

A MAGNETO-GRAVITATIONAL NEUTRON TRAP
FOR THE MEASUREMENT OF THE NEUTRON LIFETIME

Daniel J. Salvat

Submitted to the faculty of the University Graduate School
in partial fulfillment of the requirements for the degree
Doctor of Philosophy
in the Department of Physics,
Indiana University
April 2015

Accepted by the Graduate Faculty, Indiana University, in partial fulfillment of the requirements
for the degree of Doctor of Philosophy.

Doctoral Committee

Chen-Yu Liu, Ph.D.

William M. Snow, Ph.D.

Rick Van Kooten, Ph.D.

Roger Pynn, Ph.D.

March 12, 2015

Copyright ©2015
Daniel J. Salvat

Acknowledgements

Thanks Boss.

Daniel J. Salvat

A MAGNETO-GRAVITATIONAL NEUTRON TRAP FOR THE MEASUREMENT OF THE NEUTRON LIFETIME

Neutron decay is the simplest example of nuclear beta-decay. The mean decay lifetime is a key input for predicting the abundance of light elements in the early universe. A precise measurement of the neutron lifetime, when combined with other neutron decay observables, can test for physics beyond the standard model in a way that is complimentary to, and potentially competitive with, results from high energy collider experiments. Many previous measurements of the neutron lifetime use ultracold neutrons (UCN) confined in material bottles. In a material bottle experiment, UCN are loaded into the apparatus, stored for varying times, and the surviving UCN are emptied and counted. These measurements are in poor agreement with experiments that use neutron beams, and new experiments are needed to resolve the discrepancy and precisely determine the lifetime. Here we present an experiment that uses a bowl-shaped array of NdFeB magnets to confine neutrons without material wall interactions. The trap shape is designed to rapidly remove higher energy UCN that might slowly leak from the top of the trap, and can facilitate new techniques to count surviving UCN within the trap. We review the scientific motivation for a precise measurement of the neutron lifetime, and present the commissioning of the trap. Data are presented using a vanadium activation technique to count UCN within the trap, providing an alternative method to emptying neutrons from the trap and into a counter. Potential systematic effects in the experiment are then discussed and estimated using analytical and numerical techniques. We also investigate solid nitrogen-15 as a source of UCN using neutron time-of-flight spectroscopy. We conclude with a discussion of forthcoming research and development for UCN

detection and UCN sources.

Chen-Yu Liu, Ph.D.

William M. Snow, Ph.D.

Rick Van Kooten, Ph.D.

Roger Pynn, Ph.D.

Contents

1	Neutron Decay	1
1.1	Introduction	1
1.2	Neutron Decay in the Standard Model	1
1.3	Effective Theory of Neutron β -decay	4
1.4	The Predicted Neutron Lifetime	7
1.5	The Impact of an Improved τ_n Measurement	9
1.6	Conclusions	11
2	The History of τ_n	12
2.1	Introduction	12
2.2	First Measurements	15
2.3	Improved In-Beam Measurements	16
2.4	The Material Bottle Method	20
2.5	Improved Bottle Measurements	23
2.6	Corrections and Criticisms of Bottle Experiments	25
2.7	Magnetic Bottles	27
2.8	Conclusions	30
3	Experimental Design	32
3.1	Introduction	32
3.2	Permanent Magnet Trap	33
3.3	Trap Door and UCN Guides	36

3.4	Holding Field Coils	38
3.5	AFP Spin Flipper	40
3.6	UCN Cleaner	41
3.7	UCN Detectors	42
3.8	Vanadium Activation Detector	51
3.9	Automation and Data Acquisition	54
4	First Experimental Campaign	57
4.1	Introduction	57
4.2	Backgrounds	57
4.3	Determination of the Storage Time	59
4.4	Cleaner Upscatter Detectors	63
4.5	Conclusions	64
5	Second Experimental Campaign	66
5.1	Introduction	66
5.2	Vanadium Detector Characterization	66
5.3	The ^{52}V Mean Lifetime	80
5.4	Improved UCN Transport	81
5.5	Studies with the Vanadium Foil	84
5.6	Discussion and Conclusions	88
6	Systematic Effects	90
6.1	Introduction	90
6.2	Residual Gas	91
6.3	Depolarization	94
6.4	Material Losses	95
6.5	Cleaning Quasi-Bound UCN	96
6.6	Cleaner-Generated Effects	101

6.7	Microphonic Heating	107
6.8	Gain Drifts	115
6.9	Dead Time and Pileup	117
6.10	Time Dependent Backgrounds	118
6.11	Phase Space Evolution and Vanadium Activation	120
6.12	UCN Source Fluctuations	123
6.13	Summary and Conclusions	126
7	Solid Nitrogen as a UCN Converter	128
7.1	Introduction	128
7.2	Experiment	129
7.3	Results	130
7.4	Discussion	132
7.5	Conclusions	136
8	Conclusions	137
8.1	Summary and Overview	137
8.2	Outlook	139
	Appendices	142
A	Slow Neutrons	143
A.1	Introduction	143
A.2	Slow Neutrons and Nuclei	144
A.3	Neutron Scattering from Condensed Matter	145
A.4	Crystalline Solids	148
A.5	Ultracold Neutron Production	151
A.6	Ultracold Neutrons	153
B	Experimental Modeling	157

B.1	Introduction	157
B.2	Kinetic Theory Model of the Trap	157
B.3	Neutron Tracking	160
Bibliography		165
Curriculum Vitae		
 Tables		
5.1	The extracted vanadium lifetimes from each run pair. For the first run pair, the second run could not be used. The values are weighted by the square of their uncertainties and combined to obtain the final result.	82
5.2	The GV-normalized $\text{xor-}\beta$ signal and flipper-on to flipper-off contrast for each geometry, each normalized to the signal and contrast of the original geometry with the trap door and piston drive present. No flipper-off data were acquired for the bare geometry.	84
6.1	Potential systematic effects related to storing, counting, and normalizing.	91
6.2	Estimates of potential systematic effects in the current experiment, with corrections where applicable.	127
A.1	Four classes of slow neutrons, with approximate energy, velocity, wavelength, wavenumber, and temperature ranges.	143
B.1	Some solutions for the coefficients c_i and d_i . The third order Ruth integrator is from ref. [117].	163

Figures

1.1	The decay of the free neutron at quark-level.	4
1.2	The decay of the free neutron in the low energy effective theory.	5
1.3	The combined determination of V_{ud} and λ from $0^+ \rightarrow 0^+$ nuclear decays, the β -asymmetry parameter A , and the neutron lifetime τ_n	8
2.1	The mean neutron lifetime τ_n versus time.	13
3.1	A schematic of the experiment. P is the polarizing magnet, S is the spin flipper, and B is the UCN monitor detector. There is another monitor to the left of the polarizing magnet (not shown). The holding field coils (not shown) are arranged outside of and around the vacuum jacket of the trap.	33
3.2	The completed permanent magnet trap, prior to insertion into the vacuum jacket. The rectangular hole at the bottom is filled by a separate magnet plate (not shown here) which is fastened to a pneumatic actuator.	34
3.3	The magnetic field generated by a Halbach array. The red arrows represent the magnetization direction of each magnet, and the black arrows represent the magnetic field.	35
3.4	UCN are directed to or through the tubular sections to the left or right. A UCN that goes up passes through the brass spokes, and up around the steel rod and magnet plate (top). The pneumatic piston is beneath the guide cross to push the piston rod upwards or downwards.	37

3.5	A cutaway of the trap door and guide cross assembly. The cross assembly (middle) surrounds the piston rod. UCN can pass through the spoke holes up to the top of the cross assembly, past the magnet plate (door), and into the trap. The features above the guide cross assembly are encased in the vacuum vessel of the trap. . . .	38
3.6	The holding coils (shown from the side in red) produce field lines (dotted blue) which are perpendicular to the field of the Halbach array.	39
3.7	The design of the holding coil. Each L-shaped copper bundle rests in an 80/20 frame, and the frames are fastened together with joining brackets.	40
3.8	A closeup of the individual copper conductors. The bars are fanned out, and fitted with water connections for cooling. Each bar is electrically connected to its corresponding bar on the other L-split with a copper bussing clamp.	41
3.9	The component of the magnetic field parallel to the UCN beam axis produced by the AFP solenoid with 5 A of applied current.	42
3.10	The UCN cleaner. The rectangular polyethylene sheet is fastened to the Al frame, which swivels on four legs so that its height can be changed. The linkage on the back of the sheet connects to a pneumatic actuator which is fed through a bellows from outside the vacuum jacket. The two Al wing-shaped plates bolt the assembly to the frame of the Halbach array.	43
3.11	A cutaway of the apparatus showing the placement of the UCN cleaner.	44
3.12	A schematic of the UCN detector assembly. The bottom plate, aluminum window, and polyoxymethylene cap are sealed with o-rings.	45
3.13	The detector mount configuration. Lengths are not to scale.	46
3.14	Pulse-height spectra for the helium and boron-coated ionization chambers using UCN. The vertical lines represent the Li and α energies of 0.84, 1.02, 1.47, and 1.78 MeV.	47

3.15	The discriminated count rate from the boron-coated detector as a function of applied anode bias. The discrimination threshold is adjusted for each voltage as needed to remove low energy background pulses.	48
3.16	Comparison of boron-coated detector spectra for UCN and thermal neutrons. The spectra are scaled so that their integral is unity. The predictions are discussed in section 3.7.5.	49
3.17	A cutaway of the vanadium detector. A 5 cm thick rectangular Pb background γ -ray shield surrounds the array of 8 NaI detectors. The vacuum jacket between the NaI detectors houses the V foil (which can be lowered into the trap) which is surrounded by the plastic β detectors.	52
3.18	A schematic of the detector geometry. The Pb shield surrounds the stack of NaI detectors. The black line represents the vacuum break. The plastic scintillators immediately surround the V foil within vacuum. The foil is moved down into the trap using a magnetic coupler to connect to a linear actuator outside the vacuum.	53
4.1	The average beam off background rates for the ^{10}B detector over the run campaign. The red line is a linear fit for runs greater than 140, with a slope of approximately 90 mHz per day.	58
4.2	The beam off background rate for the ^{10}B detector with a lower level ADC cut of 0.65 V.	58
4.3	The timing of components during a fill and empty cycle. The beam is on (off), valves are open (closed), and the cleaner is down (up) for a high (low) signal.	60
4.4	The ^{10}B counter rate during a measurement cycle. From left to right, the vertical lines represent time t_{pre} when the shutter is closed, t_{fill} when the trap door is closed and shutter opened, and t_{empty} when the trap door is opened.	61
4.5	The signal S versus t_{store} . The storage time constant of the trap is given by τ_{store} from the exponential fit (upper). The distribution of residuals of the exponential fit are normalized to their statistical uncertainty (lower).	62

4.6	The pulse height spectrum of the ^3He tubes (left). The pulse height spectrum of the ^3He tubes (right).	63
4.7	The combined counts over time during 30 s cleaning, with exponential fit (left). The combined counts over time for fill and empty runs with no cleaning (right).	64
5.1	The measured ^{60}Co spectrum using the summed NaI detectors. The black shows the measured spectrum, the red shows the spectrum with the source removed, and the blue is a direct subtraction of the foreground and background spectra.	67
5.2	The pulse height spectrum of vanadium activation events in coincidence with plastic scintillator events. The blue line shows the gaussian fit within the fit range, denoted by vertical dotted lines (left). The relative peak position from the gaussian fits to the high-statistics vanadium activation data versus time, along with the linear fit (right).	69
5.3	The time-to-nearest-event for the beam on period of run 655, activity measurement during that run, and beam off background from run 656, each in a 200 s window. The (left) time-to-nearest-event for the two plastic scintillators and (right) plastic scintillators and NaI detectors are shown.	69
5.4	The relative uncertainty in the activity as determined by the fit to eqn 5.1 for xor β events (left) and xor β coincident with NaI events (right).	71
5.5	The singles and coincidence rates in the detector for the high statistics run (see text).	72
5.6	Digitized waveforms from the phoswich scintillators. An event with just a slow component (black) and an event with a fast and slow component (blue) are shown.	73
5.7	(left) A typical digitized waveform (solid) with the associated double exponential fit (dotted line). (right) The same waveform (solid) with vertical dotted lines representing the domain of integration for the fast part of the pulse; all times after the later dotted line form the domain of slow integration.	74

5.8	The fast and slow integrals of events from run 826, a vanadium activation run in a modified guide geometry. The xor β events (left) and xor β - γ coincidences (right) are shown.	75
5.9	The angle of the slope of fast and slow intensity for xor β and coincidence events. Run 826 (vanadium activation) is shown in black, and run 827 (background) is shown in blue. In each case, events for times after the vanadium has been raised into the detector array are included.	76
5.10	The best-fit gain for trial runs compared to the initial reference run versus average measurement time, in hours.	77
5.11	The background rates for the different run types. Rates for β singles events (upper left), xor β events (upper right), NaI singles (lower left) and xor β - γ coincidences (lower right) are shown.	78
5.12	The background in the NaI detector just after the beam was on for 400 s. The red curve is a fit to a constant background plus a decaying component, consistent with the ~ 25 m half-life of ^{128}I	79
5.13	An example of ^{52}V lifetime run pairs. The beam-on data is cut from the runs, and the foreground and background combined together with the appropriate time offset between the runs (upper), with the red showing the best fit. The lower panel shows the fit residuals.	81
5.14	The bare (upper left), elbow (upper right), elbow & plate (lower left), and box (lower right) geometries that were tested to improve UCN transport efficiency into the trap. Cu components are shown in red, and brass components in yellow. . . .	83
5.15	The extraction of the vanadium signal for a typical run. The time distribution of β -xor events (top) is shown along with the fit to extract the number of vanadium counts. The GV rate (bottom) is shown along with the window within which the normalization rate M is computed (red vertical lines).	85

5.16	The normalized xor β signal as a function of vanadium draining time in the nominal guide configuration. The best fit for an exponential approach is also shown. . . .	86
5.17	The height dependence of the detector signal after filling for 200 s.	87
6.1	98
6.2	The rate of incident UCN upon the cleaner versus time. Data are fit to two-exponential functions, and the best-fit for the time constants t_1 and t_2 are shown. The fitted amplitudes of the exponential terms are (relative to the legend) 121.5 and 46.0 in relative units.	99
6.3	A trajectory that is not cleaned by just the prototype cleaner. The black line represents the spatial trajectory, and the red box outlines the position of the cleaner.100	
6.4	The rate of cleaning using the prototype cleaner and additional polyethylene sheet. No UCN remain after ~ 200 s, suggesting that less than 6×10^{-4} . The fitted amplitudes of the three exponential terms are (relative to the plot legend) 681.6, 334, 4, and 30.2, in relative units.	101
6.5	The initial spectrum entering the trap (blue) and spectrum after the trap has been filled (red). The vertical black dotted line corresponds to the maximum trappable UCN energy.	102
6.6	Number of surviving UCN for the three cleaning heights.	103
6.7	The relative number of trapped UCN after 100 seconds for the three different cleaning heights. The black line is a linear fit.	104
6.8	Vertical linear density (normalized to unity) extracted from simulation without cleaning, after 100 s.	105
6.9	The rate of absorption of UCN on the cleaner for 3 different cleaning heights. . .	106
6.10	Plots of z vs. t for various phases. Red: 0; blue: $\pi/4$; green: π	109
6.11	Plots of E vs. t for various phases. Red: 0; blue: $\pi/4$; green: π	110
6.12	Energy transfer after 1 bounce as a function of δ to a UCN with $h = 0.44$, $f = 40$, $A = 10^{-5}$	111

6.13	Energy transfer after 1 bounce as a function of A to a UCN with $h = 0.44$, $f = 40$, $\delta = 2.0$	112
6.14	Energy transfer after 1 bounce as a function of f to a UCN with $h = 0.44$, $\delta = 2.0$, $A = 10^{-5}$	113
6.15	The initial (upper) and final (lower) spectra after $N = 3000$ bounces.	114
6.16	The VSD along a particular axis of the vacuum jacket on the UCN beamline at LANSCE.	115
6.17	The relative correction to the lifetime (for nominal 885 s lifetime) as a function of the vanadium absorption time t_m (left) and as a function of the change of draining η (right).	121
6.18	The rate of absorption on the vanadium versus time after three different storage times. Two-exponential fits are also shown with best-fit values for time constants t_1 and t_2	122
6.19	The rate in a monitor detector based on the convolution of the system response $G(t)$ with sequential pulse chains; the black curve shows the rate assuming a sequence of constant pulse chains, and the red shows a constant set of pulse chains, except for a single pulse chain with 1% lower UCN output during filling (top). The same rates, but zoomed in on the region of nearly-constant rate during filling, over which the rate is averaged to compute the normalizing factor (bottom).	125
7.1	The differential scattering cross-section versus E . $E > 0$ corresponds to energy loss, and $E < 0$ to energy gain. The vertical scale is set by integrating to the total scattering cross-section (see text).	130
7.2	The dynamic structure factor of polycrystalline $\alpha\text{-}^{15}\text{N}_2$. The color scale is set by integrating the measured differential cross-section and equating it to the total scattering cross-section. The black line corresponds to the UCN production curve given by eqn 7.2.	131

7.3	The GDOS for solid nitrogen at two temperatures in the α -phase. Peak broadening at higher temperature is observed.	132
7.4	The UCN production cross-section in $^{15}\text{N}_2$ for $E_{UCN} \leq 181$ neV versus incident neutron energy. The inelastic cross-section for $E < 2$ meV is difficult to determine due to elastic contamination, though is likely small compared to that of the energy range shown here.	133
7.5	The UCN yield of $^{15}\text{N}_2$ for a neutron flux of $10^{14} \text{ cm}^{-2} \text{ s}^{-1}$ from a cold moderator at temperature T . The yield in UCN is optimized for incident neutrons at a temperature $T = 40$ K.	134
7.6	The mean free path λ_{up} for UCN to upscatter to non-UCN energies within the solid nitrogen volume. The error bars are statistical.	135
B.1	The quantity $\omega(\epsilon) \equiv d\Omega/d\epsilon$ for typical UCN energies. The results of the Monte Carlo integration are reasonably approximated by a polynomial.	158
B.2	A procedural representation of the symplectic integration method for separable, time independent Hamiltonians.	162
B.3	The relative change in particle energy for the pendulum problem using fourth order symplectic (left) and fourth order Runge-Kutta (right) algorithms. The symplectic method demonstrates the long-term stability of the energy.	163
B.4	The spatial trajectory of a trapped UCN (left and bottom panels), and the calculated energy of the neutron versus time.	164
B.5	The spatial trajectory of a trapped UCN with energy near the maximum trappable energy (left and bottom panels), and the calculated energy of the neutron versus time.	165

1 Neutron Decay

1.1 Introduction

The free neutron undergoes β decay

$$n \rightarrow p + e + \bar{\nu}_e \quad (1.1)$$

with a Q value of approximately 782 keV. A theoretical model for β -decay was first proposed by Fermi[83], which was subsequently modified due to the discovery of parity violation, and eventually incorporated into the framework of electroweak theory. Theoretical considerations related specifically to neutron decay have been recently reviewed in the literature[18, 134]. In this chapter we briefly review the prediction of neutron decay from the standard model and effective low energy framework to describe the decay, following the notation used in standard texts[126]. We then discuss the current experimental knowledge of neutron decay, the predicted value of the neutron mean lifetime, and the impact of its measurement on the fields of particle physics and cosmology.

