of the emptying peak in order to draw conclusions on neutron spectrum. In 2014 the significance of the emptying peak has been limited by the small count rates on PF2. However it could become a powerful tool in combination with the increased filling rates we are expecting on SUN2.

Part II.

HOPE - un piège magnétique
pour neutron ultra-froid dédié à
la mesure du temps de vie du
neutron

5. Le neutron et son temps de vie en théorie et expérience

5.1. Un outil varié pour explorer le monde

Les neutrons sont devenus un outil indispensable dans une grande variété de sciences expérimentales. Les biologistes, chimistes et physiciens des matières solides s'en servent pour étudier les caractéristiques structurales et dynamiques de leurs échantillons dans des expériences de diffusion de neutron. En physique fondamentale ce sont les caractéristiques du neutron même qui sont investiguées afin de tester le *modèle standard de la physique des particules* (MS) et chercher des effets au-delà de ce modèle. Notamment la désintégration beta du neutron est une source riche en informations et a engendré un grand nombre d'expériences au cours des dernières décennies.

Une première théorie de la désintégration beta a été proposée en 1934 par Enrico Fermi [11]. Par analogie avec l'interaction électromagnétique, cette *interaction Fermi* était un processus à quatre corps purement vectoriel. Dans ce modèle, le neutron se transforme spontanément en proton en émettant simultanément un électron et un anti neutrino électronique.

Dans les années cinquante, deux expériences pionnières ont changé notre perception de l'interaction faible et amené à une reformulation de la théorie. En 1957 c'était la fameuse *expérience Wu* qui a délivré la première preuve de la violation de la parité dans la désintégration beta [45]. Wu mesurait la distribution angulaire des électrons émis par une source polarisée de ^{60}Co . Elle trouvait que la distribution était fortement asymétrique puisque la plupart des électrons étaient émis à l'encontre de la direction de polarisation. Peu après *l'expérience Goldhaber* allait confirmer le résultat et, qui plus est, démontrer que la violation de la parité dans l'interaction faible est maximale [15]. Basé sur ces découvertes Richard Feynman et Murray Gell-Mann [14], et indépendamment d'eux, George Sudarshan et Robert Marshak [42], ont établi une nouvelle formulation de l'interaction faible, prenant en compte la violation de parité de l'interaction : la *théorie V-A* (vecteur moins axial vecteur). Suivant la notation de référence [44], *l'élément de matrice de la transition* M entre l'état initial et finale prend la forme:

$$M = [G_V \bar{p} \gamma_\mu n - G_A \bar{p} \gamma_5 \gamma_\mu n] [\bar{e} \gamma_\mu (1 - \gamma_5) \nu] \quad (5.1)$$

Un diagramme de Feynman de l'interaction est montré à la droite de la figure 1.1. Contrairement à l'interaction de Fermi le processus n'est plus ponctuel. Un des deux *quarks down* du neutron se transforme en *quark up* avec émission d'un *boson W* virtuel. Quant à lui le boson se désintègre aussitôt, émettant l'électron et l'anti neutrino électronique. Les

5. Le neutron et son temps de vie en théorie et expérience

deux facteurs d'équation 5.1 correspondent aux deux vertex du diagramme de Feynman. Le second facteur est le vertex leptonique où sont créés l'électron et le neutrino. Dans le domaine léptonique la constante de couplage vecteur et axial vecteur sont identiques et la violation de parité est donc maximale. Dans le domaine hadronique, qui décrit la transformation du quark down en up, l'identité des deux couplages n'est plus exacte. Ceci est dû au fait que la transformation ne se déroule pas au niveau élémentaire mais à l'intérieur du neutron. Conséquemment les fonctions d'onde du neutron et du proton doivent être prises en compte ce qui va annuler l'identité exacte. Dans l'équation 5.1 cela s'exprime par les deux nouvelles quantités G_V et G_A , les constantes de couplage vectoriel et axial vectoriel de l'interaction faible. G_V et G_A sont des paramètres libres de la théorie et doivent être déterminés expérimentalement. Un accès possible à ces paramètres est la mesure du *temps de vie du neutron libre* τ_n . La relation entre τ_n et l'élément de matrice M est donnée par la *règle d'or de Fermi*:

$$\tau_n^{-1} = \frac{2\pi}{\hbar} |M|^2 \rho, \quad (5.2)$$

où ρ représente la densité d'état final de l'interaction dans l'espace des phases.

τ_n et la matrice CKM

Une caractéristique de l'interaction faible est qu'elle ne permet pas seulement des transitions entre les quarks up et down mais aussi entre les différentes *générations des quarks*. Le MS contient six différents quarks en tout qui sont regroupés en trois générations : *up* et *down*, *charm* et *strange* ainsi que *top* et *bottom*. On sait maintenant que les états propres de masse des six quarks $|q\rangle$ ne sont pas identiques à leurs états propres en interaction faible $|q'\rangle$. Le mélange entre les différents états est donné par la matrice *Cabibbo-Kobayashi-Maskawa* (CKM) [18]:

$$\begin{pmatrix} |d'\rangle \\ |s'\rangle \\ |b'\rangle \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} |d\rangle \\ |s\rangle \\ |b\rangle \end{pmatrix}. \quad (5.3)$$

Le carré de chaque élément de la matrice donne la probabilité qu'un certain quark soit transformé en un autre quark lors d'une interaction faible. Par conséquent, les constantes de couplage G_V et G_A ne sont pas des quantités fondamentales mais composées:

$$G_V = g_V \cdot |V_{ud}|, \quad G_A = g_A \cdot |V_{ud}|. \quad (5.4)$$

Tout comme G_V et G_A , g_V , g_A et V_{ud} sont des paramètres libres. Toutefois les neuf éléments de la matrice CKM sont liés entre eux par l'unitarité de la matrice: la somme des carrés de chaque ligne doit être égale à 1. Si la somme était inférieure à 1, cela constituerait une preuve directe d'une quatrième génération de quarks. Aujourd'hui la valeur la plus exacte de V_{ud} est déterminée à partir de désintégrations nucléaires *super-permises* [30]: $|V_{ud}| = 0.97425(22)$. Il est également possible de déterminer V_{ud} à partir de la désintégration du neutron en utilisant la formule suivante [27]:

$$|V_{ud}|^2 = \frac{4908.7(1.9) \text{ s}}{\tau_n (1 + 3\lambda^2)}, \quad (5.5)$$

où $\lambda = g_A/g_V$ est le ratio des constantes de couplage. Déterminer à nouveau V_{ud} en utilisant 5.5 constituerait une confirmation importante de l'unitarité de la matrice CKM. Une valeur indépendante de λ peut être prise, par exemple, d'une mesure de la corrélation angulaire de l'émission d'électron dans la désintégration du neutron : $A_0 = -2\lambda(\lambda + 1) / (1 + 3\lambda^2)$. Afin de devenir compétitive avec la valeur prise des désintégrations super-permises, il faudra améliorer la précision actuelle de $\tau_n = 880.3(1.1) \text{ s}$ d'à peu près un ordre de grandeur.

τ_n dans la cosmologie

La durée de vie du neutron joue également un rôle important dans de nombreux modèles cosmologiques où il va être utilisé pour calculer des taux de réactions. Un bon exemple est la *nucléosynthèse primordiale*, la formation des isotopes les plus légers - jusqu'au ${}^7\text{Li}$ - lors des premières minutes après le Big Bang. Les rapports des abondances de ces isotopes sont intimement liés au rapport du nombre de neutrons et de protons au début de la nucléosynthèse et, par ce biais, à la durée de vie du neutron. La comparaison des abondances - mesurées dans l'univers actuel - aux prédictions faites à partir du modèle de la nucléosynthèse, représente un test important de la théorie du Big Bang.

5.2. Etat actuel de la mesure de τ_n

Une vue d'ensemble des mesures les plus précises de τ_n des trente dernières années est montré dans la figure 1.2. Les mesures peuvent être réparties en deux sortes d'expérience largement différentes: les *mesures sur faisceaux* et les *mesures de piégeage*.

Les mesures sur faisceaux vont mesurer le taux de désintégration absolu d'un faisceau de neutrons ainsi que le nombre absolu de neutrons traversant le volume de désintégration. La technique nécessite deux mesures absolues de haute précision ce qui représente un défi expérimental considérable.

Les expériences de piégeage ne nécessitent aucune mesure absolue. L'idée est de piéger des neutrons pour des laps de temps prolongés, comparable à leur durée de vie. A la fin du temps de stockage le piège va être ouvert et le nombre de neutrons survivants est compté. En variant le temps de stockage il est possible de déterminer la *constante de temps de stockage* du piège en ajustant une loi de décroissance exponentielle aux données expérimentales. La relation entre cette constante de stockage et la durée de vie du neutron est la suivante:

$$\tau_{\text{piege}}^{-1}(E_{\text{tot}}) = \tau_n^{-1} + \sum_i \tau_{\text{perte},i}^{-1}(E_{\text{tot}}) \quad (5.6)$$

Les $\tau_{\text{perte},i}$ sont les *constantes de temps de stockage partielles* liées aux différents modes de pertes lors du piégeage. Le défi principal des expériences de piégeage est la limitation ainsi que la détermination précise des différents $\tau_{\text{perte},i}$.

5. Le neutron et son temps de vie en théorie et expérience

Une deuxième difficulté provient de la “volatilité” des neutrons. Ce ne sont que les neutrons d’une énergie cinétique *ultra* basse qui peuvent être stockés pour des périodes de temps prolongées. On appelle ces neutrons des *neutrons ultra froids*, ce qui est généralement abrégé en UCN (angl.: *ultra cold neutrons*). Dans le cas des pièges matériels, les énergies maximales sont limitées à $E_{kin} \lesssim 300 \text{ neV}$.