1.2 Neutron Decay in the Standard Model

At quark level in the standard model (SM), the decay of the neutron is due to the coupling of the two lightest (u and d) quarks to the electroweak gauge fields. The u and d form a weak SU(2) doublet which leads to the interaction

$$\mathcal{L}_{SM} \supset \frac{1}{\sqrt{2}}g_2 W_\mu^+ Q^{-\mu} + \frac{1}{\sqrt{2}}g_2 W_\mu^- Q^{+\mu} + \frac{e}{s_W c_W} Z_\mu Q_Z^\mu + e A_\mu Q_{EM}^\mu \quad (1.2)$$

with quark currents

$$Q^{+\mu} = \bar{d}_{(L)I} V_{IJ}^\dagger \gamma^\mu u_{(L)J} \quad (1.3)$$

$$Q^{-\mu} = \bar{u}_{(L)I} V_{IJ} \gamma^\mu d_{(L)J} \quad (1.4)$$

$$Q_Z^\mu = \frac{1}{2} \bar{u}_{(L)I} \gamma^\mu u_{(L)I} - \frac{1}{2} \bar{d}_{(L)I} \gamma^\mu d_{(L)I} - s_W^2 Q_{EM}^\mu \quad (1.5)$$

$$Q_{EM}^\mu = \frac{2}{3} \bar{u}_I \gamma^\mu u_I - \frac{1}{3} \bar{d}_I \gamma^\mu d_I. \quad (1.6)$$

Here, s_W and c_W are the sine and cosine of the weak mixing angle θ_W , e is the electric charge, $g_1/g_2 = \tan \theta_W$, and W_μ and Z_μ are (respectively) the charged and neutral weak gauge bosons. The indices $I = 1, 2, 3$ and $J = 1, 2, 3$ (repeated indices summed) represent the three different quark generations. The subscript (L) represents the fact that the quark fields are left-handed projections. That is, the fields in the weak currents are given by $u_{(L)} = P_{(L)} u = \frac{1}{2}(1 - \gamma_5)u$, where $P_{(L)}$ is the left handed projection operator. The fermion currents thus take the form of a vector current minus an axial vector current (the so-called V-A form). This has the implication that the theory is *chiral* – weak interactions maximally violate parity symmetry.

The three generations of quark fields appearing in eqns. 1.3 through 1.6 acquire a mass through the yukawa coupling to the Higgs field, and their representation is chosen so as to diagonalize the mass terms. It is, however, observed that these mass eigenstates are not simultaneously diagonal with respect to the weak interaction terms; they are related to the mass terms via unitary matrices

$$d_I \rightarrow D_{IJ} d_J \quad (1.7)$$

$$u_I \rightarrow U_{IJ} u_J \quad (1.8)$$

from which we see that the matrix V_{IJ} appearing in the above is given by

$$V = U^\dagger D. \quad (1.9)$$

This matrix is called the Cabibbo-Kobayashi-Maskawa (CKM) matrix. It has the implication that weak interactions mix the different quark flavors so that, for example, a d quark can decay into a u quark, which leads to β -decay. Writing the CKM matrix explicitly in terms of the different quark flavors, we have that

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix} \quad (1.10)$$

where the matrix on the right hand side is the CKM matrix, the vector on the right hand side contains the d, s, b mass eigenstates, while the fields on the left hand side are those that appear in the weak interaction terms.

The lepton sector also participates in charged-current weak interactions:

$$\mathcal{L}_{SM} \supset \frac{1}{\sqrt{2}} g_2 W_\mu^+ L^{-\mu} + \frac{1}{\sqrt{2}} g_2 W_\mu^- L^{+\mu} + \frac{e}{s_W c_W} Z_\mu L_Z^\mu + e A_\mu L_{EM}^\mu \quad (1.11)$$

with the lepton currents given by

$$L^{+\mu} = \bar{e}_{(L)I} \gamma^\mu \nu_{(L)I} \quad (1.12)$$

$$L^{-\mu} = \bar{\nu}_{(L)I} \gamma^\mu e_{(L)I} \quad (1.13)$$

$$L_Z^\mu = \frac{1}{2} \bar{\nu}_{(L)I} \gamma^\mu \nu_{(L)I} - \frac{1}{2} \bar{e}_{(L)I} \gamma^\mu e_{(L)I} - s_W^2 L_{EM}^\mu \quad (1.14)$$

$$L_{EM}^\mu = -\bar{e}_I \gamma^\mu e_I \quad (1.15)$$

where the generation index $I = 1, 2, 3$ represents electrons, muons, and taus. The presence of these quark and lepton currents in the SM permits the tree-level diagram shown in fig. 1.1.




Figure 1.1: The decay of the free neutron at quark-level.

1.3 Effective Theory of Neutron β -decay

The decay of the neutron happens at low energy compared to the QCD scale, where confinement makes perturbative calculations involving the quarks not feasible. We thus turn to a low energy effective theory to describe the weak interaction of the bound states of the quarks. Noting that neutron decay occurs well below the weak scale (so that the W field can be integrated out), we can write the effective Hamiltonian for neutron decay in terms of a current-current interaction between the nucleons and leptons

$$\mathcal{H}_{eff} = \frac{1}{\sqrt{2}} G_F J_n^\mu J_{l\mu} \quad (1.16)$$

with nucleon and lepton currents

$$J_n^\mu = V_{ud} \bar{p} (\gamma^\mu g_V + g_A \gamma^\mu \gamma_5 - i g_M \sigma^{\mu\nu} q_\nu / 2M + g_P \gamma_5 q^\mu) n \quad (1.17)$$

$$J_l^\mu = \bar{e} \gamma^\mu (1 - \gamma_5) \nu_e. \quad (1.18)$$

Here, the Fermi coupling $G_F = e^2 / \sqrt{8} s_W^2 M_W^2$ (M_W is the W boson mass) sets the overall strength of the interaction. The matrices γ_μ are the Dirac matrices, $\sigma_{\mu\nu} = (i/2)(\gamma_\mu \gamma_\nu - \gamma_\nu \gamma_\mu)$, M is twice the nucleon mass, and q^μ is the four-momentum transfer of the W^- . Further, the

vector and axial-vector currents now include form factors g_V and g_A (evaluated here at $q^2 = 0$) which account for the modification of the weak currents by the strong interaction. It is sufficient to evaluate the form factors at zero momentum transfer because the energy of the decay is small compared to the interaction strength set by G_F . A consequence of electroweak unification is that the weak vector current is a conserved quantity as with the electromagnetic vector current, so that the strong force cannot change the neutron's weak vector charge. This is known as the conserved vector current (CVC) hypothesis, and it has the implication that $g_V = 1$. There can be small apparent deviations in g_V for the neutron due, for example, to the difference in u and d quark mass, but such effects are expected to contribute at the 10^{-5} -level. As we will see later, this is small compared to the $\sim 10^{-3}$ experimental uncertainties in measured neutron β -decay parameters, and is typically ignored in analyses[98]. There is no such conservation law for the axial current, and spontaneous chiral symmetry breaking in low energy QCD indeed causes $g_A \neq 1$.




Figure 1.2: The decay of the free neutron in the low energy effective theory.

Due to the internal structure of the neutron, other currents in addition to the vector and axial vector currents in eqn. 1.17. The weak magnetism term g_M is the result of the nucleon exhibiting higher multipoles of weak vector charge. With the CVC hypothesis, it can be shown

that g_M is simply related to the anomalous magnetic moments κ_n of the neutron and κ_p of the proton, which are well known[61]. There is also an induced pseudoscalar form factor g_P which is negligible at neutron decay energies. With the above considerations, we write the nucleon current as

$$J_n^\mu = V_{ud}\bar{p}(g_V\gamma_\mu + g_A\gamma_\mu\gamma_5 + (\kappa_p - \kappa_n)\sigma_{\mu\nu}q^\nu/2M)n. \quad (1.19)$$

Thus, taking κ_n and κ_p as known inputs, neutron decay depends most sensitively on V_{ud} and g_A , the latter of which is often cast as $\lambda \equiv g_A/g_V$.

As with nuclear β decay, if we write the differential decay rate of the neutron as a function of the neutron spin vector σ , electron momentum vector p_e , and neutrino momentum vector p_ν , it can be expressed in terms of various scalar products with different symmetry properties[96, 18]:

$$d\Gamma = \Gamma_0 p_e E_e (E^0 - E_e)^2 dE_e d\Omega_e d\Omega_\nu \times \left[1 + a \frac{p_e \cdot p_\nu}{E_e E_\nu} + b \frac{m_e}{E_e} + 2\sigma \cdot \left(A \frac{p_e}{E_e} + B \frac{p_\nu}{E_\nu} + D \frac{p_e \times p_\nu}{E_e E_\nu} \right) \right] \quad (1.20)$$

with E_e/E_ν the electron/neutrino energies, Ω_e/Ω_ν the electron/neutrino momentum solid angles (with respect to σ), m_e the electron mass, and the overall factor Γ_0 a function of the coupling constants in the Hamiltonian. The coefficients a , b , A , B , and D represent correlations between the various three-vectors in the decay process, and can be expressed as functions of the (in general complex) ratio $\lambda = |g_A|/|g_V|e^{i\phi}$. For example, the coefficient A describes the correlation between the direction of the neutron spin and the electron momentum, an observable which violates parity symmetry; this correlation coefficient is in fact non-zero, which is expected due to manifest parity violation in the underlying theory. One can also write a differential decay rate as a function of the electron polarization and summing over the neutron spins, which gives several additional correlation coefficients.

The total decay rate Γ of the neutron can be computed by summing over neutron polarizations

and integrating over the decay phase space, which gives the lifetime $\tau_n = \Gamma^{-1}$ [121]:

$$\tau_n^{-1} = \frac{m_e^5}{2\pi^3} G_F^2 V_{ud}^2 [g_V^2(1 + \Delta_V^R) + 3g_A^2(1 + \Delta_A^R)] F_n. \quad (1.21)$$

The factor F_n is a calculated phase space factor for neutron decay (including weak magnetism and nuclear recoil effects)[128]. The corrective terms $\Delta_{V/A}^R$ are radiative corrections due to Bremsstrahlung of the final state charged particles and corrections to the weak vertex. Including these corrections, as well as the comparatively well known value of G_F from muon decay, the neutron lifetime can be written as

$$\tau_n^{-1} = \frac{V_{ud}^2(1 + 3\lambda^2)}{4908.7(1.9) \text{ s}} \quad (1.22)$$

where the uncertainty in the numerical factor is primarily due to hadronic uncertainties in the radiative corrections[107].

In all, there are more than a dozen such observables in neutron decay, which, when combined with experimental knowledge of the neutron lifetime, greatly over-constrain the number of parameters in the effective Hamiltonian. This makes measurements of neutron decay observables a powerful tool for investigating new physics which could cause small deviations SM prediction of the observables discussed here.

1.4 The Predicted Neutron Lifetime

The most precise determination of V_{ud} comes from studies of the β -decay lifetime of nuclei. Nuclear decays with spin and parity quantum numbers J^P of 0^+ in both the initial and final states only depend on the vector coupling, which is the same for different nuclei under the assumption of the CVC hypothesis. By incorporating isospin-breaking corrections, nucleus dependent radiative corrections, and nuclear structure corrections, one can relate the mean lifetimes of such nuclei to V_{ud} . This has been performed with satisfactory agreement across several nuclei, from which V_{ud} is found to be $|V_{ud}| = 0.97425 \pm 0.00022$ [90]. Quark flavor changing pion decay (in particular

$\pi^+ \rightarrow \pi^0 + e^+ + \nu_e$) also determines V_{ud} , but with somewhat larger uncertainty to date[38].

We are thus left with the contribution λ of the axial coupling to the predicted neutron lifetime. This can in principle be determined by lattice QCD calculations, though precision of these calculations is currently not competitive with other methods[1]. Currently, measurements of the A coefficient (where $A = -2\lambda(\lambda + 1)/(1 + 3\lambda^2)$) provide the most precise determination of λ . Generally speaking, measurements of A consist of detecting β particles emitted from a sample of polarized neutrons, and comparing the number of β particles emitted parallel and anti-parallel to the neutron spin[46, 32]. The particle data group average of measurements gives $A = -0.1184 \pm 0.0010$, from which we have $\lambda = -1.2723 \pm 0.0023$ [47].




Figure 1.3: The combined determination of V_{ud} and λ from $0^+ \rightarrow 0^+$ nuclear decays, the β -asymmetry parameter A , and the neutron lifetime τ_n .

Using these experimental results and eqn. 1.22 the predicted neutron lifetime is $\tau_n = 883.1 \pm 0.8$ s. This value disagrees with the experimental global average of $\tau_n = 880.3 \pm 1.1$ (see fig. 1.3). Moreover, there is a statistically significant spread in different measurements of τ_n , and some measurements have recently been re-evaluated, which has caused a $\sim 5\sigma$ shift in the mean value[27, 23, 28, 69]. For these reasons, the discrepancy with the SM prediction is thought to be

due to underestimated (or unconsidered) systematic effects[134]. One might therefore take the experimental uncertainty in τ_n to be somewhat larger than the uncertainty given in the global average, which is a primary motivation for an improved measurement of τ_n .

1.5 The Impact of an Improved τ_n Measurement

1.5.1 CKM Unitarity

The CKM matrix given in eqn. 1.10 is unitary, and as a consequence, the sum of the square modulus of any row or column must equal one. Taking the first row, it must be that

$$|V_{ud}|^2 + |V_{us}|^2 + |V_{ub}|^2 = 1. \quad (1.23)$$

A deviation of this sum from unity is evidence that either the SM does not adequately describe all possible quark transitions, or the Fermi coupling G_F that governs the overall strength of charged current weak interactions in the SM is in fact not universal; any such deviation would thus be evidence of beyond-SM physics, such as a fourth generation of quarks.

The deviation Δ_{ckm} from unitarity can be defined as

$$\Delta_{ckm} = |V_{ud}|^2 + |V_{us}|^2 + |V_{ub}|^2 - 1 \quad (1.24)$$

The second and third entries in the matrix have been experimentally determined through the study of kaon and B meson decays[55], which puts a combined constraint on unitarity violation $\Delta_{ckm} = (1 \pm 6) \times 10^{-4}$ [128]. This constrains new physics at an energy scale competitive with collider experiments[13, 14]. The V_{ud} entry could be extracted more precisely from neutron decay (which obviates the need for nuclear structure corrections as in $0^+ \rightarrow 0^+$ decays) with improved determinations of λ and a reliable determination of τ_n at the sub-second level, and is therefore an appealing model-independent way of constraining new physics.

1.5.2 Tests of V-A Theory

The effective Hamiltonian of eqn. 1.16 includes all of the non-negligible contributions to the lepton and nucleon currents predicted by the underlying V-A interaction of the SM. However, beyond-SM theories can in general predict small (but detectable) scalar (S) or tensor (T) currents which would modify the Hamiltonian, and thus modify the neutron decay rate. For example, theories with spontaneous left/right SU(2) symmetry breaking, theories with leptoquarks (i.e. particles with both lepton and baryon number), and supersymmetric theories can introduce small S and T currents in neutron and nuclear β -decay[18, 121]. Within the SM, there are in principle S/T operators induced in eqn. 1.16 by (for example) loop diagrams involving the Higgs boson, but such contributions are expected to be many orders of magnitude below what is currently observable.

The strengths of these beyond-SM interactions can be constrained using measurements of the observables from neutron and nuclear decays, and these searches are complimentary to beyond-SM interactions that can be probed in high energy experiments[13]. Fits of the relative strength of all possible effective form factors, making various model assumptions, have been performed using available data of total decay rates and decay correlations[121]. In order to translate these constraints into constraints on the true couplings of beyond-SM theories, one must estimate or make assumptions about the induced form factors for the neutron due to the new couplings. For example, the authors of ref. [70] discuss the prospects for lattice QCD calculations of the induced scalar charge, and explore its impact on beyond-SM constraints. Further, stringent limits on tensor form factors can be derived by combining $0^+ \rightarrow 0^+$ decays, the neutron lifetime, and angular correlations[97].