Restons avec les pièges matériels; on trouve que - la désintégration mise à part - la probabilité de perte durant le stockage va être dominée par les *pertes par collision avec la paroi* τ_{paroi} . Afin de déterminer τ_{paroi} les expérimentateurs ont recours à la méthode de *l’extrapolation au piège infini*: la constante de temps de stockage du piège va être mesurée à plusieurs reprises en variant la taille du piège. A la fin, une extrapolation à un piège infiniment grand est faite à partir des mesures réalisées. Pour une taille infinie, la fréquence de collision, et donc par voie de conséquence τ_{paroi} , sont tous deux égaux à zéro.

Cependant cette extrapolation n’est pas une tâche triviale et engendre généralement des erreurs systématiques considérables. De ce fait le piégeage d’UCN dans de forts champs magnétiques ($V_{mag} = 60.3 \text{ neV T}^{-1}$) représente une alternative prometteuse puisqu’elle contourne le problème des pertes par collision avec la paroi. Des mesures de τ_n avec une précision finale aux alentours de 100 ms devraient devenir accessibles grâce à cette technique.

6. Conception et aspects systématiques de l'expérience

HOPE - le *H*albach-type *O*ctupole *P*ermanent magnet *U*CN trap - est un piège magnéto-gravitationnel pour UCN dédié à la mesure du temps de vie du neutron libre. Le projet a été proposé par Oliver Zimmer il y a une quinzaine d'années [47]. Le concept général de l'expérience a été élaboré par Kent Leung, qui a également commencé à travailler sur la partie clé du projet, le système d'aimant permanent [24]. J'ai repris son travail en 2011, finalisé la conception de l'expérience et pu finir sur une première prise de données sur la source UCN PF2 en 2014. La première partie de ma thèse est dédiée à une description détaillée du piège et de ses composantes principales. Une vue d'ensemble de l'appareil est montré dans la figure 2.1. Les résultats du temps de faisceau sont présentés dans la deuxième partie.

6.1. Les aimants permanents

Le système des aimants permanents (SAP) constitue le "cœur" de l'expérience car ils apportent la partie principale de la barrière magnétique. Le SAP est un cylindre de 1.2 m de long et un diamètre extérieur d'à peu près 20 cm. Le cylindre est traversé par un trou concentrique d'un rayon intérieur de $\rho_{SAP} = 46.8$ mm. Ce trou constitue la zone de stockage d'UCN. Les régions proches des bords exclues, le champ magnétique absolu $|B|$ peut être approximé par la formule suivant :

$$|B(\rho)| = |B_{\perp}(\rho)| = |B_{max}| \left(\frac{\rho}{\rho_{SAP}} \right)^3, \quad (6.1)$$

où $|B_{\perp}|$ est la part du champ magnétique perpendiculaire à l'axe du cylindre, $|B_{max}|$ est le champ magnétique maximal (1.3 T) et ρ est la position radiale.

Techniquement ce champ magnétique est réalisé par un arrangement d'aimant dit de *type Halbach* [17]. Une illustration de la structure de cet arrangement peut être vue à la gauche de la figure 2.2.

6.2. Les bobines supraconductrices

Les aimants permanents apportent le confinement radial mais aux deux extrémités du système les UCN sont toujours libres d'entrer et ressortir du piège à leur guise. Nous avons décidé d'orienter le SAP verticalement. Ainsi, grâce au champ gravitationnel de la terre ($V_{grav} = 102$ neV m⁻¹), il est impossible aux UCN d'atteindre l'extrémité supérieure

6. Conception et aspects systématiques de l'expérience

du système. L'extrémité inférieure est quant à elle fermée grâce à une bobine supraconductrice. Cette bobine dite "bobine shutter" atteint un champ magnétique d'environ 1.7 T et peut être mise à son courant maximal en 10 secondes. Une rapide mise en route de la bobine est essentielle; le piège étant rempli et vidé par le bas, la bobine doit donc être allumée et éteinte pour chaque cycle de mesure.

La bobine shutter fait partie d'un système de trois bobines coaxiales, placées les unes à coté des autres sur un tube de support commun (figure 2.7). Le diamètre intérieur du tube de support est de 300 mm ce qui permet de placer le système de bobines autour du SAP.

Au cours des mesures faites en 2014 les deux bobines supplémentaires n'ont pas été nécessaires. Elles deviendront plus importantes lors de la détection des particules chargées émises lors de la désintégration des neutrons. Pour ces mesures un détecteur d'électron ou de proton va être installé en haut du SAP. La bobine du milieu créera un champ magnétique vertical de quelques centaines de mT qui servira comme champ de guidage aux particules chargées. La bobine du haut pourrait être utilisée comme lentille à protons qui les focalise sur un détecteur de petite taille.

6.3. Composants intérieurs et méthodes expérimentales

6.3.1. Effets systématiques

Pour la précision que nous envisageons, plusieurs effets systématiques sont non négligeables. Un effet notable est la perte d'UCN par dépolarisation. Le potentiel magnétique qui est vu par les neutrons dépend de l'orientation de leur spin par rapport au champ magnétique externe. Seuls les neutrons dont le spin est orienté à *l'encontre* du champ magnétique peuvent être stockés effectivement. Conséquemment il est primordial d'empêcher les neutrons piégés de perdre leur polarisation au cours du temps de stockage. Des simulations ainsi que des mesures préliminaires faites par Kent Leung suggèrent que notre piège devrait être plutôt insensible à ces dépolarisations [24]. Même sans précaution supplémentaire nous nous attendons à un taux de dépolarisation inférieur à $1 \times 10^{-7} \text{ s}^{-1}$.

Des pertes d'UCN par collision avec des molécules de gaz résiduel doivent également être prises en compte. Une limitation de la pression finale dans la zone de stockage en dessous de $1 \times 10^{-6} \text{ mbar}$ devrait très probablement résoudre ce problème.

L'effet dominant va venir des UCN sur-critiques. Un UCN sur-critique est un UCN ayant une énergie totale¹ au-delà de la barrière énergétique du piège. Avec la bobine du bas à 1.7 T l'énergie stockable maximale est d'environ 60 neV.

Par définition, les UCN sur-critiques ne sont pas proprement confinés dans le piège. Néanmoins ils peuvent atteindre des temps de stockage considérables, comparables à la durée de vie du neutron même. La présence de neutrons sur-critiques dans le piège abaisse le temps de stockage moyen de l'ensemble d'UCN et, de ce fait, corrompt la mesure. Il est quasiment exclu de différencier un neutron sur-critique d'un neutron sous-

¹Energie totale = énergie cinétique + énergie potentielle par rapport au minimum du potentiel de la zone de stockage.

6. Conception et aspects systématiques de l'expérience

critique dans les données finales. Ainsi, il est capital de limiter la part de neutrons sur-critiques dans l'ensemble d'UCN ainsi que de fidèlement prouver cette limitation. La contamination acceptable peut être estimée aux alentours de $N_{surcrit}/N_{souscrit} = 10^{-3}$. Nous avons inclus plusieurs dispositifs à l'intérieur de l'appareil qui devront nous aider à atteindre cet objectif.

6.3.2. L'absorbeur

Au début de chaque cycle un absorbeur en polyéthylène est descendu dans la zone de stockage afin de *nettoyer le spectre d'UCN*, c'est-à-dire enlever les UCN trop énergétiques pour être stockés durablement. Des UCN qui touchent la surface de l'absorbeur vont subir une diffusion inélastique et être "expulsés" du domaine énergétique ultra froid. L'énergie totale minimale nécessaire pour atteindre la surface de l'absorbeur dépend de sa position verticale qui peut être facilement adaptée.

Le positionnement de l'absorbeur est réalisé par un moteur pas-à-pas installé au-dessus de la zone de stockage. La course maximale du moteur est d'environ 30 cm ce qui permet de faire varier la limite énergétique de plus de 30 neV.

6.3.3. Le piston

En plus de l'absorbeur qui entre par le haut, un piston en Téflon peut être introduit dans la zone de stockage par le bas afin d'*accélérer le temps de nettoyage*. Notre potentiel de stockage est hautement symétrique (ce qui peut être vu p. ex. dans figure 2.9) ce qui est susceptible de favoriser des trajectoires d'UCN stables, circulant loin de la surface de l'absorbeur. De surcroît, les neutrons vont avoir tendance à séjourner proche du minimum du potentiel qui se trouve en plein milieu de la zone de stockage, loin de la surface de l'absorbeur.

Le piston est censé contrer ces effets indésirables: grâce à sa surface plutôt irrégulière la plupart des collisions avec le piston vont avoir un caractère diffus. A chaque fois qu'ils heurtent sa surface les UCN vont donc changer de direction avec une grande probabilité et finir par être dirigés vers la surface de l'absorbeur. Comme le potentiel de Fermi du Teflon est supérieur à la barrière énergétique du piège, le piston peut pénétrer profondément dans la zone de stockage sans pour autant mener à des pertes trop élevées d'UCN sous-critiques. En haussant la surface du piston au-dessus du minimum de potentiel du piège, la fréquence de collision des UCN devient maximale et le temps de nettoyage devrait être réduit considérablement.

Comme l'absorbeur, le piston est positionné grâce à un moteur pas-à-pas. Celui-ci a une course maximale d'environ 700 mm et atteint des vitesses de déplacement $\geq 50 \text{ mm s}^{-1}$.

6.3.4. Le switch détecteur

Le "switch détecteur" se situe en dessous de la zone de stockage. Il s'agit d'un passage pour UCN avec deux positions différentes: la "position de détection" et la "position de remplissage". Une esquisse détaillée du switch est montrée dans la figure 2.13. Il s'agit

6. Conception et aspects systématiques de l'expérience

d'une grande chambre en aluminium (517 x 152 x 190 mm) qui contient un corps en Téflon de 365 x 110 x 160 mm. Le corps en Téflon est pourvu de deux passages circulaires, un vertical et un deuxième incliné de 52.5°.

Quand le switch se trouve dans sa position de remplissage, le passage vertical est placé dans l'axe de la zone de stockage et relie ainsi le piston au piège. Dans la position de détection, le passage incliné est placé en dessous de la zone de stockage et les UCN sont libres de "tomber" dans un détecteur UCN qui se trouve à peu près 1 m en dessous du minimum du potentiel. L'inclinaison du passage comble le déport horizontal du détecteur par rapport à l'axe du piège qui est d'environ 25 cm.

Le déplacement du switch se fait par un système pneumatique, fait sur-mesure afin de remplir nos exigences de vide importantes.

7. Le temps de faisceau sur PF2

Le temps de faisceau en 2014, fut la toute première fois que HOPE a été intégralement assemblé et nous avons passé une grande partie du cycle avec sa mise en service. Cependant nous avons été entièrement récompensés pour cet investissement puisque nous avons abouti à un état de fonctionnement automatisé et stable qui nous a permis de poursuivre avec une première prise de données pendant les neuf derniers jours du cycle. Nous avons dédié ce temps à des premiers tests systématiques et à l'optimisation des paramètres de mesure.

7.1. Procédures opératoires et paramètres d'intérêt

Toutes nos mesures étaient des mesures *fill-and-empty* (remplissage-et-vidage):

- Pendant les premières 95 secondes de chaque cycle la bobine du bas était éteinte et les UCN venant de la turbine de PF2 étaient libres d'entrer et sortir du piège.
- Ensuite la bobine était mise en route et le piège fermé.
- Une fois que la bobine à son courant maximal, le piston et l'absorbeur étaient insérés dans la zone de stockage pour enlever les neutrons sur-critiques.
- Après le nettoyage commençait la période de stockage, de durée variable (entre 30 et 2500 secondes).
- A la fin du cycle, la bobine du bas était éteinte et le nombre d'UCN survivants était compté.
- Une dernière phase de 100 secondes a été prévue pour surveiller le bruit de fond.

Comme l'effet systématique dominant est dû aux UCN sur-critiques, nous avons décidé de nous concentrer sur l'investigation et l'optimisation du nettoyage du spectre.

7.2. Premières données prises avec HOPE

7.2.1. Mesures du nettoyage du spectre

Il y a trois paramètres indépendants qui vont influencer le nettoyage du spectre: La position de l'absorbeur, la position du piston et le temps de nettoyage. Lors de nos mesures, nous avons essayé neuf positions différentes de l'absorbeur. Ces neuf positions ont été combinées avec deux positions de piston - complètement retiré ou légèrement en

7. Le temps de faisceau sur PF2

dessous du minimum du potentiel - et deux temps de nettoyage - 120 et 240 secondes - pour une totalité de 24 scénarios de nettoyage différents.

Pour chaque combinaison, nous avons effectué des mesures avec 30 et 480 secondes de stockage. La comparaison des deux permet de faire des estimations sur le nombre de neutrons stockables et sur-critiques en fonction des paramètres de nettoyage. Des informations supplémentaires proviennent du *taux de fuite de neutrons* pendant le temps de stockage: pour les UCN sur-critiques, la barrière magnétique créée par la bobine shutter ne représente pas un obstacle insurmontable. Conséquemment ils peuvent quitter le piège par le bas durant le temps de stockage et trouver leur voie vers le détecteur UCN. Ce taux de fuite résulte en un taux de comptage lentement décroissant durant la phase de stockage (voir figures 3.2 et 3.8). En prenant le nombre total d'UCN étant tombés du piège jusqu'à l'instant t $\Delta N_{fuite}(t)$, la constante de temps de fuite des UCN sur-critiques τ_{fuite} et la constante de temps de stockage des sur-critiques τ_{crit} ¹ nous pouvons calculer le nombre d'UCN sur-critiques initial $N_{crit,t=0}$ en utilisant la formule suivante:

$$\begin{aligned} \Delta N_{fuite}(t) &= \int_{t'=0}^{t'=t} N_{crit,t=0} \exp^{-t'/\tau_{crit}} \frac{1/\tau_{fuite}}{1/\tau_{crit}} dt' \longrightarrow \\ &\longrightarrow N_{crit,t=0} = \Delta N_{fuite}(t) \frac{1}{1 - \exp^{-t/\tau_{crit}}} \frac{\tau_{fuite}}{\tau_{crit}}. \end{aligned} \quad (7.1)$$

Les figures 3.7 et 3.9 montrent les résultats des mesures de nettoyage. On peut voir que le nettoyage n'est pas encore très effectif. En choisissant des positions d'absorbeur suffisamment basses, le nombre initial d'UCN sur-critiques est considérablement réduit mais le nombre d'UCN stockable s'en voit tout autant affecté. Quand le piston est monté dans la zone de stockage la réduction des deux, UCN stockable et sur-critiques, devient encore plus prononcée. Manifestement, le piston ne change pas seulement la direction des UCN mais les absorbe également avec une probabilité non-négligeable. Si cela est dû à *l'échauffement Doppler* durant le mouvement du piston, des contaminations sur sa surface ou un autre effet jusqu'alors ignoré devra être exploré lors d'un prochain temps de faisceau.

7.2.2. Mesures avec des temps de stockage longs

Afin d'obtenir une première estimation du temps de stockage des UCN sous-critiques, nous avons également effectué quelques mesures avec des temps de stockage plus longs, allant jusqu'à 2500 secondes. Pour ces mesures nous avons utilisé trois combinaisons différentes de paramètres de nettoyage. Les résultats des trois ensembles de données sont montrés dans la figure 3.10. Selon le mode de nettoyage choisi nous avons stocké entre 50 et 100 UCN sous-critiques par cycle. Malgré les statistiques médiocres, les mesures ont été satisfaisantes: En ne prenant que les mesures "hautement nettoyées" - avec un

¹ $\tau_{crit}^{-1} = \tau_n^{-1} + \tau_{fuite}^{-1}$

7. Le temps de faisceau sur PF2

taux de neutrons sur-critiques d'environ 15 % - nous arrivons à une constante de temps de stockage de $\tau_{\text{piege}} = 881 \pm 45$ s. Dans sa marge d'erreur, ce résultat est en bon accord avec la moyenne mondiale reconnue de la durée de vie du neutron $\tau_{PDG} = 880.3 \pm 1.1$ s [30].

Outre ce résultat encourageant les mesures ont également révélé quelques contradictions. Entre autres, la constante de temps de stockage déduite est bien plus longue que l'on aurait supposé avec un taux d'UCN sur-critique aussi élevé. La raison des contradictions apparentes réside dans le modèle trop simplifié utilisé pour décrire la portion d'UCN sur-critiques. Dans mon analyse, je n'ai utilisé qu'une seule constante pour décrire la totalité des UCN sur-critiques. En réalité la constante de temps de stockage de chaque UCN dépend de son énergie cinétique et peut donc varier fortement de neutron à neutron.

Puisque la détermination du nombre d'UCN sur-critiques dépend de τ_{crit} , il est nécessaire de reprendre la méthode d'analyse. Pour ceci nous aurions également besoin de données supplémentaires sur le spectre d'UCN piégés et de son évolution. Eventuellement ces informations pourront être extraites de la *courbe de vidage*. Quand le piège est vidé à la fin du cycle ceci ne se fait pas instantanément mais avec un certain délai dans le temps dû à la diminution continue du courant dans la bobine shutter. Par conséquent un neutron de haute énergie pourra quitter le piège plus tôt qu'un neutron de basse énergie. Dans figure 3.11 on peut voir comment cette causalité influence la forme de la courbe de vidage et deviner l'étendue des informations sur le spectre d'UCN contenues dans les données.

7.3. Perspectives

7.3.1. Aspects statistiques

Les résultats présentés dans cette thèse sont fortement limités par la statistique très basse. Après environ 24 heures de prise de données l'erreur totale sur la constante de stockage est de 45 s. En supposant que cette erreur soit de nature purement statistique nous pouvons calculer le *gain en neutron* minimal qu'il nous faudra afin d'arriver à une certaine précision. Si nous définissons comme but minimal une précision de ± 1 s après un cycle entier - soit 50 jours - nous trouvons qu'il va falloir augmenter le nombre d'UCN piégés d'au moins un facteur 40. Même sans augmenter le nombre d'UCN nous pouvons améliorer la précision statistique quelque peu, en optimisant les paramètres des mesures (les temps de stockage choisis, le nombre de points de mesure, etc.). Ainsi nous pouvons réduire le gain nécessaire à un facteur 30 à 35.

Nous nous attendons à une hausse majeure une fois que les bobines auront été réparées. Durant le temps de faisceau nous ne pouvons pas les utiliser à leur courant maximal (300 A) à cause d'un dysfonctionnement de leurs *switchs supraconducteurs*. Au-delà de 200 A les switchs restent continuellement au-dessus de leur température de transition et il n'est donc pas possible de mettre les bobines en mode persistant. Des calculs numériques ont indiqué que, selon le spectre d'UCN piégés, l'augmentation du champ de la bobine du bas de 1.1 à 1.7 T pourrait décupler le nombre de neutrons stockés. Nous avons aussi fait quelques mesures à 300 A - avec la bobine du bas en mode

7. Le temps de faisceau sur PF2

“ouvert” - où nous avons trouvé une augmentation d’à peu près 300 %. Il est probable que le désaccord entre les calculs et les mesures soit dû à une forte déformation du spectre d’UCN lors des mesures sur PF2 causé par la grosse marche gravitationnelle entre la turbine UCN et le piège.

Quoi qu’il en soit, même dans le plus optimiste des cas, il semble pratiquement exclu que nous atteindrons une statistique suffisante sur PF2. Effectivement nous n’allons pas retourner à PF2 mais installer HOPE sur la source SUN2. SUN2 fait partie d’une nouvelle génération de source UCN, basée sur la création des UCN par diffusion inélastique de neutrons froids dans de l’Hélium superfluide. Les UCN sont accumulés dans la source jusqu’à atteindre une densité d’équilibre. Puis la source est ouverte et les UCN sont libres de se propager dans le piège. SUN2 n’a pas encore servi pour remplir une expérience mais les derniers tests ont été très prometteurs, délivrant plus de 6×10^5 UCN par cycle d’accumulation.