1.5.3 Big Bang Nucleosynthesis

The neutron lifetime plays a direct role in the predicted abundance of light nuclei in the early universe. After $t \sim 30 \mu\text{s}$, a small fraction of protons and neutrons interacted with each other via $n + \nu_e \leftrightarrow p + e^-$, $n + e^+ \leftrightarrow p + \bar{\nu}_e$, and neutron decay. Eventually (after $t \sim 1 \text{ s}$) the strength of

the weak interaction was overcome by the expansion and cooling of the universe, and thereafter only neutron decay contributed to the change in the relative number of neutrons and protons. At $t \sim 3$ minutes, the universe had cooled enough for the nucleons to form light, stable nuclei as a result of fusion. The resulting abundance of these nuclei (in particular ${}^4\text{He}$) thus depends on the neutron lifetime[18], and a comparison of the observed and predicted helium abundance is a test of cosmological models.

As discussed in section 1.4, recent measurements of τ_n are highly discrepant. The impact of discrepant lifetime measurements on primordial nucleosynthesis was explored in ref. [108]. Forthcoming astronomical measurements of light element abundances and measurements of the baryon to photon ratio will ultimately make τ_n the least well known experimental input, and new neutron lifetime measurements will therefore be needed to test predictions from cosmology.

1.6 Conclusions

Neutron decay offers several experimental observables which provide a competitive test of electroweak theory. More precise measurements of neutron decay parameters, including the neutron lifetime, are motivated by their potential to test for new physics at energy scales similar to current collider experiments. Further, primordial nucleosynthesis (as well as other charged-current processes) benefits from the resolution of τ_n at ~ 1 s precision.

The ambiguity in the current global data weakens the ability of neutron decay to discover discrepancies in the standard model or cosmology. Therefore, new methods for measuring the neutron lifetime, a careful study of potential issues with past experiments, and focus on experimental investigations of systematic effects are all needed for neutron decay to reach its full potential as a scientific tool.

2 The History of τ_n

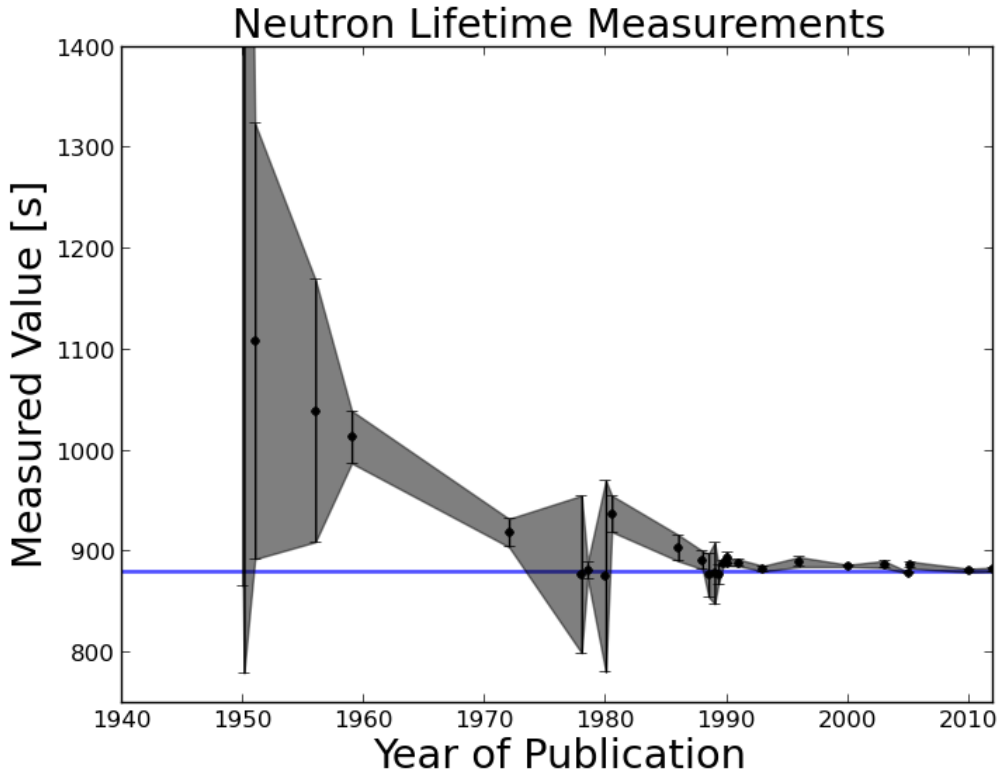
2.1 Introduction

The neutron was discovered in 1932 by Chadwick[9, 11] by bombarding a beryllium target with α -rays from polonium, inducing the reaction ${}^9\text{Be}(\alpha, n){}^{12}\text{C}$ [10]. The instability of the neutron was suggested a few years later by a determination of the neutron mass from the deuteron binding energy[12]. However, it was not until a decade thereafter, with the development of the first continuously operating nuclear reactor, that the decay of the neutron was observed by Snell *et al.*[122]. Since this time, roughly two dozen measurements of the neutron lifetime τ_n have been performed.

The precision of τ_n has improved steadily over this time, though not without inconsistency and skepticism. In fact, Snell *et al.*'s estimate placed a *lower* limit of 22 minutes on the neutron lifetime – 7 minutes longer than the current accepted value. This was perhaps a presage of the future τ_n narrative. Figure 2.1 shows experimental determinations of the lifetime over time. Several of these measurements were subsequently re-evaluated or withdrawn[47, 134].

Early measurements of the neutron lifetime used an “in-beam” method of measuring the neutron lifetime. Generally, these measurements consisted of passing a cold or thermal neutron beam through a charged particle detector (β or proton, or both). If the absolute beam density, charged particle detection efficiency, and decay volume can be determined, the neutron lifetime can be extracted.

The neutron flux in such an experiment is typically measured by activating a foil of known absorption cross section and density. For a neutron beam with flux spectrum $d\phi/dv$, the measured rate depends on not the flux, but the density $d\rho/dv = v^{-1}d\phi/dv$ of neutrons (integrated over

Figure 2.1: The mean neutron lifetime τ_n versus time.

the detection volume and neutron velocities). Conveniently, the activation rate R_f of a foil is given by

$$R_f = \int \sigma_a(v) \frac{d\phi}{dv} dv = \sigma_{th} v_{th} \int \frac{d\rho}{dv} dv = \sigma_{th} v_{th} \rho_n. \quad (2.1)$$

Thus, for a thin foil where the $1/v$ absorption cross section dependence is valid, the beam density is given directly by the activity measurement. We can write the neutron fluence in terms of the above quantities

$$R_n = \epsilon_n A \rho_{foil} \sigma_{th} v_{th} \rho_n \quad (2.2)$$

where ϵ_n is the efficiency of measuring the neutron activation of the foil, A is the beam area, and ρ_{foil} is the areal density of the foil. The charged particle detection rate is given by

$$R_p = \frac{\epsilon_p \rho_n A L}{\tau_n} \quad (2.3)$$

where ϵ_p is the charged particle (typically proton) detection efficiency L is the length of beam from which charged particles can be detected. We can thus determine τ_n :

$$\tau_n = \frac{\epsilon_p R_n L}{\epsilon_n R_p v_{th} \sigma_{th} \rho_{foil}}. \quad (2.4)$$

From this we see that a precise measurement of the neutron lifetime requires accurate metrology to determine the effective length L , and to account for geometric considerations in determining the efficiencies. Detector considerations such as back-scattering from windows, thresholds, gain drifts, and backgrounds can also make the determination of the efficiencies difficult. That said, not all experiments fall completely into this paradigm, and have in some cases found elegant ways of mitigating or bypassing some of these difficulties.

Basically all other measurements of the neutron lifetime use ultracold neutrons, or UCN (see appendix A for discussion). If UCN are loaded into a suitable trap (typically a material trap), the number of surviving neutrons can be measured after various storage times, from which the neutron lifetime can be determined. As long as the detection efficiency for surviving UCN is the same for all storage times, one need only perform a relative measurement of the UCN at different times, thus avoiding the need to determine absolute efficiencies as in beam-based experiments. However, the stored UCN may also be absorbed within the trap or inelastically scatter out of the trap; that is, the storage time τ_s of the trap is given by

$$\tau_s^{-1} = \tau_n^{-1} + \tau_{loss}^{-1} \quad (2.5)$$

where τ_{loss}^{-1} is the rate at which neutrons are lost due to, for example, interactions with the walls of the trap, or with residual gas in the apparatus. A trap-based measurement of the neutron lifetime must characterize these loss mechanisms, and extrapolate to $\tau_{loss}^{-1} = 0$. Further, these loss mechanisms may depend on the given energy of a UCN (and more generally the phase space distribution of the UCN) which leads to many subtle difficulties and many methods to characterize these effects. Insofar as equilibrium kinetic theory is valid for UCN, one can compute the wall

collision rate $\gamma(E)$ which depends on the UCN energy as well as the bottle geometry. With this we can write eqn. 2.5 as

$$\tau_s^{-1} = \tau_n^{-1} + \eta\gamma(E) \quad (2.6)$$

where η is the probability of loss-per-bounce. One can perform successive storage measurements varying η , E , or γ , and using the calculated scaling of the loss term to extrapolate to zero loss. In practice, varying the trap geometry (and hence γ) is common, and is often known as *dimensional extrapolation*.

In addition to the scientific interest in a precise determination of τ_n , the criticism of past experimental methods and analyses motivates new experiments. It is therefore worthwhile to review the landscape of neutron lifetime measurements; this provides the context for the experimental design and theoretical considerations that will be discussed in later chapters.

2.2 First Measurements

The first measurements by Snell consisted of passing the neutron beam through a thin-walled vacuum can, with a cylindrical electrode (4 kV) which deflected the decay protons into an electron multiplier. Proportional counters to detect the β particles in coincidence were implemented to reduce the background, and a series of measurements were performed with a boron shutter to stop the neutron beam, as well as aluminum foils to shutter the β particles and protons. A hydrogen leak was introduced and the resulting rate in the proton detector for comparison. This demonstrated neutron decay, but the authors only estimated a lower bound on the lifetime[122, 123].

At the same time, Robson performed a similar measurement at Chalk River Laboratories. There were some important differences: the protons were detected by a magnetic spectrometer, and the apparatus also included a β spectrometer, which demonstrated that the decay product was a proton, and determined the β endpoint energy to be 782 ± 13 keV. In addition, the intensity of the neutron beam is measured by activating manganese foils, and the author notes

the convenience of a thin foil with $\sigma_a \propto 1/v$ for measuring the beam density. The measurement, combined with the decay rate (and an estimate of the absolute detection efficiencies) leads to an estimated neutron lifetime of 1108 ± 216 s[114, 115].

Also in the 1950's, Spivak *et al.* performed similar measurements at the Atomic Energy in the USSR. They note that the spatial dependence of the electric field profile used to extract the decay protons can lead to uncertainties in determining the effective beam length L , and the authors use favorable electric field geometries to mitigate this effect. Sodium and gold samples are activated to determine the neutron flux, from which the neutron lifetime is determined to be 1012 ± 26 s[59, 21].

A somewhat different technique was used by D'Angelo at Argonne National Laboratory. A neutron beam is passed through a cloud chamber, and decay events photographed at a rate of ~ 2 frames per second with a stereoscopic camera system. Gold foils are used to characterize the beam density, with an uncertainty of $\sim 7\%$. While the detection efficiency for the cloud chamber is well understood, the experiment suffers from environmental, neutron beam generated, and neutron capture generated γ -ray backgrounds. Materials are chosen carefully in the construction of the apparatus to mitigate neutron related backgrounds, and considerable lead shielding is employed. The result is a neutron lifetime of 1099 ± 164 s, with a statistical uncertainty contributing to much of the total uncertainty[16].

2.3 Improved In-Beam Measurements

At the Risö reactor, Christensen *et al.* used a technique which greatly reduced the uncertainty in determining the effective beam length. A pair of scintillator paddles were immersed in a uniform, 0.7 T magnetic field perpendicular to the neutron beam. The magnetic field served to guide the decay electrons to the scintillators, so that ideally any decay event between the scintillators would be detected. This, combined with a 0.4% measurement of the neutron density using a calibrated ^3He -based neutron counter (cross-checked with a gold foil activation measurement), gave a lifetime of 918 ± 14 s. The authors discuss the effects of the gyroradius of the electrons

(which would permit some fraction of electrons near the edge of the decay volume to miss the scintillators), as well as the reduction in detection efficiency due to magnetic mirror effects and detector thresholds, and these effects contribute substantially to the uncertainty[34].

As discussed in the previous section, proton counting is made difficult due to electric field profiles needed to accelerate them into a detector, which can make the effective determination of L difficult. Bondarenko *et al.* addressed this by utilizing a decay volume that is free of applied voltage[54]. There is an aperture adjacent to the decay volume which contains the focusing and acceleration electrodes. In this way, the protons that reach the aperture are collected with near 100% efficiency, which was checked using H^+ and α sources, and the effective decay length is easily determined because it is field-free. The solid angle collection efficiency into the aperture is determined in a Monte Carlo study, including model uncertainties. Gold foil activation is used to determine the beam flux, and backgrounds are investigated with a Cd shutter for the neutron beam, and electrostatic mirror in front of the proton detector. The authors find a lifetime of 877 ± 8 s. A later publication re-analyzed this experiment[124], considering corrections to the neutron flux determination such as scattering from the gold foils and non- $1/v$ behavior. The authors perform separate experiments investigating the transparency of the high voltage grids used for proton collection, ultimately finding a corrected value for the neutron lifetime of $\tau_n = 891 \pm 9$ s.

One method eliminating the need to determine L was investigated in the PERKEO experiment. The apparatus consists of a solenoidal magnetic field (parallel to a neutron beam) which bends upwards at either end towards scintillator paddles to detect the decay β s. A careful determination of the threshold efficiency and source calibrations were performed. To reduce backgrounds, two PMTs viewed one scintillator, and a coincident signal was required in both PMTs to reduce backgrounds. The unique feature of the experiment was that the neutron beam was *pulsed*, and the pulse bunches were ~ 1.5 m in length, whereas the decay volume was about 20 cm larger. For this reason, knowledge of the decay length L was unnecessary because the signal could be measured when the bunch was completely in the decay volume. Low statistical sensitivity

was a side effect of this method, contributing a 10 s uncertainty to the measurement of $\tau_n = 876 \pm 21$ s[51]. Gain drifts, detector resolution and calibration, and neutron beam characterization contributed to the systematic uncertainties.

Significant technical improvements were made by Byrne *et al.* starting in 1980[48]. The authors counted protons using a penning trap (a 1 T solenoidal magnetic field capped by 1 kV mirror potentials) placed perpendicular to the neutron beam. The neutron beam was passed through the trap, and protons were allowed to collect in the trap. The neutron beam was thereafter shuttered, and the trap opened to view a silicon surface barrier detector biased to -30 kV. Because the neutron beam was not passing through the apparatus while protons were counted, backgrounds were greatly reduced, and the surface barrier detector provided excellent signal-to-noise. In addition, the neutron beam was actively monitored using a ^{10}B foil. The reaction $^{10}\text{B}(n,\alpha)^7\text{Li}$ produces energetic α particles which were counted by an arrangement of surface barrier detectors. The final result is $\tau_n = 937 \pm 18$ s.

A revised version of the experiment reoriented the penning trap parallel to the neutron beam, and an additional improvement was introduced which greatly reduced the systematic effect of determining L . The apparatus consisted of a 5 T solenoidal magnetic field. There were a series of sixteen electrodes along the length of the solenoid magnet. For a given measurement, different electrodes could be used as the electrostatic caps (1 kV). The length of the penning trap could thus be varied with all other experimental parameters being equal. In this way, edge effects are eliminated via linear extrapolation to $L^{-1} \rightarrow 0$. The authors also demonstrate long (100 ms) trapping times for the protons to check for sources of loss. The beam is characterized with a ^{10}B foil as in the previous experiment, for which the $1/v$ law is valid to 0.03%. The largest systematic uncertainty was in determining the areal density of the foil (0.3%), and there was a -3.6 ± 0.5 s correction due to non-uniformity of the foil. Other sources of systematic uncertainty were knowledge of the ^{10}B cross section, uncertainty in proton detection efficiency, and scattering from the boron foil substrate. The authors found $\tau_n = 893.6 \pm 5.3$ s[49].

This result was corrected by a later analysis. A Monte Carlo study of the trapped protons

was performed for different electrode configurations, including inhomogeneities in the magnetic field. These inhomogeneities caused slight deviations from linearity, with short trap lengths being slightly elongated, and long trap lengths being slightly shortened. This introduced a -4.4 s correction in τ_n . In addition, the authors find an error in proton detector deadtime corrections used in previous work. The corrected result is found to be 889.2 ± 4.8 s[50].

The basic technique of Byrne *et al.*[49] was used in a more recent experiment at the National Institute of Standards and Technology (NIST). As in the previous experiment, a proton penning trap (4.6 T, 800 V caps) with sixteen variable electrodes were used. The same linear length extrapolation was performed, with corrections from a Monte Carlo study.

The authors investigate proton loss and back-scattering in the dead layer of the silicon surface barrier detector. Estimates were performed using SRIM, and the bias potential and dead layer thickness were varied. The authors compute the lifetime extrapolating to 0 deadlayer thickness.

The neutron beam was incident upon a ${}^6\text{LiF}$ foil on a silicon substrate. Silicon surface barrier detectors with precision-machined apertures viewed the foil, and the α particles counted to determine the beam density. The solid angle efficiency for this detector array was determined to within 0.1% using a calibrated α source, as well as contact metrology. The determination of the areal density of the foil contributed a 2.2 s uncertainty, and knowledge of the ${}^6\text{Li}(n,t)\alpha$ cross section contributed a 1.2 uncertainty. Effects such as neutron beam divergence, finite LiF foil thickness, proton trap non-linearity, and loss in the LiF foil substrate contributed 1 to 5 s corrections with uncertainties ranging from 0.1 to 1 s. Proton counting dominated the statistical uncertainty, and the authors find $\tau_n = 886.8 \pm 1.2_{\text{stat}} \pm 3.2_{\text{sys}}$ s[58, 52].

More recently, Yue *et al.* have performed separate experiments to calibrate the LiF-based neutron monitor. A neutron detector was used to measure the efficiency of the neutron monitor used in the original experiment to greater precision. The stability of the LiF foils over time was established, so that the improved determination of the efficiency can be used in the determination of the lifetime from the 2005 data, and the result does not depend on the absolute value of the Li absorption cross section. From this analysis, the authors find $\tau_n = 887.7 \pm 1.2_{\text{stat}} \pm 1.9_{\text{sys}}$

s[31].

This experiment is the most precise determination of τ_n using a neutron beam. It should be noted that because the neutron density measurement contributed the most to the uncertainty, multiple independent determinations of the beam density could increase the precision of the measurement. With this motivation, a continuation of the NIST beam is planned, which aims to measure τ_n with a sub-second total uncertainty.

2.4 The Material Bottle Method

The first determination of τ_n using UCN was performed by Kosvintsev *et al.* at the SM-2 reactor[80, 81]. An upright, cylindrical aluminum container was used to store UCN for various storage times. A series of aluminum discs could be inserted into the trap in order to vary the total surface area, and thus the loss rate τ_{loss}^{-1} . The authors measure the storage time for different numbers of discs, and use this to extrapolate to $\tau_{loss}^{-1} \rightarrow 0$. The authors find $\tau_n = 903 \pm 13$ s, with systematic uncertainties due to UCN leakage through shutters, loss due to residual gas, uncertainty of the measured UCN spectrum (which is needed to average the loss rate over different UCN energies), and the presence of Al_2O_3 on the surface of the trap. A preliminary measurement with a D_2O coating on the walls gave $\tau_n = 892 \pm 20$ s, but no subsequent measurements were reported.

An experiment by Alfimenkov *et al.* used a spherical trap with a single opening. The trap could be rotated so that the opening could face a UCN guide to load the UCN, and could then subsequently be rotated so that the opening was at different heights. Different UCN energy bins could therefore be selected by rotating the trap. The aluminum trap walls were coated with a 0.3 to 0.5 μm thick Be layer, and a 3 to 7 μm layer of solid oxygen at 15 K. The quoted losses per bounce are $(28.1 \pm 4) \times 10^{-6}$ and $(6.1 \pm 0.6) \times 10^{-6}$, respectively. This lead to storage times of 8 to 10 hours (excluding β -decay). A second measurement campaign used a cylindrical trap (72 cm diameter by 15 cm long). This geometry has a different calculated UCN-wall collision rate, which provides a means of extrapolating to zero collision rate, and hence zero loss.

The experiment is explained in detail in ref. [74]. The authors discuss tests of the quality of the surface coatings, the effect of residual gas, as well as the effect of partial overlap of the selected energy bins. The effect of UCN heating due to the movement of the trap, uncertainties in the optical potential, and uncertainty in the knowledge of the UCN spectrum within the energy bins. The final result is $\tau_n = 888.4 \pm 3.1_{stat} \pm 1.1_{sys}$.

Several material trap-based measurements have used Fomblin oil or other fluorocarbon oil, with a quoted loss-per-bounce $\eta = (2 - 3) \times 10^{-5}$ at 20° , and an optical potential $V_O = 106.5$ neV[75]. In addition to a reasonable optical potential and low loss-per-bounce, Fomblin has a low vapor pressure, and Fomblin coatings can be replenished *in situ* to provide a uniform surface.

The first experiment to use Fomblin was performed by Mampe *et al.* at the Institut Laue Langevin[75]. The apparatus consists of a rectangular glass-walled $30 \times 40 \times x$ cm³ volume, where x can be varied. Part of the roof is corrugated to assure that the loaded UCN are distributed isotropically on a short time-scale (the authors quote a time of a few seconds). The general measurement technique is to vary x so as to vary the mean free path for wall interactions $\lambda = 4V/S$ in a way that is proportional to the difference in times for which UCN are trapped. The loss (due to changing the surface area) is thus varied while maintaining the same average number of bounces during storage for each choice of x . In this way, changes in the UCN spectrum (due to the energy dependence of the loss-per-bounce) have the same effect for all measurements, and an extrapolation $\lambda^{-1} \rightarrow 0$ can be performed to eliminate the effect of wall loss. There is a $\sim 0.6\%$ correction due to the fact that gravity causes the collision rate with the ceiling to be lower than that of the floor, somewhat spoiling the assumption that the UCN sample the whole surface evenly. There is also a correction ($+0.3\%$) due to differences in initial UCN spectra for different x , which can induce a non-linearity in the storage time versus λ^{-1} .

The authors also discussed the direct assessment of potential systematic effects. The vacuum pumps were changed to increase the effect of microphonics, which could cause slightly inelastic collisions with the walls, thus “evaporating” some small amount of UCN from the trap, and no reduction in the storage time was observed here. The authors also measure the lifetime while

changing the temperature of the walls, which changes the probability of upscattering per bounce, and find that the $\lambda^{-1} \rightarrow 0$ extrapolation is reliable at all temperatures chosen. In order to investigate the effect of residual gas on the storage time, the measurement was performed at somewhat higher vacuum pressure, which also caused no noticeable effect.

Finally, the authors replaced the corrugated surface with a smooth glass surface, and repeated the measurement. A reduction in the storage time constant for holding times less than 450 s was observed, but for holding times $t > 900$ s, typical storage times were again observed. The authors hypothesize that the completely smooth surface compromises the validity of the expression for λ . Generally speaking, this suggests that the UCN in the completely smooth trap do not rapidly fill their accessible phase space – this can lead to the general failure of kinetic theory as a tool for understanding UCN trap-based neutron lifetime measurements, and can cause there to be UCN with $E > V_O$ which persistently remain in the trap for long times (we hereafter refer to this type of UCN as *quasi-bound*). The final result is the weighted average of different measurements, and gives $\tau_n = 887.6 \pm 3$ s.

A general feature of the above experiments is that they rely on the calculation of the wall collision rate for different geometries or energies. This lead Mampe *et al.* to propose a means of measuring a quantity that depends on the collision rate. A fomblin-coated trap is used, and the quasi-bound UCN are removed by lowering a polyethylene cleaner into the trap at the beginning of storage cycles. An array of ^3He -filled drift tubes is arranged outside of the trap to measure UCN that are lost due to upscattering, thus providing a measurement of the collision rate of UCN with the walls. Storage measurements are performed with two different trap geometries, and at three different temperatures (to vary the upscattering rate). The neutron lifetime is taken as a combination of the measurements at each temperature, correcting for the collision rates for the two geometries. There are systematic effects due to leakage through the valve that closes the trap and temperature gradients on the walls of the trap. In addition, there are systematic effects related to the validity of the analysis and measurement of the collision rate. In particular, the upscattering detection efficiency may be different for the two geometries, which can spoil the

determination of the collision rate and make the extrapolation to zero loss less reliable. Including estimates and corrections of all effects, the authors' final result is $\tau_n = 882.6 \pm 2.7$ s[76].

2.5 Improved Bottle Measurements

The first measurement of τ_n with a quoted total uncertainty of less than one second was an improved version of the experiment in ref. [76]. The trap consisted of a cylindrical, double-walled, Fomblin-coated vessel in a horizontal orientation (inner cylinder 90 cm long by 33 in diameter, outer cylinder larger by a 2.5 cm margin on all sides). A shutter connected the inner vessel to the guide, and there is another shutter to connect the inner and outer vessels. The vessels could be rotated in order to dip them into a Fomblin bath to re-coat the surfaces, and the temperature of the walls could be varied between -26 and 20 Celsius and pumped to a vacuum of $(1 - 5) \times 10^{-6}$ Torr. As in ref. [76], an array of ^3He -filled drift tubes was arranged outside the vessels in order to detect upscattered UCN.

The storage times were measured at different temperatures over the course of 100 days. The analysis of the data incorporated a small linear drift in the observed decay rate to incorporate (to first order) the effect of the UCN spectrum evolving in the trap versus time. Spectrum and time averaged quantities are computed in this model to extract τ_n . The authors correct for different emptying times of the trap, which introduces a -3.1 ± 0.4 s correction. The difference in thermal neutron efficiency for the two vessels is corrected using Monte Carlo studies and separate experiments. The linear drift was estimated by measuring the time dependence of the thermal neutron count rates, which introduced a -2.0 ± 0.3 s correction. In addition, a 0.2 s uncertainty was included to account for loss due to residual gas, a 0.3 s uncertainty for the thermal neutron detection efficiency, a 0.2 s uncertainty in characterizing the number of UCN with $E > V_O$ for Fomblin, and a 0.15 s uncertainty due to temperature gradients across the vessels. The end result is $\tau_n = 885.4 \pm 0.9_{stat} \pm 0.4_{sys}$ s.

The measurement of τ_n with the lowest quoted uncertainty used a method similar to ref. [2]. Two rotatable traps were used – a cylinder (76 cm diameter by 14 cm long), and a quasi-spherical

trap composed of two truncated cones connected with a cylindrical section. The cylindrical trap had a calculated collision rate 2.5 times higher than the quasi-spherical trap.

The traps were coated with low temperature Fomblin which was evaporated onto the trap surface, and could be replenished *in situ* with an evaporator on top of the apparatus. To test the coating quality, storage measurements were performed also using a titanium trap coated with the oil; titanium has a negative optical potential, so a long storage time shows that the trap was completely covered with the oil, leaving no exposed regions of low potential. The traps used for the experiment were beryllium. The trap was cooled to 113 Kelvin, and the estimated loss-per-bounce of the trap surface was 2×10^{-6} . After filling the trap, it is rotated to remove UCN at the high-energy end of the spectrum, at which point it is brought upright. The neutrons are then stored for either 300 or 2000 s. The trap is emptied by successively tipping the trap to select different bands of the UCN spectrum. This permits dimensional extrapolation (using the two different trap shapes) as well as energy extrapolation of the calculated loss factor $\gamma(E) \rightarrow 0$.

A Monte Carlo of the experiment was performed to reproduce the data, with the probability of diffusive scattering in the trap left as a free parameter. Diffusive scattering probabilities of 10% and 1% were chosen, from which the authors conclude that the true probability is at least 10%. The simulation reproduces the experiment to within 0.236 s, which is dominated by the statistics of the simulation. In order to estimate the effect of loss due to residual gas, the trap pressure was increased to 8×10^{-4} mbar, and the storage time at two pressures used to extrapolate to zero pressure. The loss-per-bounce versus energy is varied within reasonable limits from which the authors find a 0.144 s uncertainty. The quoted uncertainties due to uncertainty in the Fomblin potential is 0.004 s, and 0.058 s uncertainty in the dimensions of the traps. The uncertainties in the UCN spectrum determination amounted to 0.104 s, for a final result of $878.5 \pm 0.7_{stat} \pm 0.3_{sys}$ s. At the time of publication, this result was 5.6σ lower than the next most precise result (that of Arzumanof *et al.*), and 6.5σ lower than the world average. The authors note that their result in combination with measurements of λ is in better agreement with the standard model. In addition, the authors retract the result of ref. [2], citing poor statistical sensitivity and coverage

and uniformity of the oxygen coating[27].

More recently, a measurement was performed with a technique similar to ref. [75]. As before, a rectangular, variable-volume, Fomblin-coated, glass-walled trap is used to store UCN ($43 \times 42 \times (0 \text{ to } 75) \text{ cm}^3$). The time-scaling method is used, so that each storage time (for each trap volume) gives the same number of UCN-wall collisions. The trap incorporates an outer “pre-storage” volume with a polyethylene roof to rapidly upscatter any quasi-bound UCN before storage in the primary volume. The fomblin coating can be replenished through the piston that moves the trap wall. The trap vacuum was 8×10^{-6} mbar or better, and the vacuum was constantly monitored with a mass spectrometer.

The authors combine results from measurements with different wall temperatures. There are contributions to the uncertainty from the determination of the trap length, loss due to residual gas, model uncertainty in the loss-per-bounce as a function of UCN energy, temperature gradients, and variations in the emptying time of the UCN from the trap. The authors find $\tau_n = 880.7 \pm 1.3_{stat} \pm 1.2_{sys} \text{ s}$ [23].

2.6 Corrections and Criticisms of Bottle Experiments

As of the publication of Serebrov, *et al.*[27], the two most precise measurements of the neutron lifetime were in significant disagreement, and this has lead to discussion of potential sources of new or under-estimated systematic effects in these material bottle experiments. Further, some of the experiments in the previous section were re-analyzed in light of such systematic effects, and new determinations of τ_n have been deduced from the data. In particular, the results of Arzumanov *et al*[66] and Mampe *et al.*[75] have been re-evaluated, causing a substantial downward shift in the global average of τ_n . Here, we will review some of the discussion of these effects and the re-evaluations of these experiments.

Serebrov and Fomin[120] argue that the work of Serebrov *et al.*[27] exhibited the lowest loss probability of any material bottle experiment, and it is therefore reasonable to expect that systematic corrections in that experiment cannot explain the discrepancy with other neutron

lifetime measurements (in particular that of Arzumanov et al[66]). A Monte Carlo study of the experiments in refs. [66, 75] is therefore carried out to address potential systematic effects not accounted for in the original work. The authors use a simulation which computes UCN trajectories, incorporating loss, reflection, and diffusive scattering from material surfaces, with the diffusive scattering probabilities were left as a tunable parameter.

For the experiment in Arzumanov et al[66], this simulation is used to investigate the effects of the heating of UCN by shutters, leading to a -2.8 ± 2 s correction, and the effect of UCN not being completely emptied from the inner vessel of the trap (which was not observed in the simulation, giving a 0 ± 1 s effect). In addition, the authors argue that the original correction due to unequal thermal neutron detection efficiencies for the inner and outer volumes is incorrect. A -2.1 ± 1 s correction is given, for a corrected neutron lifetime of $879.9 \pm 0.9_{stat} \pm 2.4_{sys}$ s.

This experiment was later re-analyzed by the original experimenters in ref. [69]. The authors retract their original positive correction due to thermal neutron detection efficiency, and use measurements to estimate a -1.6 ± 1.86 s correction instead. In addition, they argue that differences in the stored UCN spectra (and possibly valve-related heating) in the inner and outer vessels leads to different UCN detection efficiencies for the two volumes, giving a -3.1 ± 0.3 s correction not included in the original work. This, and with a minor correction due to revised data reduction, gives a new value $881.6 \pm 0.8_{stat} \pm 1.9_{sys}$ s which replaces the original result.

Serebrov and Fomin (also in ref. [120]) study the effect of quasi-elastic scattering of UCN from Fomblin-coated surfaces to reanalyze the work of Mampe et al[75]. The authors argue a -6.0 ± 1.6 s correction due primarily to the creation of quasi-bound UCN for long holding times as a consequence of quasi-elastic scattering, and provide a revised result $\tau_n = 881.6$ s with the original quoted uncertainty. Serebrov and Fomin, from their corrected values of these other experiments, present a revised global average of 879.9 ± 0.9 s, arguing that this lower value of τ_n is in good agreement with the standard model and with primordial nucleosynthesis predictions.

Steyerl *et al.* responded to the work of Serebrov and Fomin in ref. [28], and provided a more robust model of quasi-elastic UCN scattering. This model is used in a Monte Carlo study of the

experiment in Mampe et al.[75]. Steyerl *et al.* note minor discrepancies in the originally presented work which are incorporated into the revised analysis. A negative correction to the original result is found, though not as severe as that suggested by Serebrov and Fomin. This is in part due to the presence of UCN with $E > V_O$ which lose energy due to quasi-elastic scattering, thus becoming trapped. They recommend a new value $\tau_n = 882.5 \pm 2.1$ s, replacing the original result.

In a separate work, Steyerl *et al.* present a model of surface roughness relevant to Fomblin-coated surfaces, and compute the UCN-surface scattering intensity to second order in perturbation theory. This results in a reduced wall-loss probability for UCN, and the result is applied to a more realistic Monte Carlo simulation of the work of Serebrov et al.[27]. The authors call into question the validity of UCN-surface scattering models used in the simulations in ref. [27], and therefore question the reliability of the extrapolation methods used to extract τ_n . They argue that the chosen surface roughness correlation model changes the interpretation of the experiment, though the results of their Monte Carlo study do not suffice to recommend a particular model or explicit correction to the experiment.

2.7 Magnetic Bottles

As discussed in section A.6, magnetic field gradients apply a force $F = \pm\mu_N\nabla|B|$ to UCN. In this way, an inhomogeneous magnetic field in a suitably designed apparatus can confine UCN. The use of this technique to measure the neutron lifetime was suggested more than fifty years ago by Vladimirskii[132], and magnetic confinement was first achieved by Abov *et al.* using a magnetic trap formed by a superconducting Halbach array (Halbach arrays will be discussed in chapter 3)[78, 79].

Since this time, only Paul *et al.* have produced a measurement of τ_n using magnetic confinement. This is achieved by radially confining slow neutrons using a magnetic storage ring[77]. The confining radial magnetic force is achieved using a 1 m diameter superconducting sextupole magnet. To deduce the neutron lifetime, UCN are loaded into the trap using a curved guide and stored for various times, after which a ^3He -filled neutron counter is inserted into the neutron

beam to count the surviving neutrons.