7.3.2. Aspects systématiques – comment continuer avec les UCN sur-critiques

Nous avons vu jusqu’alors que notre mode de nettoyage n’est pas encore très efficace mais nous avons plusieurs idées pour résoudre ce problème.

Premièrement nous allons couvrir la surface du piston avec une couche de Fomblin afin de couvrir des contaminations éventuelles. Afin d’étudier la possibilité d’échauffement Doppler nous allons varier la vitesse de translation du piston. Si nous trouvons que l’effet Doppler pose un réel problème nous envisageons d’échanger le piston par une plaque verticale. Comme la direction de mouvement d’une plaque verticale serait *parallèle à sa surface* - et non plus perpendiculaire comme pour le piston - l’échauffement Doppler devrait être réduite considérablement.

Quant à la détermination du rapport $N_{surcrit}/N_{souscrit}$, le taux de fuite de neutron durant le stockage s’est avéré être une bonne source d’informations. Sa précision finale dépend du nombre total d’UCN stockés. Il va donc falloir attendre les premières mesures sur SUN2 pour voir si elles nous permettront de déterminer le rapport avec la précision nécessaire ($N_{surcrit}/N_{souscrit} < 10^{-3}$).

8. Conclusion

Dans le passé, les deux techniques dominantes pour la mesure du temps de vie du neutron ont été les mesures sur faisceau et les pièges à UCN matériels. Tandis que les précisions respectives des deux techniques ont été continuellement améliorées au cours des dernières décennies, les résultats qu'elles délivrent sont inconciliables. Le piégeage magnétique permet de mesurer le temps de vie d'une façon systématiquement indépendante et pourrait bien jouer un rôle décisif. Les sources d'erreurs systématiques dominantes des pièges magnétiques sont les pertes d'UCN dues à leur dépolarisation et aux collisions avec les molécules du gaz résiduel ainsi que les UCN sur-critiques. Dans cette thèse j'ai présenté le piège magnéto-gravitationnel HOPE, montré quelles mesures ont été prises pour gérer les effets systématiques prévus et discuté les premières données expérimentales prises lors de la mise en service de l'appareil sur la source d'UCN PF2.

Du point de vue technique le temps de faisceau a été un succès sur toute la ligne. Pour la toute première fois nous avons entièrement assemblé l'appareil, mis en service toutes les parties mécaniques et sommes arrivés à mettre l'ensemble dans un mode de fonctionnement stable contrôlé à distance. Les quelques problèmes persistants concernent surtout la communication numérique. Une fois qu'ils seront résolus l'expérience pourra être maniée par une seule personne puisque la seule interaction nécessaire sera le remplissage régulier des fluides cryogéniques ainsi que le changement occasionnel des paramètres de mesure.

Du point de vue expérimental il y a encore des améliorations à faire. Par le passé, notre handicap majeur était le taux de comptage très faible. En fonction du mode de nettoyage choisi nous nous retrouvons avec 20 à 100 UCN stockables par remplissage et une précision statistique $\delta\tau_{\text{piège}} \gtrsim \pm 45\text{ s}$ après 24 heures de mesure. Une des raisons principales pour ces taux faibles est la défectuosité des switches des bobines supraconductrices. L'échange des switches combiné à des nombreuses petites adaptations devraient augmenter le taux d'UCN initial d'un facteur 5 à 10, peut-être même un peu plus dans un scénario optimiste. La hausse décisive devrait suivre une fois que HOPE sera installé sur la source d'UCN SUN2. Dernièrement, SUN2 a délivré plus de 6×10^5 UCN par cycle d'accumulation. Si nous ne transférons que 1 % des UCN disponibles nous pourrions facilement atteindre notre but de précision provisionnelle, une erreur statistique de ± 1 s en 50 jours de mesure.

Nous avons pu obtenir une première estimation de la constante de temps de stockage du piège $\tau_{\text{piège}} = 881 \pm 45\text{ s}$ qui est compatible avec la moyenne mondiale reconnue de la durée de vie du neutron $\tau_{PDG} = 880.3 \pm 1.1\text{ s}$. Le taux de fuite associé est de $\tau_{\text{pertes}}^{-1} = -0.4 \pm 4.4 \times 10^{-5}\text{ s}^{-1}$. A cause de l'amplitude de l'erreur statistique, il n'a pas été possible de décomposer τ_{pertes} dans les différentes sources de perte d'UCN connues. En substitut, j'ai donné des estimations des décalages de la constante de temps de stockage

8. Conclusion

dû à la dépolarisation $\Delta\tau_{sys,depol}$ et aux collisions avec le gaz résiduel $\Delta\tau_{sys,gaz}(p)$. La structure de notre champ magnétique mène à un taux de dépolarisation négligeable qui résulte dans un décalage $\Delta\tau_{sys,depol} < 0.1$ s. Les pertes d'UCN dues au gaz résiduel sont moins bien définies du fait de notre méconnaissance de la pression dans la zone de stockage. Nous visons à réduire la pression finale davantage, bien en-dessous de 10^{-6} mbar, afin de garantir des taux de pertes acceptables.

La plus grosse partie des mesures a été consacrée à l'investigation du nettoyage de spectre et nous avons pu réduire le nombre d'UCN sur-critiques considérablement. Malheureusement, une haute efficacité de nettoyage a toujours été liée à des pertes fortes d'UCN stockables. Lors de la prochaine prise de données nous devons équilibrer la balance entre l'efficacité du nettoyage et les pertes d'UCN totales.

En résumé, les résultats de nos premières mesures ont été plutôt prometteurs. Il va falloir confirmer expérimentalement les estimations des décalages $\Delta\tau_{sys,depol}$ et $\Delta\tau_{sys,gaz}(p)$. Il semblerait que nous allons être capables de les limiter toutes deux à quelques 100 ms, voire moins. En ce qui concerne les UCN sur-critiques, il va falloir redéfinir nos modes de nettoyage et de détermination du spectre. Le défi principal sera de déterminer le rapport $N_{surcrit}/N_{souscrit}$ avec la précision nécessaire. Une possibilité sera d'investiguer la courbe de vidage dépendante de l'énergie des neutrons afin d'en tirer des conclusions sur le spectre des UCN. En 2014, les informations de la courbe de vidage ont été limitées par les faibles taux de comptage. Toutefois, elle pourrait devenir un outil performant une fois combiné au taux de neutrons élevé que nous attendons trouver sur SUN2.

A. Cryostat

While all of HOPE's inner installations are in-house developments, the cryostat was designed and manufactured by the British enterprise *AS Scientific*. A 3D-model of the entire apparatus is shown in figure A.1.

The cryostat's centerpiece is the LHe tank. Its effective volume - the total volume *minus* the volume of the superconducting coils - is of about 200 liter, 120 liter to cover the coils and another 80 liter in a reservoir above them. In order to minimise the helium consumption the tank is equipped with a recondensation stage: a massive copper block with a series of fins on its lower end that maximise the cooling surface (see detail 1 of figure A.1). The block is connected to the second stage of a *Cryomech PT415* pulse tube with a cooling capacity of 1.5 W at 4.2 K [6].

A second *Cryomech* cryocooler, an *AL200* cold head is used to cool the coils' current leads. The current leads are a delicate part as they easily turn into a major heat leak. In our case the problem is addressed by six *CryoSaver* current leads from the New Zealand company *HTS-110 Ltd*. The leads consist of the current leading filaments that are made from some high-temperature superconducting (HTS) material and a fiber glass matrix that provides the mechanical stability. The combined heat leak of all six of them has been specified to be around 600 mW.

In order to reduce the radiative heat load on the helium tank it is surrounded by an aluminum radiation shield. The shield is cooled by the first stage of the pulse tube and also loosely connected - by six copper braids - to the cold head and wrapped into several tens of layers of superisolation foil. Both, the shield and the foil, not only cover the bath from the outside but also from the *inside* in order to thermally decouple it from the bore tube. Thus, the cryostat can be run with the bore tube being at room temperature without a considerable increase of the radiative heat load.

Apart from the current leads the cryostat works quite well. During a test run without current we measured a LHe boil-off rate $< 0.2 \text{ l h}^{-1}$. This corresponds to a net heat load - after subtraction of the cooling power of the recondensation stage - below 110 mW. However the helium consumption increases drastically when the coils are used: during the 2014 beamtime we had to refill about every two days which corresponds to an average LHe consumption of roughly 2 l h^{-1} .

An increase in the helium consumption during the ramping process is unavoidable. Nonetheless, the boil-off rate we encountered is higher than expected. Figure A.3 shows the systems first and second iteration of the current lead system. The light green bars are the HTS current leads. On their top and bottom end the HTS leads are thermalised via two copper disc that are thermally anchored to the *AL200* cold head and the helium bath respectively. The top plate is directly screwed to the cold head while the bottom plate is connected to the bath via some copper braids and a copper bar that dips into

A. Cryostat




Figure A.1.: *SolidWorks* sketch of the cryostat. Detail 1 shows the first and second stage of the pulse tube and the copper block that serves as recondensation stage. Detail 2 shows some technical particularities. Feature A indicates a short bellows on the bottom end of the bore tube. The bore tube is not actively heated and once the permanent magnets are installed and cooled down, the bore tube's temperature will evenly drop to some floating point between the PMS' temperature and RT. The bellows allows for the thermal contractions that will occur. Feature B is the "support structure" of the superconducting coils: a ring that is tack-welded to the helium tank. The foot C has no supportive function but centers the helium tank in radial direction. It is made from fiber glass in order to minimise the conductive heat load. Detail 3, the current leads, is shown in the separated figure A.4.

A. Cryostat




Figure A.2.: The cryostat's top flange. In the front we see the pulse tube and in the back the cold head. On the left, equipped with a red handle bar, is the pumping port for the helium bath. The liquid helium entrance is hidden behind the pulse tube but one can see a LHe transfer tube arriving from the right hand side. Three 12-pin *Jaeger* plugs are situated below the transfer line. They allow for the installation of eight temperature sensors (read out by 4-wire sensing) and one heater on the bottom of the helium bath. Later we added two T-pieces in order to allow for up to six more probes.

A. Cryostat




Figure A.3.: First and second version of the current lead assembly. The left picture shows the entire assembly from the arrival of the primary copper braids on top to the feedthroughs into the helium tank in the bottom. The right picture shows an attempt by *AS* to improve the bottom part of the system. From the HTS leads' bottom end to the feedthroughs the current is led through LTS wires that are doubled with some copper braidings. The LTS leads are wrapped in spirals around the copper braids in order to add some flexibility. Unfortunately, the wires were bend way past their minimum radius, partially they were even visibly broken.

the liquid helium. The detail in the middle of figure A.3 shows how the HTS leads are linked to the copper discs: it is a sandwich structure, pressed together by two M4 screws. The leads are electrically isolated from the rest of the system by some fiber glass jackets that surround the screws and annealed aluminum plates that are placed in between the copper discs and the HTS leads.

Unfortunately the system proved to be very error prone. The screws, used to connect the HTS leads to the copper block, regularly became loose resulting in a bad thermal contact and HTS temperatures around 90 K, about 30 K above their transition temperature. Furthermore we realised that the lower section of the assembly was not working properly either. From the lower copper disc on, the current is led through low-temperature superconducting (LTS) wire and the disc is supposed to constantly stay beyond the wires' critical temperature (~ 9 K). However its temperature barely dropped below 10 K and quickly rose to 20 K and more once the current was switched on. As it was virtually excluded to perform proper measurements under the given conditions we decided to radically revise the entire system.

A. Cryostat

Figure A.4 shows a 3D sketch of the final design. Detail 1 shows the thermal and electrical connection of the HTS leads upper end. As before, we find a compressed sandwich structure but the M4 screws have been exchanged by M10 stud screws. The screws press against a thin copper plate followed by a sapphire plate, the connector block of the primary copper braid, the HTS lead and finally another sapphire plate. The sapphire plates replace the annealed aluminum plates and the thin copper plate serves as protective layer for the first sapphire plate. The stud screws are no standard pieces either but custom made screws from BeCu. Compared to classical steel screws they have a much higher thermal conductivity and should thus help to further improve the thermal anchorage. The copper disc has been exchanged by a disc made from T6 aluminum, a material that combines favorable thermal characteristics with a high mechanical hardness. It should prevent the M10 thread holes from becoming worn out as the M4 threads have been.

On the HTS leads' bottom end the system is mostly identical as can be seen in detail 2. However we could not completely exchange the lower copper disc like we have done with the upper one as it is hard soldered to the copper braids that link it to the helium bath. In the new system the copper disc is surrounded by a T6 ring that holds the BeCu screws. The connectors between the HTS leads' bottom end and the feedthroughs have been replaced by some straighter ones.

In order to reduce the thermal and electrical contact resistance at the various intersections we payed close attention to the surface condition. Wherever possible we gold-plated the surfaces in order to prevent the formation of a resistive corrosion layer. The already polluted surfaces that could not be exchanged - e.g. the feedthroughs, etc. - were thoroughly cleaned and polished as good as possible. Finally, I conducted the entire assembly wearing latex gloves in order to circumvent last minute contamination.

The revised current lead system is running stable even though we could not remedy all problems. Figure A.5 shows temperature curves from the 2014 beamtime. The top plot shows the cryostat's cooldown in the beginning of the beamtime. The timeline starts shortly before the cryocoolers have been switched on, after the LN2 precooling phase. The two steps in the temperature curves occurred when the coolers were switched on and when the LHe was filled into the tank.

The bottom plot shows the temperature behaviour during the shutter coil ramping. The chosen data is special as it comes from the 300 A runs. As at 300 A it was not possible to go into persistent mode due to the dysfunctional superconducting switches I had to leave the current running during the entire storage phase. The short bumps correspond to a 30 seconds, the longer ones to a 480 seconds storage time. Without current, the whole lower section of the current lead assembly stays < 6 K and the LTS leads should thus be superconducting. However once the current is switched on the system transits - almost instantaneously - to a new equilibrium temperature above the leads transition temperature. This indicates that the resistive heating that occurs at the numerous normal conducting junctions is still too high. It seems that the only way to avoid the undesirable temperature rise would be to completely eliminate any normal conducting passages, from the bottom of the HTS leads on until the terminals on the superconducting coils.




Figure A.4.: 3D model of the new current lead assembly.

A. Cryostat

Contrarily to the lower section, the upper section of the current lead assembly does not quickly equilibrate but is continuously rising as long as the electrical current is flowing. While the temperature rise gets slower over time it is not sure if the system would reach its equilibrium state before the HTS leads transit. In any case, once the superconducting switches have been repaired the problem should vanish.

A. Cryostat




Figure A.5.: Temperature curves of the cryostat during the cooldown and some ramping cycles. The legend enumerates the shown temperature sensors in order of their maximum temperature on the bottom plot. *Shutter Coil* and *Bias Coil* are two PT100 probes that are installed next to the top end of the respective HTS lead. *LTS Lead* is a Cernox probe that sits at the bottom end of one of the LTS leads that connects the HTS leads to the feedthroughs. The *4K Stage* Cernox sensor sits on the T& aluminum ring. *Feedthrough* takes the temperature on the flange shown in pink in detail 2 of figure A.4.

B. UCN source SUN2

The reliable conditions on PF2 have been an important factor in the success of the 2014 beamtime since they allowed me to fully focus on the assembly and commissioning of the trap. In the long run the UCN density on PF2 will most probably not be sufficient to reach our ultimate accuracy goal (see section 3.3.2) and we will therefore return HOPE to its originally designated source: SUN2. By today it seems that its potential is still not fully maxed out but the last test runs have already been very promising. In the following I will give a brief summary of SUN2 and its characteristics before presenting some of the most recent measurements. An extensive discussion including the theoretic background, technical details and first measurement results can be found in [24].

SUN2 is based on the accumulation of UCN in a superfluid helium bath that is traversed by a *cold neutron* (CN) beam. By exciting phonons in the liquid some of the CN will lose their energy and eventually end up in the UCN energy regime. The technique notably profits from a resonance in the neutron-phonon-scattering amplitude that occurs at a neutron wavelength of 8.9 Å. At that point the dispersion relations of the phonons and the neutrons cross and the incoming CN can transfer their entire kinetic energy in a single-phonon interaction. The UCN production rate in dependency of the UCN's kinetic energy $R(E_{kin})$ is given by *Fermi's Golden rule*:

$$R(E_{kin}) = \frac{2\pi}{\hbar} |M|^2 \rho_{PS}(E_{kin}) \propto \sqrt{E_{kin}}. \quad (\text{B.1})$$

The energy dependence of B.1 is entirely determined by the energy dependence of the *phase space density of the final states* $\rho_{PS} \propto \sqrt{E_{kin}}$. The energy dependence of the *transition matrix element* M can be neglected since its variations are small over the UCN energy range.

The maximum UCN density in the source is given by the product of the production rate and the storage time $\tau_{source}(E_{kin})$:

$$\rho_{source} = c \times \int_E \tau_{source}(E_{kin}) \sqrt{E_{kin}} dE_{kin}. \quad (\text{B.2})$$

τ_{source} inherits its explicit energy dependency from the wall losses that will occur during the accumulation phase. Those losses limit the integral's boundaries to $E \leq E_{wall} - E_{He4}$, since $\tau_{source}(E > E_{wall} - E_{He4}) \approx 0$, where E_{wall} and E_{He4} designate the Fermi potential of the wall material and of the helium bath respectively. Also they will modulate the final UCN spectrum as the *wall loss probability* $P_{wall}(E_{kin})$ generally increases for higher kinetic energies.

Recent results from SUN2 can be seen in the figures B.1 and B.2. Figure B.1 shows the *build up* of the UCN density ρ_{source} and the source's *emptying* peak for a converter volume

B. UCN source SUN2

made from polished aluminum plates with a Beryllium coating and a layer of Fomblin grease on top ($V_{Fomblin} - V_{He4} \approx 70$ neV). From the build-up curve we can extract a build-up time $\tau_{buildup} = \tau_{source} \approx 230$ s. The integral over the emptying peak gives a total UCN number $N_{source} = 642000$ and thus a UCN density of $\rho_{source} = \frac{N_{source}}{V_{source}} \stackrel{V=4l}{\approx} 160 \text{ cm}^{-3}$.

Figure B.2 shows two time-of-flight spectra for the same converter configuration. The more interesting plot is the lower one that shows the UCN spectrum *after* a 200 seconds accumulation phase. Compared to the upper plot the spectrum has been significantly shifter towards lower UCN energies. It peaks at $E_{kin} \approx 81$ neV.

In combination both plots reveal the potential of SUN2 as a UCN source for HOPE. During the filling of the trap the total UCN density will be diluted and the total UCN number will decrease due to losses in the trap and the guide system. However we need only about 1 % of N_{source} in order to reach a significant statistical sensitivity (see section 3.3.2). This seems to be an accessible goal. The results from the TOF-measurements support this conclusion as they show that the majority of N_{source} are low energy UCN and should therefore be storable in HOPE.




Figure B.1.: *Build up* and *emptying* phase of SUN2 at 0.66 K. During the build up the extraction pipe of the converter volume is not completely closed but allows for a small UCN leakage rate. The time evolution of the leakage rate reflects the evolution of the UCN density in the source and can be used to determine the build-up time $\tau_{buildup}$. At $t \approx 1500$ s the extraction pipe is completely opened and the source is emptied into an UCN detector. The integral over the emptying peak gives the total UCN number in the source N_{source} and thereby the maximum density ρ_{source} .