Neutrons exhibiting radial motion around the sextupole magnet experience a linear restoring magnetic force, giving both radial and vertical oscillations of the neutrons about their equilibrium orbiting trajectories. Due to even minor imperfections in the magnetic field profile, betatron oscillations of the neutrons may slowly mix the radial and vertical components of the oscillation. Therefore, a neutron initially experiencing vertical oscillations may slowly evolve into a trajectory with radial oscillations, strike the wall of the apparatus, and escape the storage ring. This is similar to the quasi-bound UCN in trap-based experiments. In this experiment, trajectories with large oscillations are removed before the storage interval with the temporary insertion of a neutron-absorbing material. This process, in the context of UCN traps, is often referred to as “cleaning” or “pre-conditioning” (cf. ref. [23]).

This introduces a loss mechanism due to this slow evolution of the occupied phase space of the neutrons in the experiment, leading to two important features in the apparatus. First, a decapole field is superimposed on top of the sextupole field; this introduces a non-linearity in the restoring force which more rapidly mixes the modes of oscillation, thus making this effect only prevalent for short storage times. Second, neutron-absorbing beam-scrapers are used to remove neutrons that experience large oscillation amplitudes. The scrapers remained in a fixed position for 20 s after the filling of the apparatus was completed, at which time they were moved outward. Long storage times are observed to be insensitive to the effect of loss due to the evolution of the beam phase space. An exponential fit to the number of surviving neutrons, discarding data at short storage times, gives $\tau_n = 877 \pm 10_{\text{stat}}$ s.

Since then, several magnetic-storage-based neutron lifetime experiments have been proposed or initiated. Ezhov *et al.* demonstrated the storage of UCN in a vertically oriented, cylindrical, high order multipole magnet trap composed of NdFeB magnets with FeCo poles as flux returns. The bottom of the trap is closed using an electromagnetic solenoid, and a polyethylene sheet is used at the top of the trap to clean quasi-bound neutrons. The authors estimate that, given the field gradients in the trap, the probability of depolarization is on the order of 10^{-4} s^{-1} . The

storage time was determined by counting the survivors from fill-and-empty cycles of the trap, which gave a storage time of $882 \pm 16_{\text{stat}}$ s over a 143 hour run campaign at the ILL[72].

A subsequent iteration of this experiment utilized a fomblin-coated aluminum pre-filling volume with an absorbing roof as a cleaner. After UCN from the source are filled into this volume, it is adiabatically lowered into the trap and opened. In addition, an outer solenoid surrounds the entire multipole trap to assure that $\mathbf{B} \neq 0$ everywhere. A guide beneath the trap leads to a UCN monitor, which, in addition to detecting surviving neutrons after storage, can also detect neutrons that depolarize during the storage time interval. Depolarization was noted as the dominant source of loss during storage, and the size of this loss mechanism could be changed by changing the current or orientation of the outer solenoid. The authors quote a statistical sensitivity of about 2 s, but no subsequent storage times or measurements of τ_n have been reported[73]. It is worth noting that the high rate of depolarization is possibly due to the accidental cancellation of the solenoidal field and the multipole field in a (presumably small) region of the trap.

An ongoing effort at NIST offers a novel alternative to the fill-and-empty paradigm of other material or magnetic bottle experiments[63, 36]. The apparatus consists of a superconducting, quadrupolar, Ioffe-Prichard trap filled with superfluid ^4He . A collimated cold neutron beam passes through the apparatus, downscattering UCN in the superfluid, thus filling the trap with UCN. With temperature $T < 250$ mK and an estimated ^3He concentration less than 5×10^{-13} , the absorption and upscattering rates in the helium are estimated to be $5 \times 10^{-6} \text{ s}^{-1}$ and 10^{-6} s^{-1} .

An additional feature of the experiment is that neutron decays are actively detected. When a trapped UCN decays, the β -ray ionizes the helium atoms which form excimer molecules. The singlet molecular state decays in about 10 ns and emits UV light (the triplet state with lifetime of 13 s is too long lived to be observed). A layer of wavelength-shifting TPB coats the walls of the helium volume, and this light is guided out of the trap by acrylic light guides and into two PMTs. Candidate events must produce a minimum of two photoelectrons in each PMT, thus reducing backgrounds related to neutron capture in the shielding materials. Experimental campaigns have demonstrated long storage times, and a forthcoming increase in the size of the trap (3.1 T field

strength, 8 liter volume) is expected to reach statistical sensitivity of 2.4 s over a 40 day reactor cycle.

In spite of the lack of improved measurements of τ_n using magnetic traps, storage times consistent with past τ_n measurements have motivated additional efforts. Namely, a vertical octupolar trap coupled to a dedicated, high density UCN source at the ILL[103], and a large-volume superconducting trap designed to detect protons from neutron decay[68].

2.8 Conclusions

In this chapter, we have reviewed the history of measurements of the neutron lifetime. The now mature neutron beam-based technique can be expected to achieve a few tenths of a second precision, but there remains a disparity between current beam- and bottle-based experiments[31] as well as a more general need for independent determinations of τ_n to improve the confidence in the global average of measurements. Though the material bottle method has produced the lowest quoted uncertainties, there is continuing discussion in the literature of the validity of model assumptions and extrapolation methods used in these experiments.

No magnetic-trap-based experiment has produced a measurement with precision comparable to material bottle or beam measurements, but these efforts have generally produced storage times that are in agreement with past measurements of the neutron lifetime. This suggests that a magnetic trap can be made to have low losses, and thus little or no extrapolation to determine τ_n . The efforts described in this chapter will generally utilize large trapping volumes for high statistical sensitivity, and high sensitivity will be necessary to vary and study potential systematic effects in these experiments.

As was the case even in the neutron storage ring of Paul *et al.*, and studied further in other experiments using magnetic confinement[53, 65], control over the occupied phase space of the UCN is important to mitigate the effects of quasi-bound neutrons (and the time dependence of the UCN phase space more generally). We note that all of the ongoing or proposed magnetic trap experiments mentioned here exhibit at least approximate cylindrical symmetry. This could

lead to quasi-bound UCN orbits and thus long cleaning times. The geometry of a magnetic trap and the dynamics of neutrons in the trap therefore play an important role in the reliability of the experimental methods for this class of experiment.

Thus, this discussion of past and current experimental efforts motivates many of the features of the experiment described in this dissertation. In particular, a large volume, low-depolarization, rapidly-cleaned trap is desirable.

3 Experimental Design

3.1 Introduction

The experimental design of the UCN τ prototype apparatus consists of a number of separate components. Coarsely, the components of the apparatus are the permanent magnet trap, holding field coils (to assure $|B| \neq 0$ throughout the trapping volume), the UCN cleaner (to remove high energy UCN from the trap), and the vanadium activation detector (to detect surviving neutrons). Beneath the trap is the UCN polarizer, spin flipper, guide system, and UCN monitor detectors. The trap geometry was motivated by neutron tracking studies[4], and the original experimental design was presented by Walstrom *et al*[62]. It is important to note, however, that some design parameters have changed since that publication.

Fig. 3.1 shows a schematic of the experiment. UCN from the source (to the left of the apparatus, not shown) pass through the polarizing magnet and spin flipper. At this point, they can be directed up through a piston-actuated door into the UCN trap, or directed to a UCN monitor at the end of a vertically curved section of guide. The cleaner can be moved into the trap via a pneumatic piston that is fed through into the vacuum system. An array of ^3He filled drift tubes line the outside of the jacket near the cleaner to detect upscattered neutrons. The vanadium foil can be lowered into the trap by a linear actuator driven by a stepper motor. In its raised position, the vanadium foil rests between two plastic scintillator paddles and an array of NaI detectors.

The performance of the UCN source at LANSCE has been reviewed, and measurements of UCN densities and spectra have been compared to comprehensive Monte Carlo models[24]. An understanding of the UCN source itself has important implications for understanding the




Figure 3.1: A schematic of the experiment. P is the polarizing magnet, S is the spin flipper, and B is the UCN monitor detector. There is another monitor to the left of the polarizing magnet (not shown). The holding field coils (not shown) are arranged outside of and around the vacuum jacket of the trap.

experimental methods and interpretation of data for the measurement of τ_n . The UCN output rate, energy spectrum, and time structure from the source are needed to interpret the rates observed in various detectors, and the (in)stability of the source can also cause systematic effects related to determining the number of UCN in the trap when it is filled. This will be discussed in sec. 6.12.

3.2 Permanent Magnet Trap

The magnetic trap consists of ~ 4000 aluminum-coated NdFeB permanent magnets arranged in an asymmetric bowl shape. Each magnet (1 in. deep, 0.5 in. wide, 2 in. long, ~ 1.35 T remnant field strength) is adhered to aluminum plates of varying sizes (typically $\sim 30 \times 30$ cm²) using West System 105 epoxy resin. In order to overcome the strong and rapidly varying forces of nearby magnets, custom jigs are used to insert each magnet into place. Once a plate has been populated with the permanent magnets, it is bolted to an aluminum frame. Fig. 3.2 shows the

constructed magnet trap. There is a rectangular region at the bottom of the trap that is absent: the magnet plate that fits into this position is not part of the trap frame, but is instead fastened to a pneumatic actuator which can raise the plate upwards to close the trap, or downwards to open it (this is described in more detail in sec. 3.3).




Figure 3.2: The completed permanent magnet trap, prior to insertion into the vacuum jacket. The rectangular hole at the bottom is filled by a separate magnet plate (not shown here) which is fastened to a pneumatic actuator.

The trap is in the shape of two piecewise torii. In cartesian coordinates x, y, z (with $+z$ being up), the inner material surface of the bowl is given by

$$x, y, z : (R + \rho)^2 + x^2 = r^2 \quad \text{and} \quad z < -1.0 \quad (3.1)$$

where $\rho = \sqrt{y^2 + z^2}$. The major and minor radii R and r are defined piecewise such that $R = 1.0$ m and $r = 0.5$ m if $x > 0$, and their values swapped otherwise. In words, each half of the trap consists of a slice near the edge of a torus centered at the origin, with azimuthal symmetry about the x axis, and the $x > 0$ side curving more steeply than the other half. Fixing the sum of the major and minor radii on either side of the $x = 0$ plane ensures that the trap

surface and its tangent vectors are continuous. In these coordinates, the bottom of the trap is at $(x, y, z) = (0.0, 0.0, -(R + r))$.

The desired field profile for the trap is achieved by arranging the permanent magnets into a Halbach array[89]. The magnetization direction of each magnet is alternated in a spatially periodic pattern. The salient features of this type of arrangement are a (nearly) vanishing field on one side of the array, and a strong magnetic field near the array surface on the other side. The field strength diminishes roughly as $\exp(-\kappa d)$, where d is the distance from the array surface and κ^{-1} is on the order of the size of the individual magnets. A visualization of a Halbach array and the resulting field is shown in fig. 3.3.

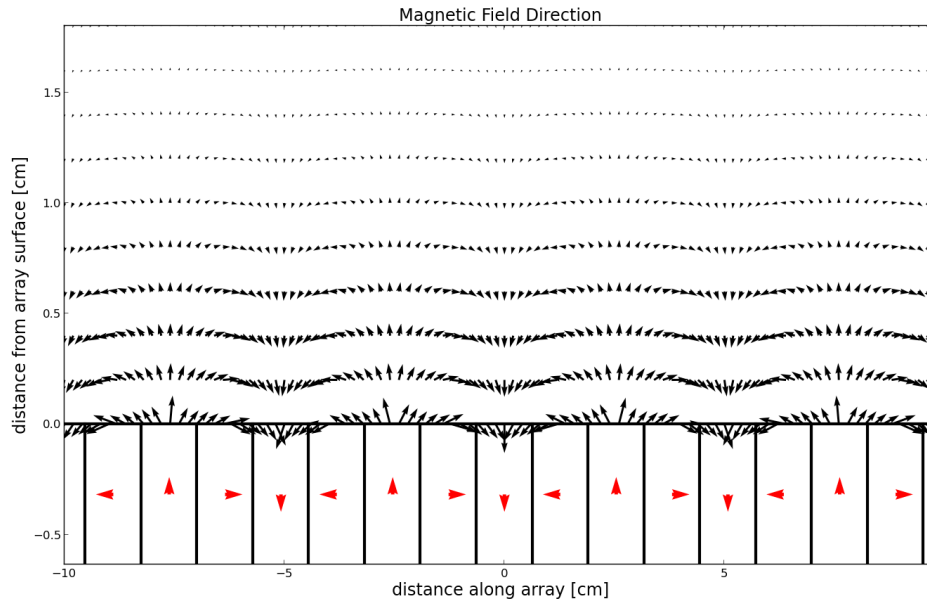


Figure 3.3: The magnetic field generated by a Halbach array. The red arrows represent the magnetization direction of each magnet, and the black arrows represent the magnetic field.

For an ideal, flat Halbach array with $\hat{\zeta}$ pointing normal from the surface, and $\hat{\eta}$ pointing along the direction of variation of the magnet orientations, we have that

$$\mathbf{B} = \frac{4B_r}{\pi\sqrt{2}} \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{4n-3} (1 - \exp(-k_n d)) \exp(-k_n \zeta) \left(\sin k_n \eta \hat{\eta} + \cos k_n \eta \hat{\zeta} \right) \quad (3.2)$$

with $k_n = 2\pi(4n - 3)/L$, and $L = 2$ in. is the periodic distance of the Halbach array. For the trap, the curvature is gradual, so the expression for the flat Halbach array is approximately true in the local toroidal coordinates η, ξ, ζ , given by

$$\zeta = \tan^{-1}(y/z) \quad (3.3)$$

$$\zeta = r - \sqrt{(\rho - R)^2 + x^2} \quad (3.4)$$

$$\eta = r \tan^{-1} \left(\frac{x}{\rho - R} \right). \quad (3.5)$$

Upon assembly, some of the populated magnet plates came into contact with plates that were already installed, causing some of the magnet plating, and in some cases small fragments of the magnets, to break off from the array. While the removal of the magnet coating doesn't pose a problem in terms of performance of the trap, the missing fragments of the magnets might be detrimental: if the field strength in the vicinity of the chipped surface is much less than the nominal ~ -0.8 T field, it acts as a "hole" for the low-field-seeking UCN, and thus can be a source of loss. While these missing fragments can eventually be replaced or repaired, these defects have been remedied by covering them with a single layer of Cu tape. This provides a high optical potential non-depolarizing surface to greatly mitigate the effect of any potential holes. We will discuss the implications of this in sec. 6.4.

3.3 Trap Door and UCN Guides

In order to open the trap, the bottom-most 6×6 in² magnet plate is separate from the rest of the trap. The plate is fastened to a stainless steel rod, and can be raised into a closed position so that it is flush with the rest of the trap, or opened by lowering the plate by approximately 4 inches. A brass wheel-shaped frame holds the rod so that it remains laterally aligned in the presence of large (~ 400 lb) lateral forces on the trap door due to the rest of the magnet array. UCN can transmit through the spokes of the brass wheel and around the trap door magnet to enter the trap. The rod is connected at the bottom to a pneumatically actuated piston. The position

can be moved by sending a proportional analog voltage to the piston controller electronics. The piston can be lowered or raised in about 1 second.




Figure 3.4: UCN are directed to or through the tubular sections to the left or right. A UCN that goes up passes through the brass spokes, and up around the steel rod and magnet plate (top). The pneumatic piston is beneath the guide cross to push the piston rod upwards or downwards.

The piston interfaces with the trap door (which is within vacuum) through a double o-ring seal on the rod, and this volume is pumped independently of the main trap vacuum. A modest amount of vacuum grease was applied to this seal to facilitate smooth movement, but the rod was cleaned of this grease using soap and alcohol. The top of the rod is exposed to the UCN volume, and any greased surfaces exposed to UCN make that surface have a high loss-per-bounce for UCN. Aside from the brass parts, most of the rest of this assembly is composed of stainless steel, which, while mechanically robust and in general UCN-friendly, has a somewhat higher

spin-flip-per-bounce probability than that of copper or carbon.




Figure 3.5: A cutaway of the trap door and guide cross assembly. The cross assembly (middle) surrounds the piston rod. UCN can pass through the spoke holes up to the top of the cross assembly, past the magnet plate (door), and into the trap. The features above the guide cross assembly are encased in the vacuum vessel of the trap.

All of the guides leading from the trap door assembly to the pre-polarizing magnet are made of 3 in electropolished copper, except for a 1 m section of quartz tube coated with a layer of diamond-like carbon. This guide is used in the vicinity of the spin-flipper because it is non-conductive. The RF resonator and magnetic field coil for the AFP spin-flipper surround this guide. Downstream of the trap door, past the aluminum shutter, the guides that curve downward to the UCN detector are electropolished stainless steel.

3.4 Holding Field Coils

Because the field strength of the Halbach array is negligible except near the array surface, a UCN in the trap could undergo a Majorana spin-flip when not near the surface. To remedy this, electromagnetic coils are placed outside of the apparatus to generate a ~ 100 Gauss magnetic

field throughout the trapping volume. To assure that this field does not cancel the Halbach array field, the coils are oriented along a circle in the y, z plane, so that the holding field is perpendicular to the Halbach field everywhere. A schematic of this arrangement is shown in fig. 3.6.




Figure 3.6: The holding coils (shown from the side in red) produce field lines (dotted blue) which are perpendicular to the field of the Halbach array.

The ten water-cooled, rectangular, copper coils are fastened to the outside of the vacuum jacket of the experiment. Each coil is split into two L-shaped pieces, each of which consists of eight hollow-core bars of copper. Each bar is wrapped in epoxy tape for insulation, and bent to the L shape on a bending table. The bundle of bars is then wrapped in another layer of epoxy tape, and a layer of shrink tape. The bars are cut to length, placed into an 80/20 frame, and baked to constrict the shrink tape. The ends of the bundles are fanned out so that each bar can be fastened with a bussing clamp to its corresponding bar on the other L split, and the ends are fitted with a hose to route water to each bar. Silver conducting paste is used for electrical connections to assure a large contact area between each L split. Detailed views of the coils are shown in figs. 3.73.8.