B. UCN source SUN2




Figure B.2.: Time-of-flight spectra of the SUN2 source. The upper plot shows the spectrum for a permanently opened converter volume while the lower plot shows the spectrum after a 200 seconds accumulation phase. The difference of both spectra demonstrates the effect of the spectrum modulation during the accumulation phase: The shape of the the initial spectrum is mainly given by the production rate $R(E_{kin})$ from equation B.1 and has its maximum density at $E_{kin} \approx 144 \text{ neV}$. The wall losses during the accumulation phase considerably decrease the ratio of high energetic UCN in the spectrum and shift the peak density to $E_{kin} \approx 81 \text{ neV}$.

C. Technical details on the 2014 beamtime on the EDM beamline of PF2

In order to have enough space for the installation of the piston mechanics and the detector switch HOPE has to be mounted on a support structure with about 2 m clearance below the cryostat. Fortunately the EDM platform has a second floor that is perfectly adapted to our needs. It is situated about 2.5 m above the ground floor and about 1 m above the level of the actual beamline and is equipped with a rectangular hole that allowed us to access the cryostat's bottom plate. The plug mechanics, detector shutter, the detector itself, all necessary installations for the cooling of the PMS as well as the UCN guiding system were accessible from the lower, main platform. Most of the "heavy machinery" - the compressors of the two cryocoolers, the power supplies for the superconducting coils, the LHe dewars - were placed on the upper platform.

A bigger drawback of this position was the total length of the guiding system. It was ~ 4 to 5 m long and consisted of nine straight stainless steel guiding pieces, two 25°-bends in order to compensate for the height difference of about 70 cm between the trap's UCN entrance and the EDM main shutter, two reducer pieces (one from 80 to 70 and a second one from 70 to 53 mm outer diameter), one aluminum foil and one UCN shutter. The high number of individual pieces increases the UCN loss probability during transport due to inevitable gaps between the tubes. Another problem is that we had to use a total of 24 vacuum joints throughout the guide system. This not only results in a high initial effort to seal off the system but also leads to an ultimately higher baseline of the final vacuum in the trapping region.

We pumped on the system simultaneously with two turbo pumps, one *Pfeiffer HiPace 80 DN 63* sitting right on top of the cryostat, next to the absorber mechanics, and a similar model from *Alcatel* on the bottom. While the *Pfeiffer* pump was directly flanged to the vacuum chamber, the *Alcatel* pump was connected to the main vacuum chamber at five different points via multiple bellows of different diameters between 16 and 40 mm (see figure C.2). The long and narrow pumping lines limited our pumping capacity but were necessary as the lower part of the vacuum system comprehends several dead volumes that had to be bypassed. We monitored the vacuum with two *Pfeiffer* full range gauges. The one on the bottom was not directly linked to the vacuum chamber but to the bypass system of the turbo pump. The gauge on top was situated on a side port right next to the upper turbo pump. Throughout our measurements the two gauges showed a pressure gradient of about an order of magnitude. The upper gauge usually was in a low 10^{-6} or even a high 10^{-7} mbar region while the gauge at the bottom showed values around 1×10^{-5}

C. Technical details on the 2014 beamtime on the EDM beamline of PF2

mbar. It is not possible to directly conclude on the vacuum in the trapping region as the gauge on the bottom was too far away while the gauge on top was much too close to the turbo pump. Probably the actual pressure was somewhat higher than the pressure measured with the upper gauge, somewhere in between 10^{-6} and 10^{-5} mbar.

A quick and easy way to simplify the main vacuum chamber would be to introduce a supplementary aluminum foil. The first foil is placed directly after the main beamline shutter and is a safety barrier between the experimental vacuum and the UCN turbine. A second foil placed right before the first bend would separate most of the straight guide pieces as well as the UCN shutter from the main vacuum. A quick access to the shutter, i.e. without having to break the main vacuum, would have been of great relief since, about two weeks before the end of the beamtime, it suddenly stopped closing. While we measured a background rate of about 30 mHz when the shutter was perfectly closed, it rose to several hundreds of Hz when the shutter remained open and the signal-to-noise ratio would fell to < 0.1 . Most probably the cause of the malfunction was a misalignment of the shutter's pneumatic drive and its rotating body. When we first tried to solve the problem we could unblock the shutter, at least for a while but ended up with a new problem as now the axis of the shutter's rotator started slowly slipping through the coupling. Consequently both, the opening and closing point of the shutter migrated over time resulting in a steady variation of the filling and the background rate.

Finally the shutter could be repaired in-situ and without breaking the vacuum. However it emphasised a major experimental weak point of the system. Once the PMS was cooled down the main vacuum chamber had to remain evacuated in order to prevent ice formation on the bore tube's inner walls. We had deliberately accepted that disadvantage when we decided not to thermally isolate the tube from the PMS. Unfortunately I had underestimated the thermal coupling of the permanent magnets and the cryostat. During the cooling tests presented in section 2.1.1, the PMS warmed-up automatically from 140 to 290 K within ~ 2 days due to radiative heating. Inside the cryostat the effect was reversed as the PMS was cooled down by the cryostat's inner walls. By applying a constant stream of gaseous nitrogen passing through the copper spirals we could counter-balance the effect and stabilise the PMS temperature above 0° C. The procedure worked passably well even though we did not manage to hold the magnets above the dew-point and I had to manually wipe out the water from the inner walls before we could close the system for pumping. It was clear that once we had started the LN2 cooling we would not be able to reopen the main vacuum chamber before the end of the beamtime as heating up the PMS again would have taken too long. Prior to future beamtimes some electrical heater on the PMS should be installed. At a thermal capacity of $\sim 550 \text{ J kg}^{-1} \text{ K}^{-1}$ [43] and a total mass of 200 kg we should be able to heat up the PMS within 1 1/2 to 2 days with a 200 W heater.

Besides setting up the hardware we also had to implement a measurement program for automated and synchronised operation. The *LabView* program I wrote allows for both, manual and automated operation of the entire set-up (see figure C.3). In addition it displays the current PMS temperature and the values from our magnetic field monitor¹

¹We placed a 3-axis Hall-probe from *LakeShore* next to the cryostat in order to monitor the development




Figure C.1.: HOPE on the EDM platform. The cryostat sits on the first floor while the plug mechanics, the detector switch and the detector itself are accessible from the ground floor. The beam line enters the picture from the right and goes through a pneumatic UCN shutter before making a 25° upward bend towards the UCN entrance of the cryostat. The light blue can sitting next to the cryostat is one of the innumerable helium dewars we emptied throughout the beam time. The two gray “boxes” in front of the cryostat are the compressor units of the pulse tube and the cold head.



Figure C.2.: Close-up of the mechanics below the cryostat. The big ice block in the upper left corner has formed around the entrance of the LN₂ spirals that are used to cool the permanent magnets. The big aluminum box to its left is the casing of the detector switch and below the switch one can see the mechanics of the Teflon piston. The whole system is connected to a turbo pump sitting on the aluminum table at the lower right corner of the picture and equipped with multiple bypasses in order to avoid dead volumes. In the background we see the blue UCN turbine and the main vacuum shutter of the EDM-beam line. The total guide length from the shutter to the cryostat is about 4 to 5 m.

C. Technical details on the 2014 beamtime on the EDM beamline of PF2

and stores the information in two synchronised logfiles. A third logfile registers any kind of interaction between the program and the peripheral devices. The communication with the UCN turbine as well as the main and the monitor UCN detector was done using some standard PF2 program that was running on a separate computer. When switched to automatic measurement mode our program was synchronised with the PF2 software via TTL signals emitted by the latter one. Communication the other way round, sending signals from our end to the PF2 computer, is not provided. On the coding level the focus was to be as generic as possible in order to allow for fast adaption to new ideas. The program's main part is an infinite loop that contains all accessible actions, such as "move-absorber", "ramp-coil" or "switch-on-heater". Each functionality is linked to an on/off-flag that can be controlled either manually or automatically to start/stop the desired action. The synchronisation with the PF2 computer is implemented as a case-structure that will execute a given case whenever one or more of the TTL signals change. In order to change the measurement scheme one only has to change the arrangement of the on/off-flags within the case-structure. Additionally, numerous safety loops are included in order to avoid abusive use such as powering the bottom-heater when the bottom coil is in persistent mode and the power supply is switched off.

Once the common "teething troubles" were resolved the program was working fine with one, unfortunately quite substantial, exception: we could not establish a stable communication between the program and the *Cryogenics* power supply that was designated to run the bottom coil. A possible explanation for the misbehaviour could be that the signal strength was not strong enough for the relatively long USB line (5 - 6 m).

I resolved the problem by exchanging the *Cryogenics* power supply by our second one, the *Heinzinger TNSU 8-300*, a donation from the PERKEO2 collaboration. The *Heinzinger* supply has a maximum current rating of 300 A and a maximum ramping voltage of 10 V. Compared to the *Cryogenics* power supply (375 A @ 40 V) its power rating is rather low and we won't use it in the long run as it would drastically limit our maximal ramping speed. However, due to the dysfunctional superconducting switches, we did not intend to go to high ramping speeds during the 2014 beamtime. The communication with the *Heinzinger* power supply is done via a 0 to 10 V analog input that regulates the 0 to 300 A current output. The ramping voltage can be manually adjusted. Figure C.4 shows a typical ramping cycle of the bottom coil as it was seen by our magnetic field monitor. At the low ramping speeds we used ($\sim 6 \text{ A s}^{-1}$) both power supplies perform almost identically. However it seems that the *Heinzinger* unit settles at a slightly lower final current. For my analysis I only considered data that has been taken once the two supplies had been exchanged. Thus there is no risk of any systematic effects due to the deviance between the two power supplies.

of the magnetic field. The measured values strongly depended on the probe's position but as long as we did not move it the relation between the applied current and the measured field remained stable.

C. Technical details on the 2014 beamtime on the EDM beamline of PF2




Figure C.3.: The *LabView* program allowed online communication with the piston- and the absorber-stepper motors, the pneumatic drive of the detector switch, the heaters of the bottom and the middle coil and the bottom coil's power supply. The left-hand side of the program contains all relevant setting options (e.g. for the movement speed of the steppers or the ramping speed of the power supply) and the manual controls. On the right-hand side one can set the parameters for the automated measurement mode. The two graphs on the right-hand side showed the development of the PMS' temperature and the Hall probes we had installed next to the cryostat for monitoring.




Figure C.4.: Magnetic field monitoring curve for a typical ramping cycle from 0 to 200 A. As the Hall probe only measures the stray fields outside the cryostat it delivers no absolute measure of the B -field but still gives a reliable feedback on the coils' current status. The two lines show the slightly different behaviour when using either of our two power supplies.

D. Monte-Carlo Simulations

The aim of my Monte-Carlo simulations was to get an estimate for the statistical accuracy $\delta\tau_{stat}$ of our measurements in dependency of the initial filling rate N_0 and the measurement parameters (e.g. the total length of the beamtime). Additionally I included a second UCN ensemble with a shorter storage time and variable ratio that is meant to mimic the presence of an overcritical UCN contamination. The work-flow of the code is the following:

1. Define the measurement parameters. The relevant parameters are the total measurement time T_{tot} , the dead time in between to storage phases T_{dead} , the total number of different storage times used N , the minimum and the maximum storage time T_{min} and T_{max} and the particular storage-time distribution given by the two arrays T_i and N_i . The T_i and N_i -values are calculated using the formulas given in figure D.1 under the boundary condition

$$\sum_{i=0}^N N_i(T_i + T_{dead}) \leq T_{tot}. \quad (\text{D.1})$$

2. Define the number of initially stored UCN N_0 , the number of initially trapped overcritical UCN $N_{0,crit}$ and the storage-time constant of the overcritical UCN τ_{crit} . By default the storage-time constant of the *undercritical* UCN τ_n is set to 880 seconds.
3. For $i = 0$ calculate the *average* number of surviving UCN in the trap using the formulas

$$N(i) = N_0 \exp^{-T_i/\tau_n} \quad (\text{D.2})$$

and

$$N_{crit}(i) = N_{0,crit} \exp^{-T_i/\tau_{crit}} \quad (\text{D.3})$$

4. The actual Monte-Carlo step: draw the number of surviving UCN $N_{Monte-Carlo}(i)$ from the appropriate distribution. For every single neutron in the trap the procedure of waiting for a specific time T_i and checking if it is still there afterward corresponds to a Bernoulli-experiment with a success probability $p = \exp^{-T_i/\tau_n}$ for the undercritical and $\exp^{-T_i/\tau_{crit}}$ for the overcritical UCN. The corresponding distributions are Binomial distributions with mean values of $N(i)$ respectively $N_{crit}(i)$. However this approach does not take into account the variance of the initial filling rates N_0 and $N_{0,crit}$. Fortunately statistics grant us an elegant access to the filling-rate distributions: we can consider the filling process as another random experiment

D. Monte-Carlo Simulations

with a very high number of trials (the number of UCN entering the guide system from the turbine) and an extremely low success probability. Such an experiment can be described using a Poisson distribution which has the major advantage of solely depending on its mean value, in our case N_0 and N_{crit} . Indeed we can even consider the entire measurement, from the filling phase until the measurement of the surviving UCN, as a single random experiment and get the final formula

$$N_{MonteCarlo}(i) = Pois(N(i)) + Pois(N_{crit}(i)), \quad (D.4)$$

where $Pois(x)$ stands for the action of drawing a random number from a Poisson distribution with a mean value x (a detailed review of the statistics used can be found e.g. in [23]). In practice the randomisation was done using the ROOT programming language and particularly the Poisson method of the TRandom3 class. Each time a new instance of the random generator was created the seed was randomised by setting the constructors argument to 0, for more information please refer to [4].

5. Repeat the last step N_i times and store all values in memory.
6. Average over all values and calculate the corresponding error:

$$N_{T_{stor,i}} = \frac{\sum_{n_i} N_{MC}(t)}{n_i} \quad (D.5)$$

$$\delta N_{T_{stor,i}} = \sqrt{\frac{\sum_{n_i} (N_{MC}(t) - N_{T_{stor,i}})^2}{n_i(n_i - 1.5)}} \quad (D.6)$$

7. For $i < N - 1$ repeat steps 1 - 4 for T_{i+1} .
8. Fit an exponential and a double-exponential decay curve to the entire data set. Store the fitting values together with the simulation parameters and dump the residual data before starting allover with a new “beamtime”. This last step is a little worrying as it is surely dangerous to completely automatise the fitting procedure without even having a look at the original data. Unfortunately it was not feasible to implement a more sophisticated analysis scheme. The number of interesting parameters and necessary repetitions per setting is too huge to go for anything else than an entirely automated processing. In the case of the single-exponential fits this does not seem to be a bigger problem. Even for high overcritical UCN shares the fitting values show no tremendously growing spread. For the double-exponential fits I added a supplementary testing-loop that would relaunch the fit with different initial parameters if the storage-time value is evidently corrupted ($\tau_{fit} > 2000$ s) or the fitting routine reported a crash. The testing-loop helped to improve the final output but even after the maximum repetition number of 100 loops the double-exponential fits still produce some extreme outliers.

D. Monte-Carlo Simulations




Figure D.1.: Overview on the different storage-time distributions I used for my MC-simulations. The different dots indicate the storage times, the lines are only to guide the eyes. Left plot: **TcNc**: the distance dT_i in between two storage times is constant and all N_i are the same. **TrNc**: $T_i = T_{i-1} + const * r^{i-1}$, all N_i are the same. **TrNr**: as TrNc with $N_i = const * T_i$. **TcdNconst**: as TcNc with $N_i = const * \exp^{T_i/\tau_n}$. **TrdNconst**: combination of TrNc and TcdNconst. Right plot: **TNrand**: T_i and N_i are randomly chosen. **TcTNconst**: as TcNc with $N_i = \frac{const}{T_i + T_{dead}}$. **TrTNconst**: combination of TrNr and TcTNconst.

In section 3.3.2 I already presented some results from those simulations that had emphasised the influence of the chosen parameter setting on $\delta\tau_{stat}$. Figure D.2 shows a more elaborated study of the problem. The lower most of the plots shows the fitting-error distribution of the entire data set. While this distribution looks fairly arbitrary, the upper two plots show that it can be decomposed into numerous clearly defined peaks that are characterised by their particular measurement-parameter settings. Indeed we see that, dependent on the chosen measurement protocol, the fitting error varies by over 100 %. From all the simulated parameter sets the best result was achieved with the *TrdNconst* measurement-point distribution, a maximum storage time of 1800 seconds and a total of 12 measurement points. As can be seen in figure D.1 the *TrdNconst* distribution has a high concentration of measurement points at low storage times. The number of measurements per storage time is constantly growing with the storage time and thereby compensates for the lower statistics at higher storage times. However it is only slightly better than the more conservative approach of constant dT_i and N_i (called *TcNc* in figure D.1). Of course we will have to consider several additional aspects before defining our final measurement protocol. E.g. we should bear in mind that the emptying time has a finite time span. Very close-packed storage-time points, as they occur in the case of the *TrdNconst* distribution, could eventually get indistinguishable. Consequently a more sophisticated analysis will have to attribute some measurement error to the storage-time value.




Figure D.2.: Simulated distribution of the fitting errors for solely well stored UCN and an average initial filling rate of 1500 UCN. The bottom plot shows the entire data set while the two top plots show the distributions for some particular parameter settings.

D. Monte-Carlo Simulations

Interpreting the results from the simulations *with* an overcritical component is more complicated. The optimal measurement scheme depends strongly on the choice of N_{crit} and τ_{crit} and it is not possible to draw general conclusions. Moreover the significance of the results is dubious due to the very simplified description of the overcritical UCN share by a single storage-time constant. I will therefore content myself with some general remarks on the simulations with overcritical UCN.

Figure D.3 shows the distribution of the τ_{fit} -values for different N_{crit} - τ_{crit} pairs. The plots show the entire simulated data set, no difference is made between different measurement protocols. The blue and the red curves show the respective fitting results of the single- and the double-exponential fits. When we only consider the single-exponential fits we see that under certain conditions - a very small overcritical UCN share and/or a high τ_{crit} - the obtained storage time is mostly identical with the mean storage time $\bar{\tau}$ as calculated with equation 2.12 from section 2.2.1. If these conditions are not fulfilled the results get much more arbitrary and the extraction of a useful τ_n -value can be excluded. The double-exponential fits yield results that are closer to the predefined τ_n -value. However they still show a large spread with a lot of outliers and are surely not suited to extract a meaningful τ_n -value.

The difference between the advantageous and disadvantageous fitting conditions also appears in the χ^2 -values of the fits. The expectation value of χ^2 is equal to the number of degrees of freedom of the fit. For a single-exponential fit with two fitting parameters, τ_n and N_0 , a total of eight measurement points and ideal conditions we should thus get a mean χ^2 -value of 6. Indeed if I average over all single-exponential fits with $N_{crit} = 0$ and $N = 8$ I get $\bar{\chi}^2 = 6.05$. In the case of an overcritical UCN ratio of 0.2 % and a τ_{crit} -value of 200 s the average only increases slightly to $\bar{\chi}^2 = 6.29$. However if I increase the overcritical UCN share to 3.3 % the mean χ^2 -value decouples to 68.88 indicating a miserable fit.

The double-exponential fits yield results that are closer to the predefined τ_n -value. However they still show a large spread with a lot of outliers and are surely not suited to extract a meaningful τ_n -value. In the best case it can be a backing solution helping us to identify and roughly characterise an overcritical residuum in the data. I insist on this point in order to emphasise that there will be absolutely no possibilities to *postclean* our data. If we are not able to reduce the overcritical UCN contamination down to some acceptable rate not even the most sophisticated analysis will enable us to extract a useful τ_n -value.