The coils are bolted to the outside of the vacuum jacket, and electrically connected in series using 4/0 AWG cable. Water lines are connected to each coil in parallel, with de-ionized water at 80 – 90 PSI supply pressure. A high current power supply is used to generate the magnetic field, and the maximum current is 300 A with a 40 V drop across all of the coils.



Figure 3.7: The design of the holding coil. Each L-shaped copper bundle rests in an 80/20 frame, and the frames are fastened together with joining brackets.

In the coordinate system used in section 3.2, we can write an expression for the holding field, assuming that it is an ideal toroidal field:

$$\mathbf{B}_{\text{hold}} = B_{h0} \frac{r + R}{\sqrt{y^2 + z^2}} \hat{\xi} \quad (3.6)$$

This field is everywhere perpendicular to the field from the Halbach array.

3.5 AFP Spin Flipper

High-field seeking UCN that pass through the polarizing magnet (see fig. 3.1) are polarized by an Adiabatic Fast Passage (AFP) spin flipper. The design of the spin flipper used here is a scaled-down low-field version of that used for the UCNA experiment[30]. The resonating cavity consists of a printed circuit board with sixteen ~ 8 cm long parallel copper strips wrapped into a cylindrical shape. Adjacent strips are connected by $1 \mu\text{F}$ capacitors, which gives a resonant frequency of approximately 400 kHz, corresponding to a resonant field for the neutron of 140 G.

The ambient field is generated by an electromagnetic solenoid which has two segments with a different number of turns of copper wire. This produces the necessary resonant field with the small gradient needed for high efficiency adiabatic fast passage spin flipping. The field parallel to the UCN guide axis along the solenoid is shown in fig. 3.9. The spin flipping efficiency is expected to be at least 98% for UCN of typical velocity coming from the LANSCE source.




Figure 3.8: A closeup of the individual copper conductors. The bars are fanned out, and fitted with water connections for cooling. Each bar is electrically connected to its corresponding bar on the other L-split with a copper bussing clamp.

3.6 UCN Cleaner

The UCN cleaner consists of a $36.5 \times 66.0 \text{ cm}^2$ low-density polyethylene sheet fastened onto an aluminum frame, mounted horizontally on one side of the trap. The sheet is fastened using double sided tape on the frame while under tension. A pneumatic piston raises the sheet to the top of the trap, or lowers it 7 cm downwards into the trap. Nine cylindrical aluminum drift tubes (5.1 cm diameter, 1.2 m long) were installed outside the vacuum jacket above the cleaner, and each tube was filled with 1.8 bar of ^3He to measure the upscattered neutrons from the polyethylene of the cleaner. Figs. 3.10 and 3.11 show the design and placement of the cleaner. The use of the cleaner in the experiment is discussed in chapter 4, and the systematic effects associated with the cleaner discussed in chapter 6.

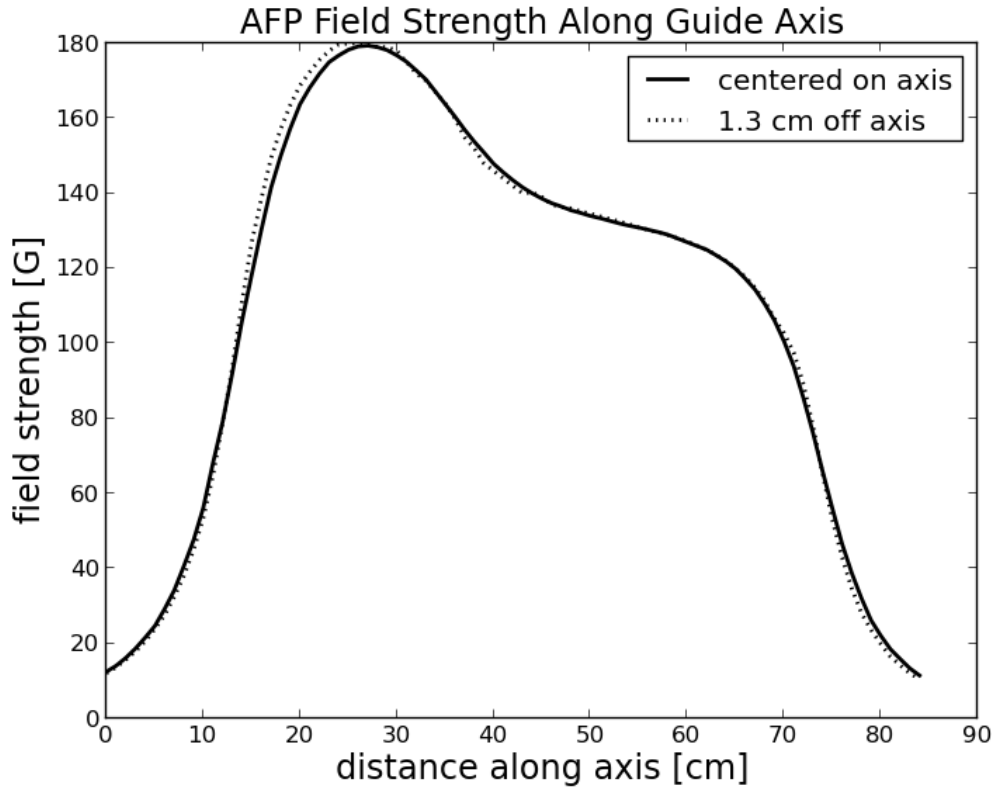


Figure 3.9: The component of the magnetic field parallel to the UCN beam axis produced by the AFP solenoid with 5 A of applied current.

3.7 UCN Detectors

3.7.1 Introduction

To monitor the flux of UCN delivered to the experiment, UCN counters are incorporated into the guide system beneath the apparatus. UCN are typically counted with ^3He -filled detectors [35]. Neutrons undergo the reaction $^3\text{He}(n,p)t$ with a Q value of 0.764 MeV, and the resulting proton and triton ionize a stopping gas with a small neutron absorption cross section such as CF_4 . The current shortage of ^3He , however, demands the development of alternative high-efficiency UCN detection schemes [101].

^{10}B provides an alternative to ^3He . Ions are produced through the neutron-capture reaction $^{10}\text{B}(n,\alpha)^7\text{Li}$, with a Q value of 2.8 MeV. This reaction leaves the ^7Li nucleus in an excited




Figure 3.10: The UCN cleaner. The rectangular polyethylene sheet is fastened to the Al frame, which swivels on four legs so that its height can be changed. The linkage on the back of the sheet connects to a pneumatic actuator which is fed through a bellows from outside the vacuum jacket. The two Al wing-shaped plates bolt the assembly to the frame of the Halbach array.

state with a 94% branching ratio, which decays to the ground state with the emission of a 0.48 MeV photon [100]. Recently, boron coatings for UCN detection have been used in the form of track-etch detectors and Gas Electron Multipliers [67, 60].

UCN are monitored in this experiment using a boron-lined detector. The inside of a cylindrical ion-chamber is coated with a boron that was previously used for multi-layer thermal neutron detectors at LANSCE[133].

In this section, we investigate the performance of this design. Its efficiency is directly compared to a duplicate ion chamber filled with a $^3\text{He}/\text{CF}_4$ gas mixture. In addition, the UCN interaction with the layer is characterized by separately detecting UCN and thermal neutrons, and comparing the pulse height spectra. The pulse height spectra of detectors with solid neutron-reactive layers can be understood using a simple Monte Carlo technique [39, 109]; we perform this analysis for




Figure 3.11: A cutaway of the apparatus showing the placement of the UCN cleaner.

the boron-coated ion chamber, and the result indicates that UCN mostly interact with the outer surface of the boron layer.

3.7.2 Design

The ion chamber is a 76.2 mm long stainless steel tube, with a 76.2 mm outer-diameter and 1.7 mm thick walls. UCN transmit into the detector through a 0.3 mm thick aluminum window fastened to one side of the tube, and a polyoxymethylene cap electrically isolates the detector from UCN guides and vacuum components. The ion chamber anode is a single 4.8 mm-thick copper conductor, protruding 50.8 mm into the cylindrical volume. The detector is pumped and filled with gas through a VCR feedthrough (see fig. 3.12).

The coating is made with $58 \pm 30\%$ mg of $> 97\%$ isotopically pure boron powder. The average grain size is measured with an optical microscope to be less than $1\mu\text{m}$. The powder is mixed with an acetone (15 ml) and polystyrene (4 mg) solution, which acts as an adhesive. The mixture is then sprayed manually in three coats over the interior of the detector volume. The smoothness and uniformity of the coating are not characterized. We estimate the layer thickness




Figure 3.12: A schematic of the UCN detector assembly. The bottom plate, aluminum window, and polyoxymethylene cap are sealed with o-rings.

to be $1.5 \pm 0.5 \mu\text{m}$, assuming that the coating is distributed evenly on the detector walls. The ^{10}B capture cross section is 3835 barns at 2200 m/s; at a number density of $1.31 \times 10^{23}\text{cm}^{-3}$, the mean-free path $\ell \equiv (n\sigma)^{-1}$ is 27 nm for 3 m/s UCN.

The use of isotopically enriched boron is necessary for UCN detection: natural boron has a high optical potential ($V \approx 200 \text{ neV}$) that would reflect UCN away from the layer, greatly diminishing the absorption efficiency. Pure ^{10}B has $V_f = -3.4 - 36.4i \text{ neV}$. With 3% of ^{11}B ($b = 6.65$), the net potential is $3.5 - 36.4i \text{ neV}$. For a flat surface with this potential, the probability of UCN reflection per bounce averaged over all angles of incidence is 15% for 50 neV

UCN. When operated in a mode where the detector is below mean beam-height, this reflection probability is much lower. In addition, the UCN within the chamber have more than one chance to capture on the wall. It is thus likely that the detector window and upscattering of the UCN within the detector gas are the limiting factors for the total efficiency of the detector.

3.7.3 Characterization Methods

We compare the performance of the boron-coated detector to the same detector housing filled with 10 mbar of ^3He gas. Both detectors are filled with 500 mbar of CF_4 gas. The anodes are biased to 500 Volts, and pulses are collected using Amptek A250CF charge-sensitive preamplifiers [3]. The pulses are then amplified by spectroscopy amplifiers with a $6\ \mu\text{s}$ time constant and read into a multi-channel analyzer. For CF_4 gas at a pressure of 500 mbar and temperature of 300 Kelvin, the ranges of the charged particles (estimated using SRIM2008 code [127]) are no larger than 7.8 mm. The full energy of the ions (less the energy loss in the solid layer) can therefore be deposited in the gas.




Figure 3.13: The detector mount configuration. Lengths are not to scale.

To acquire data using UCN, the helium and boron-coated detectors are fastened to either side of an electro-polished guide tee and mounted on top of a gate valve, below which is a UCN guide leading to the UCN source. The gate valve is closed to measure the background rate of the

detectors. A small aperture on the underside of the UCN guide leads downward to a multi-wire proportional counter used to monitor the incoming flux of UCN. This configuration is shown in fig. 3.13. Thermal neutron data are acquired using a ^{252}Cf neutron source moderated with room temperature polyethylene. An energy calibration is performed for the boron-coated detector by introducing an additional 5 mbar of ^3He gas into the chamber and using the known full energy peak of 0.764 MeV.

3.7.4 Results

Figure 3.14 shows a comparison of the pulse height spectra for the helium and boron-coated detectors. Because neutrons capture on the wall of the boron detector, only one of the Li or α will ionize the gas, and there is no full-energy peak. We are thus left with two prominent edges corresponding to the full energy of each ion. There is a less intense higher energy edge due to α ions from the 6% decay branch. The discriminated count rate versus anode voltage for the boron-coated detector is shown in fig. 3.15. The saturated ion region is reached for applied biases above 300 V.

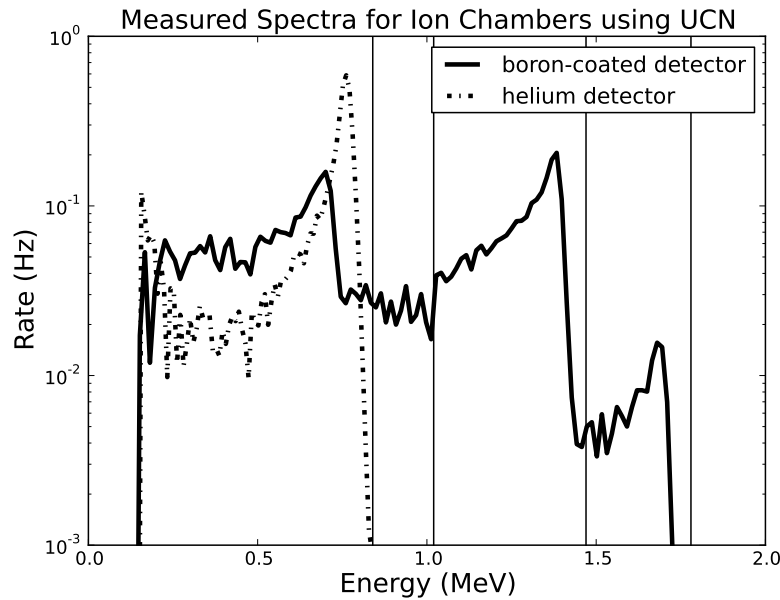


Figure 3.14: Pulse-height spectra for the helium and boron-coated ionization chambers using UCN. The vertical lines represent the Li and α energies of 0.84, 1.02, 1.47, and 1.78 MeV.




Figure 3.15: The discriminated count rate from the boron-coated detector as a function of applied anode bias. The discrimination threshold is adjusted for each voltage as needed to remove low energy background pulses.

To measure the relative detector efficiencies of the helium and boron-coated designs, data are acquired simultaneously for both detectors. The spectrum for each detector is measured with the gate valve open, and then measured with the gate valve closed to determine the background spectrum. The signal spectrum is formed by normalizing all counts to the monitor rate, and subtracting the closed-valve spectrum from the open-valve spectrum. In addition, counts below 0.15 MeV are discriminated, as γ -ray backgrounds from neutron capture are potentially high in this energy range. The background rate for both detectors was $700 \text{ mHz} \pm 15\%$, primarily due to thermal neutrons that escape from the UCN source.

The ratio of the helium detector signal to the boron-coated detector signal is $(94 \pm 8) \%$. The error includes the statistical uncertainty of the background spectra as well as an estimated uncertainty in establishing the discrimination threshold for the two detectors.

Finally, fig. 3.16 compares the pulse-height spectra of the boron counter using UCN and thermal neutrons. The full energy peaks are significantly broader when detecting thermal neutrons: faster neutrons can penetrate further into the boron layer than UCN, and the subsequent

ions must traverse a significantly larger portion of the coating, thus losing more energy prior to ionization in the gas. This will be discussed quantitatively in the next section.

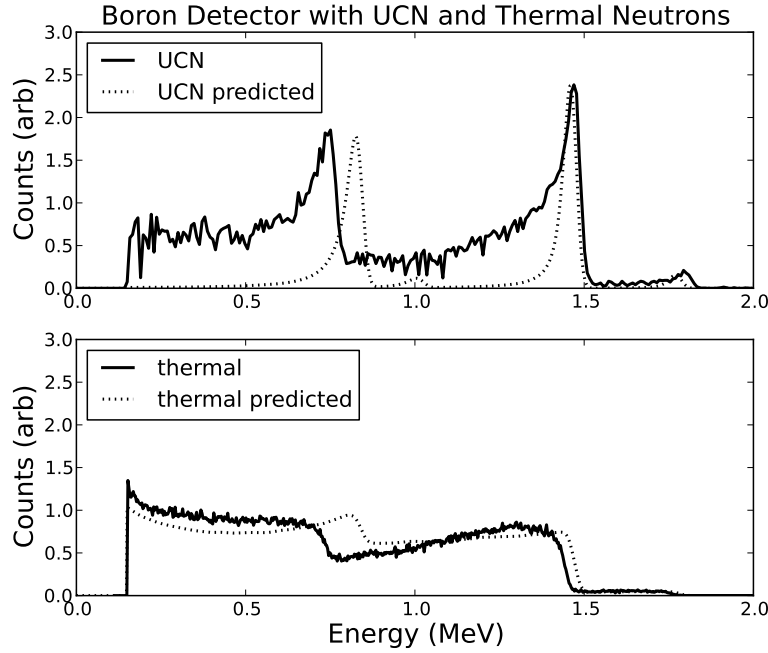


Figure 3.16: Comparison of boron-coated detector spectra for UCN and thermal neutrons. The spectra are scaled so that their integral is unity. The predictions are discussed in section 3.7.5.

3.7.5 Discussion

Because the detector windows and stopping gas are the same for the ^{10}B and ^3He detectors, only a difference in the neutron absorption efficiency and ion collection efficiency can cause a discrepancy in the relative detector efficiency. The capture probability for the ^3He detector is close to unity: typical UCN (50 neV) in the gas have $\ell = 11$ mm, which is several times smaller than the detector dimensions. In addition, the charge collection efficiency is expected to be high.

The efficiency of the boron detector depends on the probability of a UCN transmitting into the layer, as well as the probability that either of the resulting α or ^7Li ion escapes the layer with sufficient energy to ionize the stopping gas. The probability of the latter process is given by

[133]:

$$\epsilon_{boron} = \frac{1}{2} \left(1 - \exp \left(-\frac{T}{\ell} \right) \right) \left(1 - \frac{\ell}{R_0} \right) + \frac{T}{2R_0} \exp \left(-\frac{T}{\ell} \right). \quad (3.7)$$

Here, $T = 1.5 \mu\text{m}$ is the layer thickness, and R_0 is the range of either the α or Li ion in ^{10}B . The respective ranges of the α and Li are 3.6 and $1.9 \mu\text{m}$ for the 94% decay branch, and 4.4 and $2.2 \mu\text{m}$ for the 6% branch. In all cases, $\ell \sim 27 \text{ nm}$ is much smaller than R_0 and T , so that the efficiency is close to unity for UCN that enter the layer. The measured relative efficiency of the ^3He and ^{10}B detectors therefore indicates a reasonable efficiency for UCN to transmit into the layer.