D. Monte-Carlo Simulations




Figure D.3.: Distribution of the τ_{fit} -values for single- and double-exponential fitting in dependency of the number of overcritical UCN and their storage time. The dashed vertical lines indicate the mean storage times of the respective UCN ensembles.

E. Stray field measurements



Figure E.1.: Experimental physicist measuring the stray fields of his experiment using a cord and a pair of scissors.

Bibliography

- [1] Arepoc. *High Linearity Hall Probes for Room and Cryogenic Temperatures*, 2008.
- [2] S. Arzumanov et al. Neutron life time value measured by storing ultracold neutrons with detection of inelastically scattered neutrons. *Physics Letters B*, 483, Iss. 1:15 – 22, 2000.
- [3] S. Arzumanov et al. Analysis and correction of the measurement of the neutron lifetime. *Journal of Experimental and Theoretical Physics Letters*, 95, Iss. 5:224 – 228, 2012.
- [4] CERN. <https://root.cern.ch/drupal/>.
- [5] A. Chao and M. Tigner. *Handbook of Accelerator Physics and Engineering, 3rd edition*. World Scientific Publishing, 1999.
- [6] Cryomech. *Model PT415 Cryogenic Refrigerator - Installation, operation and routine maintenance manual*.
- [7] M. Daum et al. First observation of trapped high-field seeking ultracold neutron spin states. *Physics Letters B*, 704:456–460, 2011.
- [8] S. Dzhosyuk et al. Determination of the Neutron Lifetime Using Magnetically Trapped Neutrons. *Journal of Research of the National Institute of Standards and Technology*, 110:339 – 343, 2005.
- [9] ESRF. <http://www.esrf.eu/accelerators/groups/insertiondevices/software/radia>.
- [10] V. Ezhov et al. Measurement of the neutron lifetime with ultra-cold neutrons stored in a magneto-gravitational trap. *arXiv:1412.7434*, 2014.
- [11] E. Fermi. Versuch einer Theorie der beta-Strahlen. *Zeitschrift für Physik*, 88 (3-4):161 – 177, 1934.
- [12] E. Fermi. Sul moto dei neutroni nelle sostanze idrogenate. *Ricerca scientifica*, 7 no. 2:13 – 52, 1936.
- [13] E. Fermi et al. Artificial Radioactivity Produced by Neutron Bombardement. *Proceedings of the Royal Society A*, 146:483 – 500, 1934.
- [14] R. P. Feynman and M. Gell-Mann. Theory of the Fermi interaction. *Physical Review*, 109, no. 1:193 – 198, 1958.

Bibliography

- [15] M. Goldhaber et al. Helicity of Neutrinos. *Physical Review*, 109 (3):1015 – 1017, 1958.
- [16] R. Golub, D. Richardson, and S. K. Lamoreaux. *Ultra-Cold Neutrons*. IOP Publishing Ltd, 1991.
- [17] K. Halbach. Design of permanent multipole magnets with oriented rare earth cobalt material. *Nuclear Instruments and Methods*, 169, no. 1:1–10, 1980.
- [18] M. Kobayashi and T. Maskawa. CP-Violation in the Renormalizable Theory of Weak Interaction. *Progress of Theoretical Physics*, 49 no. 2:652 – 657, 1972.
- [19] LakeShore. *Gaussmeter Hall Probes - Catalog*.
- [20] LakeShore. *Model 218 Temperature Monitor User's Manual*, 2012.
- [21] S. K. Lamoreaux and R. Golub. Experimental Searches for the neutron electric dipole moment. *Journal of Physics G*, 36, no. 10:104002, 2009.
- [22] G. Le Bec, J. Chavanne, and C. Benabderrahmane. Simulation of NdFeB permanent magnets at low temperature. *Proceedings of PAC09, Vancouver, BC, Canada*, T15:327–329, 2009.
- [23] W. R. Leo. *Techniques for Nuclear and Particle Physics Experiments, 2nd edition*. Springer, 1994.
- [24] K. Leung. *Development of a new superfluid helium ultra-cold neutron source and a new magnetic trap for neutron lifetime measurements*. PhD thesis, Technische Universitaet Muenchen, 2013.
- [25] K. Leung et al. A Comparison of Two Magnetic Ultra-Cold Neutron Trapping Concepts Using a Halbach-Octupole Array. *Proceedings of the 2012 Workshop "Next Generation Experiments to Measure the Neutron Lifetime", Santa Fe, New Mexico, 9 - 10 November 2012*, page 192, 2014.
- [26] K. Leung and O. Zimmer. Proposed neutron lifetime measurement using a hybrid magnetic trap for ultra-cold neutrons. *Nuclear Instruments and Methods A*, 611, no. 2-3:216–218, 2009.
- [27] W. J. Marciano and A. Sirlin. Improved calculation of electroweak radiative corrections and the value V_{ud} . *Physical Review Letters*, 96, Iss. 3:032002, 2006.
- [28] D. C. Meeker. Finite element method magnetics (FEMM), version 4.0.1 (build 03 dec). <http://www.femm.info>, 2006.
- [29] J. S. Nico et al. Measurement of the neutron lifetime by counting trapped protons in a cold neutron beam. *Physical Review C*, 71, Iss. 5:055502, 2005.
- [30] K. A. Olive et al. 2014 Review of Particle Physics. *Chinese Physics C*, 38:090001, 2014.

Bibliography

- [31] Oswald Elektromotoren. *Dokumentation für supraleitende Magnetfeldspulen*, 2012. Custom manual for our superconducting coil system.
- [32] W. Paul et al. Measurement of the neutron lifetime in a magnetic storage ring. *Zeitschrift für Physik C*, 45, Iss. 1:25 – 30, 1989.
- [33] Pfeiffer Vacuum. *Compact FullRange Gauge PKR 251*, 2005.
- [34] A. Pichlmaier et al. Neutron lifetime measurement with the UCN trap-in-trap MAMBO II. *Physics Letters B*, 693:221 – 226, 2010.
- [35] J. M. Robson. The Radioactive Decay of the Neutron. *Physical Review*, 83, Iss. 2:349, 1951.
- [36] A. Sakharov. Violation of CP Invariance, C Asymmetry and Baryon Asymmetry of the Universe. *Journal of Experimental and Theoretical Physics Letters*, 5, 1967.
- [37] A. Serebrov et al. UCN anomalous losses and the UCN capture cross section on material defects. *Physics Letters A*, 335, Iss. 4:327 – 336, 2005.
- [38] A. Serebrov et al. Neutron lifetime measurements using gravitationally trapped ultracold neutrons. *Physical Review C*, 78:035505, 2008.
- [39] A. Serebrov et al. Ultracold-neutron infrastructure for the PNPI/ILL neutron EDM experiment. *Nuclear Instruments and Methods A*, 611, no. 2 - 3:263 – 266, 2009.
- [40] M. Simson. *Measurement of the electron antineutrino angular correlation coefficient a with the neutron decay spectrometer aSPECT*. PhD thesis, Technische Universität München, 2010.
- [41] A. Steyerl et al. A new source of cold and ultracold neutrons. *Physics Letters A*, 116, no. 7:347 – 352, 1986.
- [42] E. C. G Sudarshan and R. E. Marshak. Chirality invariance and the universal fermi interaction. *Physical Review*, 109, no. 5:1860 – 1862, 1958.
- [43] Vacuumschmelze. Rare-Earth Permanent Magnets: Vacodym - Vacomax. <http://www.vacuumschmelze.de>, 2012.
- [44] F. E. Wietfeldt and G. L. Greene. Colloquium: The neutron lifetime. *Reviews of Modern Physics*, 83, no. 4:1173 – 1192, 2011.
- [45] C. S. Wu et al. Experimental tests of parity conservation in beta decay. *Physical Review*, 105(4):1413 – 1415, 1957.
- [46] A. T. Yue et al. Improved Determination of the Neutron Lifetime. *Physical Review Letters*, 111:222501, 2013.
- [47] O. Zimmer. A method of magnetic storage of ultra-cold neutrons for a precise measurement of the neutron lifetime. *Journal of Physics G*, 26, no. 1:67 – 77, 2000.

Bibliography

- [48] O. Zimmer et al. Superthermal source of ultracold neutrons for fundamental physics experiments. *Physical Review Letters*, 107, no. 13:134801, 2011.

Acknowledgments

First of all I want to thank Oliver Zimmer for his support and never ending optimism, seeing challenges instead of issue and always keeping on track. I wish him best luck with all his future endeavors and steadily increasing neutron numbers for him and HOPE.

Special thanks goes to all those who have helped me pull through the long months spent on level D during the data taking. Fabien Lafont for pushing innumerable helium and nitrogen bottles. Sergey Ivanov for always being available and ready to help. Peter Geltenbort for sharing almost a lifetime of experience in UCN physics. Thomas Brenner for the technical support and the daily craning routine. Florian Martin for his support but also for the shared trips to Abingdon. Loris Babin who only had a few days to get to know the apparatus but who will have so much more time to enjoy it during the next few years. And also Thorsten Lauer for demonstrating his in-situ skills in UCN instrumentation.

Of course the work did not start at the level D gate and there are many more people who deserve my gratitude. Martin Simson for his enduring support and valuable advice but also for the ski trips and just having a good time. Didier Berruyer for still standing in the workshop when even the scientific staff had already left. Kent Leung for his groundwork and unbroken interest in the project. The entire NPP group for all the helping hands, the discussions and the occasional distraction. My fellow PhD students, not only from the ILL but also all the expats from PF1 and PF2, for making even night and weekend shifts a fun experience. Innumerable people all over the ILL who helped out whenever they could. And a whole lot of other people in and outside of science who have made these last years a success and an enjoyable time.