The pulse-height distribution for a single ion is given by the number of detected ions created at an initial depth in the layer x and angle to the normal of the surface θ that escape the layer such that their final energy is E . That is,

$$\frac{dN}{dE} = N_0 \int_0^1 d(\cos \theta) \int_0^\infty dx \times P(x, \cos \theta) \delta(E - \eta(x/\cos \theta)). \quad (3.8)$$

Here, P is the probability of an ion starting at a depth x and angle θ within the layer, and is given by

$$P dx d(\cos \theta) = \ell^{-1} \exp(-x/\ell) dx d(\cos \theta) \quad (3.9)$$

where ℓ is the mean free path of the neutron incident on the boron layer.

Further, η is the final ion energy given an initial depth x and angle θ , and is given by

$$\eta = E_0 - \int_0^{x/\cos \theta} \frac{dE}{dL} \cdot dL \quad (3.10)$$

where E_0 is the initial ion energy, and dE/dL is the ion energy loss per distance traveled in the boron layer. We calculate η using dE/dL for the lithium and α ions in solid boron-10 using SRIM2008 (see for example fig. 6 in [39]).

We perform a Monte Carlo integration of equation 3.8 to find dN/dE , adding the contribution of both ions and decay branches. The resulting spectrum is convolved with a gaussian distribution

of width 0.015 keV, which is commensurate with the width of the full energy peak for the ^3He pulse height spectrum; this is a simple means of including the instrumental width into the theoretical prediction [109]. This calculation is performed separately for $\ell = 27$ nm (UCN) and $\ell = 20\mu\text{m}$ (thermal neutrons), and the results are represented by the dashed lines in fig. 3.16.

The model captures the qualitative shape of the spectra. There are, however, an excess of counts below the full energy peaks in the UCN data which are not predicted by the model. The ^{10}B powder grain size is similar to the ion ranges, so that surface roughness of the layer may cause additional energy loss for detected ions with large angles with respect to the surface normal. In addition, the upscattering of UCN in the detector may introduce a thermal component to the UCN spectrum. While this observation does not effect the total efficiency, it demands future investigation. Furthermore, it is evident from fig. 3.16 that the the observed energies of the Li peaks are measured to be 10% lower than their expected values. This has been reproduced under several operating conditions, and it is not explained by a discrepancy between the theoretical and observed detector gain and offset. The effect is therefore not understood at this time.

3.8 Vanadium Activation Detector

To count the surviving UCN in the trap, the trap door can simply be opened and the UCN drained back into the guide system where they are detected by a UCN counter. As discussed in chapter 2, this is the approach used in past neutron lifetime experiments using material traps. However, draining times can be long, and the efficiency to transport UCN from the trap to the counter may depend on the UCN spectrum and occupied phase space of the trap. It is therefore desirable to detect UCN in the trap, because it offers the potential to make a detector which rapidly and efficiently detects all of the neutrons with less sensitivity to the phase space of the UCN in the trap. To this end, the apparatus incorporates a prototype detector that uses vanadium activation to count the surviving UCN. After a storage interval, a $127\mu\text{m}$ -thick, $10.2 \times 20.3\text{ cm}^3$ rectangular foil is lowered into the trap to absorb the UCN. The foil is then raised out of the trap into a detector array to measure the activity of the foil and deduce the number of survivors.




Figure 3.17: A cutaway of the vanadium detector. A 5 cm thick rectangular Pb background γ -ray shield surrounds the array of 8 NaI detectors. The vacuum jacket between the NaI detectors houses the V foil (which can be lowered into the trap) which is surrounded by the plastic β detectors.

Vanadium is chosen because it has a slightly negative optical potential ($V_O \approx -7$ neV), and thus a high UCN transmission probability. In addition, the upscattering probability has been shown to be approximately 1% of the total loss (absorption and upscattering), so that UCN that strike the foil have a high probability of nuclear absorption[42]. The activated ^{52}V ($T_{1/2} = 3.743(5)$ m) emits a β -ray with a 2.5 MeV endpoint that is coincident with a 1.4 MeV γ -ray with a more than 99% branching ratio. For the prototype, a detector array for the coincident detection of the β and γ rays is employed, providing a low-background method for detecting ^{52}V decays.

A cutaway of the vanadium foil and detector array is shown in fig. 3.17. The foil is mounted in a steel frame, and fastened to a magnetically-coupled linear actuator. The actuator can move the foil (in vacuum) from the detector array above the apparatus to the bottom of the trap in less than one second. The β detectors consist of two 4 mm thick, 10.2×30.5 cm² polyvinyl toluene scintillators placed on both sides of the foil. Acrylic light guides couple to one edge of each scintillator, leading out of the vacuum jacket to 5 cm diameter PMTs. The sides of the vacuum jacket are 0.3 cm-thick Al to assure that γ -rays transmit outside of the apparatus. A




Figure 3.18: A schematic of the detector geometry. The Pb shield surrounds the stack of NaI detectors. The black line represents the vacuum break. The plastic scintillators immediately surround the V foil within vacuum. The foil is moved down into the trap using a magnetic coupler to connect to a linear actuator outside the vacuum.

stack of $5.1 \times 10.2 \times 20.3 \text{ cm}^3$ hermetically-sealed NaI crystals coupled to PMTs are used for γ -ray detection. In the February 2013 experimental campaign, two NaI detectors (one on each side of the vanadium foil) were operational, and the plastic scintillator was of type BC404 (2 ns decay time). In the subsequent run campaign, 8 NaI detectors were implemented, and the plastic scintillators changed to a phoswich sandwich consisting of a 0.5 mm layer of EJ212 scintillator (2 ns decay time) backed by a 3.5 mm layer of EJ240 (285 ns decay time). The phoswich geometry is used (in conjunction with a waveform digitizer) to discriminate pulses which only contain the slow component: β -rays from the vanadium decay must necessarily pass through the fast scintillator

and stop in the slow scintillator, so that events with only a slow component cannot be due to ^{52}V β -decay. A diagram of the detector geometry is shown in fig. 3.18

3.9 Automation and Data Acquisition

For the Feb 2013 run campaign, a FAST ComTec MPA3 was used to acquire signals from the UCN detectors, upscatter detectors, and vanadium detector array. A global 20 MHz clock provides timing information for each detector channel, and the initial CPU time of each run is also recorded. Shaped and amplified signals are inputted, and the pulse height recorded for each event. In the case of the β detectors, a hardware anti-coincidence is performed, and hardware threshold is set, and the resulting digital pulse is sent to the MPA3 system to record the event time.

In the second run campaign, this system was replaced with a 12 bit 250 MHz waveform digitizer. While only the waveform acquisition for the β detectors was necessary, all detector channels were read into the waveform digitizer to preserve precise relative timing across all detectors. Each digitizer channel is triggered independently by a preset minimum voltage threshold for the given channel. Upon triggering, a set number of pre-samples is stored, and the remaining waveform is recorded. Signals are sent to a PC over ethernet, and each event contains the trigger time from the 250 MHz clock and the packet time from the ethernet transfer. Timing filtering amplifiers are applied to the analog signals before input into the digitizer, which provides cleaner signals and good matching of each signal to the dynamic range of the digitizer with a negligible reduction in timing resolution.

In order to perform a measurement cycle, many devices must be activated in a repeatable time sequence. Valves must be open and closed, the cleaner and vanadium actuators raised and lowered at particular times, and runs must be synchronized with the detectors. To do this, a custom-made program is used to create individual POSIX threads which set timers, and execute commands for the various devices. For most of the controls, a central measurement and automation device (LabJack U3) is used.

To actuate the cleaner, shutter, and gate valve, the control pins on the U3 are configured for

digital output. The output is sent to the control side of a solid state relay. On the load side of the relay, a 24 V power supply connects in series with the relay and the actuator solenoid. Thus, when the U3 sends a 5 V signal, the relay sends power to the solenoid, actuating the device. The LANSCE proton beam enable is a TTL signal sent to the accelerator's central control room. When no voltage is applied, a kicker diverts the hydrogen ions away from the RF cavities in the accelerator. When a high signal is sent, the kicker is deactivated, and protons are incident on the UCN target, providing arbitrary control over beam delivery to the spallation target for the UCN source.

A 14 bit DAC module that connects to the U3 and communicates via i^2c protocol is used to send an analog position signal to the trap door piston. The position encoder output of the PCS is sent to an analog input of the U3, and the control program logs the piston position in regular (typically 10 s) intervals.

The PC communicates with the vanadium actuator controller via RS232 connection. The GUI sends serial commands to the controller to set the acceleration and speed parameters to the stepper motor, and instructs the stepper motor to raise and lower the vanadium foil using the same type of timing threads as the labjack.

For synchronization of the slow control and detector DAQ, the MPA3 has a digital output which sends a high signal when a run has commenced. This is inputted into the U3, and the control GUI waits until the detectors are acquiring data before the control sequence is started. To perform a run, the user inputs the delay and spacing between each actuator event into the control GUI. When the run is started, the GUI waits for the signal from the MPA3 system, and the run commences. The timing parameters are written as a data entry into an HDF5 file along with a date time stamp and user-inputted description of the run. When the control run is complete and the MPA3 data acquisition run is stopped, the MPA3 software writes a time-stamped file of the detector events. In this way, the control runs can be correlated in time with detector data acquisition runs for later analysis. Upon implementation of the waveform digitizer in later run campaigns, the digitizer acquisition program was integrated into the slow control program, and

runs synchronized in software.

The reliability of the timing of each device can have implications for a high precision measurement of the neutron lifetime. Prior to taking data with UCN, the above system was tested by performing ~ 200 run cycles wherein the various devices were actuated for various times (typically 20 to 30 s). Each event was timed independently by a high frequency clock and a stop watch timer. The delay and duration of each event varied by typically 0.01 s.

4 First Experimental Campaign

4.1 Introduction

The apparatus was assembled as in fig. 3.1 in February 2013 at LANSCE. A one-week measurement campaign was performed to test the components described in chapter 3. During this time, background data were acquired, spin contrast measurements performed to tune the spin flipper, and UCN were stored in the trap. Data were acquired using the V activation detector, but the signal-to-background was poor due to limited NaI detector solid angle efficiency and low UCN density. Therefore, these data were not used to characterize the apparatus in this campaign. The upgrade and commissioning of the V detector will be discussed in chapter 5. In this chapter, we determine the storage time of the trap, examine detector backgrounds, and study the signal from the upscatter detectors.

4.2 Backgrounds

The backgrounds in the monitor detector were studied by computing the count rate while UCN are stored in the trap and the beam is off. These conditions are used since they are most similar to those for the emptying process. The time integration window for each run is also limited to assure that UCN that are being emptied into the detector from the guides are not included. Figure 4.1 shows this rate over the course of the run campaign.

Between runs 70 and 140 there is a notable increase in the background rate, most likely due to repeated re-assembly of ^{10}B -loaded rubber shielding around the outside of the detector, and due to a (subsequently discovered) intermittent ground loop which caused excess noise in the detector preamplifier. For runs after 140, It was observed that the low-energy tail of the pulse

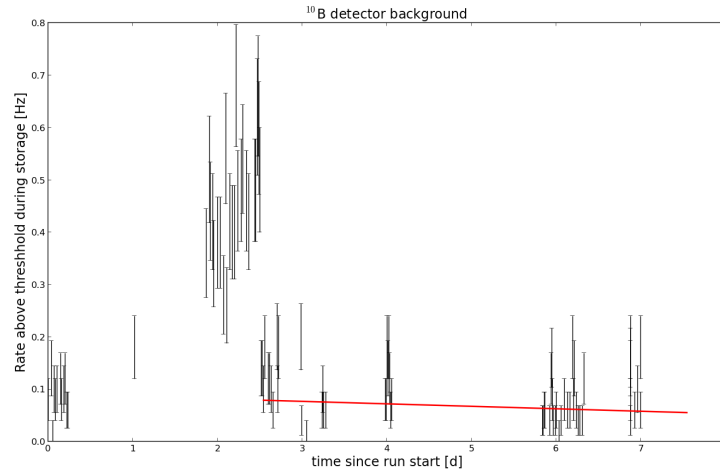


Figure 4.1: The average beam off background rates for the ^{10}B detector over the run campaign. The red line is a linear fit for runs greater than 140, with a slope of approximately 90 mHz per day.

height spectrum contributed a substantial amount to the rate. It's therefore worth considering whether an appropriate pulse height cut improves the signal-to-background: figure 4.2 shows the background versus time with a lower level cut on the ADC events of 0.65 V which removes much of the observed noise tail, and this improves the stability for later runs by a factor of three.

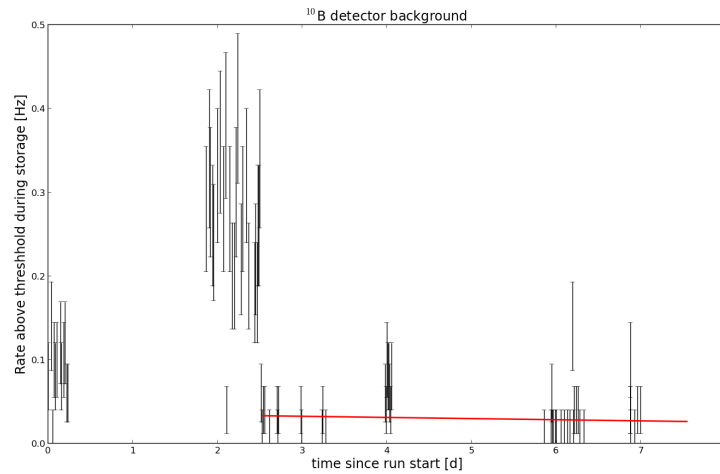


Figure 4.2: The beam off background rate for the ^{10}B detector with a lower level ADC cut of 0.65 V.

The boron detector (analyzed for its signal during fill-and-empty measurements) is not consis-

tent with a flat background, even accounting for the increase in the rate from run 70 to run 140. Care must therefore be used to subtract the background either on a run-by-run basis, or in some way account for drifts in the background over time. As discussed in the next section, a run-by-run fit of the data including a background term accounts for this background drift. A single cause for the ostensible background drift was not identified, but it is likely due to a combination of changes in detector shielding (mentioned above), slow contamination of the detector gas due to diffusion or leakage through the o-ring seals of the detector chamber, and changes in the environmental background such as beam-induced activity.

4.3 Determination of the Storage Time

In this section we present a determination of the storage time of the trap. UCN are loaded into the trap, cleaned, stored for various times, and emptied into the ^{10}B counter to determine the storage time constant τ_{store} . A typical filling and emptying cycle of the apparatus is shown in fig. 4.3. At the beginning of a cycle, all UCN valves are open, the cleaner is lowered, and the proton beam is turned on. After $t_{\text{pre}} = 30$ s, the shutter is closed and filling continues for another 180 s until the time t_{fill} . Once filling is complete, the shutter is then opened to drain UCN from the guides, the proton beam is turned off, and the trap door and main gate valve are shut. The cleaner remains in a lowered position for an additional 30 s; it is raised at time t_{hold} , and the UCN are stored for variable amounts of time (100 to 2000 s). The trap door is then opened at time t_{empty} to measure the number of surviving UCN and measure the detector background. The storage time is then given by $t_{\text{store}} = t_{\text{empty}} - t_{\text{hold}}$.

Fig. 4.4 shows the UCN monitor rate during a typical measurement cycle. The count rate increases as the density in the guide system saturates. Once the shutter is closed at t_{pre} , the count rate reduces due to deflecting away neutrons with $E < V_{\text{O}}^{(\text{Al})}$. At the end of the filling cycle, the count rate rapidly increases due to re-opening the shutter, and then diminishes with time as the UCN drain from the guide system. Upon re-opening the trap door at t_{empty} , the surviving neutrons are counted.




Figure 4.3: The timing of components during a fill and empty cycle. The beam is on (off), valves are open (closed), and the cleaner is down (up) for a high (low) signal.

After filling the trap, some UCN remain in the guide system for an average time of $\sim 5 - 10$ s before being lost or detected. For short storage times, this affects the otherwise constant background in the monitor detector. To correct for this, the counting rate from time t_{fill} until t_{empty} is fit to the function

$$B(t) = B_0 \exp(-\beta t) + B_1 \quad (4.1)$$

where B_0 , B_1 , and β (typically $\sim 0.15 \text{ s}^{-1}$) are free parameters. As an example, for the 2000 s storage time runs $B(t)$ is dominated by $B_1 \approx 0.02$; the average signal to background (integrated over the signal window) for these runs was approximately 1.7. This incorporates the detector background B_1 along with the UCN draining from the guide system at time t_{hold} .

For sufficiently stable operation of the UCN source, the initial number of trapped UCN is proportional to the detector rate from time t_{pre} until t_{fill} . The mean rate R during this time window is used as a normalizing factor for the emptying signal. The ratio of counts P with the shutter open to the rate R provides an indication of changes in source performance, which could cause fluctuations in the ratio of trappable to countable UCN as discussed above. Runs with




Figure 4.4: The ^{10}B counter rate during a measurement cycle. From left to right, the vertical lines represent time t_{pre} when the shutter is closed, t_{fill} when the trap door is closed and shutter opened, and t_{empty} when the trap door is opened.

P/R fluctuating by more than $\sim 20\%$ from the nominal range were rejected. This amounted to rejecting 9% of the experimental cycles. This cut improves the goodness-of-fit, but does not cause a significant change in the central value or standard error; thus, these fluctuations do not have a significant effect for the results in this chapter, but need further examination in order to interpret high precision data that will be used to extract τ_n in future run campaigns. The potential causes and treatment of these fluctuations are discussed in more detail in chapter 6.

With the above, we define the signal S of stored UCN to be

$$S \equiv \frac{1}{R} \frac{1}{\Delta t} \int [D(t) - B(t)] dt \quad (4.2)$$

where the limits of integration run from t_{empty} to $t_{\text{empty}} + 100$ s, Δt is the integration bin width, and $D(t)$ is the measured counter rate while emptying. The signal for various t_{store} is shown in fig. 4.5, for a total data acquisition time of approximately twenty hours. We perform a least squares fit of $N \exp(-t/\tau_{\text{store}})$ to the data, from which τ_{store} is deduced. The measurement is repeated

with the holding field deactivated, which reduces the storage time due to depolarization of the trapped UCN. We find that $\tau_{\text{store}} = 860 \pm 19 \text{ s}$ ($\chi^2/\nu = 0.87$) with the holding field activated, and $\tau_{\text{store}} = 470 \pm 160 \text{ s}$ ($\chi^2/\nu = 1.17$) with the field deactivated. The fit value of τ_{store} with this method is consistent with a determination by computing the log-ratio of the signal of long and short storage times (as done, for example, in refs. [27] and [23]).

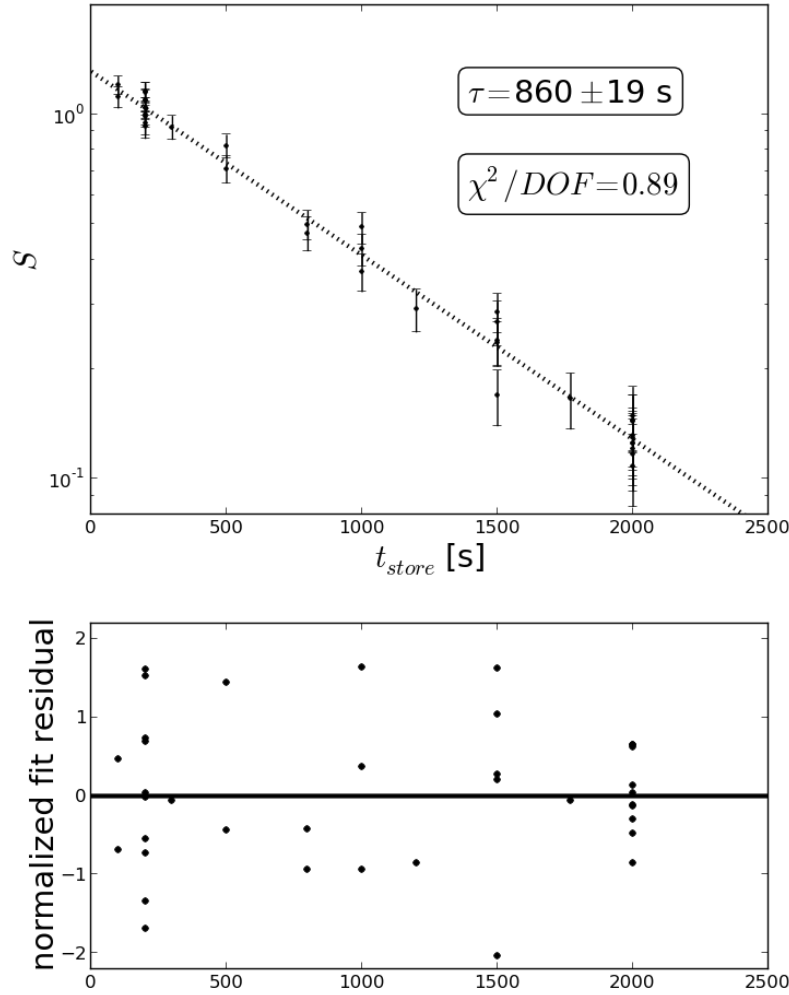


Figure 4.5: The signal S versus t_{store} . The storage time constant of the trap is given by τ_{store} from the exponential fit (upper). The distribution of residuals of the exponential fit are normalized to their statistical uncertainty (lower).

4.4 Cleaner Upscatter Detectors

In this section, the ^3He drift tubes on top of the vacuum jacket (near the cleaner) are analyzed. An appropriate set of runs is identified, and the data are combined to observe the time dependence of the counts during the cleaning process. Runs wherein the trap is filled for a fill and empty measurement with 30 seconds of cleaning were used.

Fig. 4.6 shows the average drift tube rate during filling. The prompt background from the beam pulses is evident while the trap is being filled. Fig. 4.6 shows a typical pulse height spectrum for the drift tubes during cleaning. The data from the ^3He tubes are extracted, and an ADC cut of 1.8 to 3 V is applied. While some events below the full energy peak are also neutron-generated (though deposit less energy due to wall-effect), rejecting data outside of the full energy peak offers improved signal-to-background.

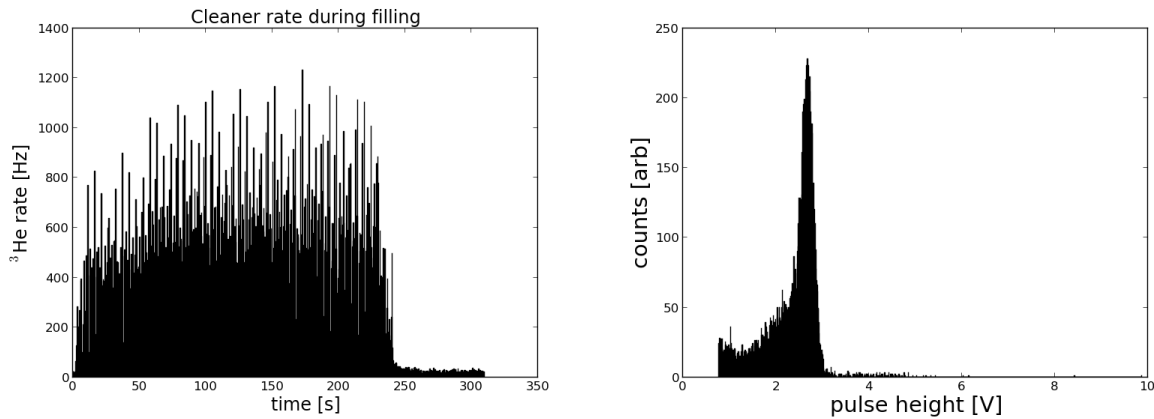


Figure 4.6: The pulse height spectrum of the ^3He tubes (left). The pulse height spectrum of the ^3He tubes (right).

Data for all of the runs are aligned in time, and counts for each run are extracted starting 5 s after the trap door shuts, and ending 50 s later. From the above, we plot the counts over time during the 30 s cleaning cycle averaged over all such runs. This is shown in the left panel of fig. 4.7 along with a $N \exp(-kt) + C$ fit. This gives a time constant $k^{-1} \approx 8$ s. We can perform the same analysis of fill and empty runs where we did not clean. This is shown in the right panel of fig. 4.7, and the signal is not observed.

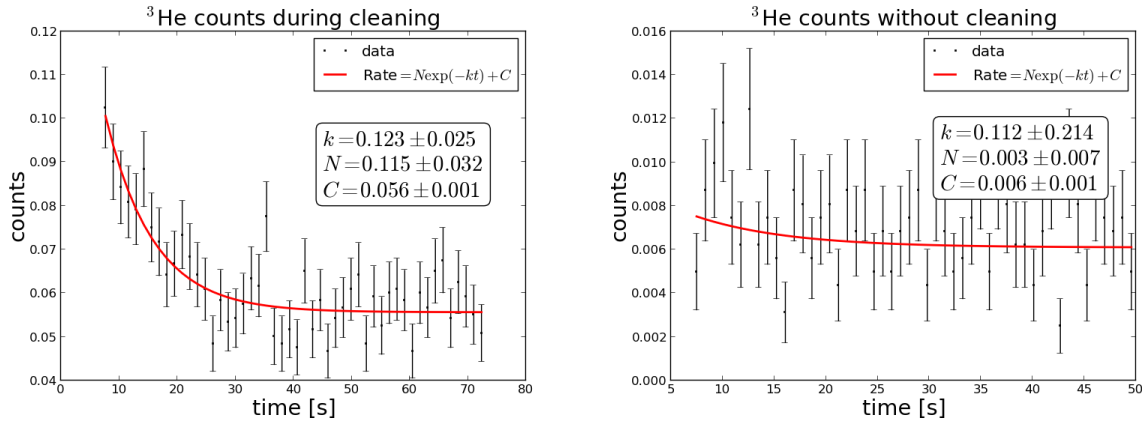


Figure 4.7: The combined counts over time during 30 s cleaning, with exponential fit (left). The combined counts over time for fill and empty runs with no cleaning (right).

In order to truly conclude that the cleaning curve shown in figure 4.7 is due to upscattered UCN in the trap, it will ultimately be desirable to compare runs wherein all valve/actuator timings are equal, *except* that the trap door never opens. Though no such data were taken during the limited time of this run campaign, future experiments are planned to examine the nature of the signal from these detectors, and potentially improve the detector efficiency by placing detectors closer to the cleaner.

4.5 Conclusions

The measured storage time of 860 ± 19 s is within 1σ of current averages of τ_n measurements. This demonstrates that the storage time is long compared to the neutron lifetime, which is a critical for a magnetic-trap-based measurement of τ_n .

However, the (statistical) uncertainty is quite large. The data included in the storage time measurement comes from approximately one day's worth of run cycles. However, in order to perform additional systematic studies in the future, many such cycles must be performed, and improving UCN transport into the trap is necessary to perform these characterizations in a reasonable time-span. In addition, the polarization of the UCN is observed to be poor due to the presence of stainless steel components used in and near the trap door. We estimate $\sim 10^4$ trapped neutrons per fill, but substantially less are detected due to the poor transport efficiency into the

monitoring detector. The statistical uncertainty per measurement can be greatly improved by moving the trap door further out of the way to make for a larger UCN conductance into the trap, replacing the stainless steel components with copper or other non-depolarizing material, and implementing an in situ detector (e.g. the vanadium activation detector) which doesn't rely on high transport efficiency from the trap to an external detector.

In this campaign, data were acquired with the vanadium activation detector, but exhibited very low solid angle efficiency for the NaI detectors and high background rate in both the plastic scintillators and NaI detectors. This is addressed with greater NaI detector coverage and improved shielding, and will be discussed in chapter 5.

For the storage time measurement in this chapter, a determination of the background on a run-by-run basis was sufficient to provide consistent signal data, in part due to the high signal-to-background of UCN counters. However, any changes in shielding geometry, changes in the environmental background (e.g. proton beam spill, activation due to the accelerator), and other effects such as gain drifts, will complicate precision measurements, and it will be necessary to devote dedicated studies to the source of backgrounds and their effect on the measurement of τ_n .

The normalization procedure used here of measuring the saturated rate in the monitoring detector while filling, was sufficient to provide a consistent normalization for each run well within the statistical uncertainty. However, as will be discussed in chapter 6, this may not be sufficient for a precision measurement in the presence of source output fluctuations that are on similar time-scales to the trap filling time.

Finally, the data from the ^3He cleaner detectors shows a time-dependent signal correlated with the trapped UCN. However, there may be other sources of background in these detectors, such as upscattered UCN from the guide system, or accelerator-generated thermal neutrons present in the experimental hall. This motivates more detailed study of these data in future campaigns, as well as separate experiments to study upscattered neutrons from the Cu guides and other sources.

5 Second Experimental Campaign

5.1 Introduction

A second run campaign was performed in a few runs from the fall of 2013 through February of 2014. The goals of this campaign were to commission and characterize the V detector, investigate the stability of monitoring the number of loaded UCN, and improve UCN polarization and transport efficiency. The guide system was initially configured as in the previous run campaign. This configuration was modified during the campaign to investigate new guide configurations to improve the trap density.

A custom waveform digitizer was used to acquire data from all detectors, with 12 bit digital resolution and a 250 MHz sampling rate. This is sufficient to identify the peak from the fast component of the plastic scintillators coming directly from the PMTs. The UCN counters were read out using integrating preamplifiers, and spectroscopy amplifiers with 2 μ s shaping times. Attenuators were placed on the output of the UCN counters as needed to reduce the signal voltage in order to fit within the allowable range for the waveform digitizer inputs.

5.2 Vanadium Detector Characterization

The vanadium detectors were configured as discussed in section 3.8, with two stacks of four NaI detectors surrounding the plastic phoswich scintillator paddles. The NaI PMT signals were summed together into one data acquisition channel, and the plastic scintillators are each read into a data acquisition channel. The gate valve monitor and downstream ^{10}B detector were also implemented.

5.2.1 NaI Characterization

Because the NaI detectors are summed into a single channel, it is important that the gain of each of the 8 detectors is well matched, so that events outside of the region of interest for ^{52}V decay can be excluded. In order to gain match the detectors, a ^{60}Co source was placed in the middle of the detector array, and one NaI detector biased at a given time. The 1.1 and 1.3 MeV γ -ray peaks from the source were identified in each detector, and the PMT bias adjusted to make the peaks appear in the desired location of the pulse-height spectrum. The detectors are then all biased to verify that the spectrum summed over all detectors provides the desired resolution. Figure 5.1 shows the measured pulse height spectrum of ^{60}Co from the NaI detectors.

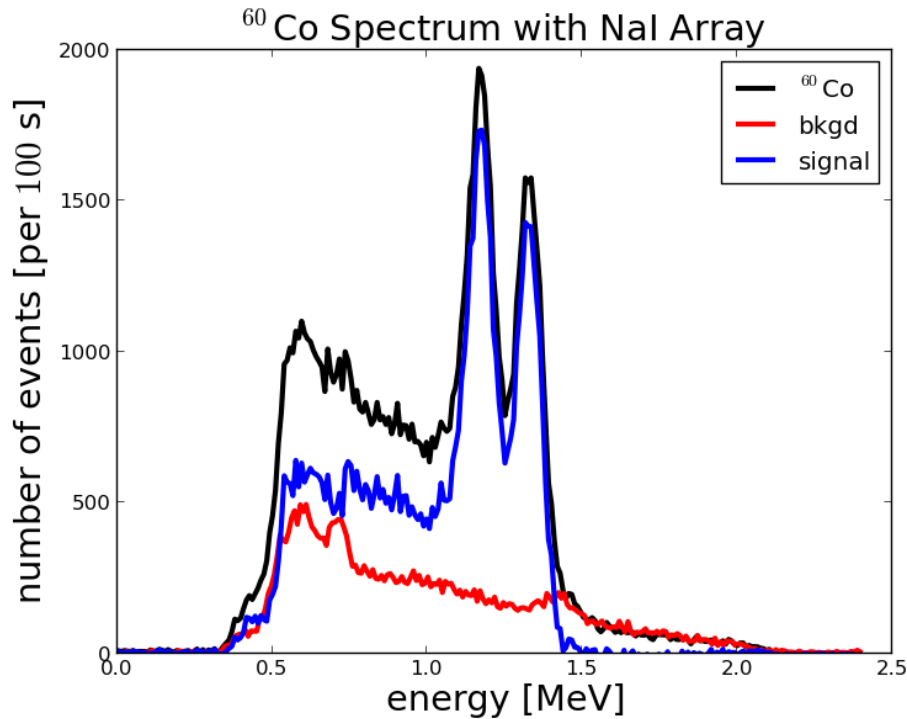


Figure 5.1: The measured ^{60}Co spectrum using the summed NaI detectors. The black shows the measured spectrum, the red shows the spectrum with the source removed, and the blue is a direct subtraction of the foreground and background spectra.

The expected photopeaks and compton edge are readily discernable in the spectrum, as are pileup events from the two photons. There is not a significant $1.1 + 1.3$ MeV full energy peak due to the low solid angle efficiency of a sole detector in the array. From a gaussian fit to

the photopeaks in the ^{60}Co , the FWHM resolution is approximately 10%. Because the goal is to count decays (with the aid of coincidence with the plastic scintillators), energy resolution is not critical aside from being able to reject events with energy much greater than the 1.4 MeV photopeak for ^{52}V . Regardless, the energy resolution could be improved in future run campaigns using spectroscopy amplifiers and longer pedestal sampling for each waveform. In addition, the detectors could be read out independently to obviate the need for careful gain matching.

This gain matching procedure was performed with the holding field both off and on, and two different gain matching routines used for either configuration due to differences in PMT gain with magnetic field. The change in gain between these two configurations is as large as 10-20% for the detectors closest to the holding field coils. Ultimately, the holding-field-on gain matching was used (as no holding-field-off runs were performed in this campaign). The overall gain was set so as to fit events with energy up to ~ 2 MeV. Any events with more energy saturated the digitizer input, and (as they are above the ^{52}V photopeak) were rejected.

Also in fig. 5.1, an additional peak is evident near the ^{60}Co peaks (and hence near the ^{52}V peak). This is readily identified as the electron capture decay branch of ^{40}K : an isotope that is prevalent in the experimental hall.

We can use vanadium activation spectra to investigate the gain stability of the NaI detectors during typical run cycles. As will be discussed in more detail in section 5.2.2, the V detector signal-to-background ratio is considerably improved by forming coincidences between the NaI and plastic scintillators. While measuring the activity of the vanadium, The NaI spectrum for coincident events clearly resolves the ^{52}V photopeak.

Thus, to investigate gain stability, coincident NaI spectra were obtained from a series of four high-statistics vanadium activation runs. The spectra were formed for events during the time shortly after the vanadium was raised into the detector array, to maximize the number of signal events compared to background. In each run, the ^{52}V photopeak was fit to a gaussian distribution of variable height, mean, and width, and the resulting mean taken as the location of the photopeak. The mean versus time (normalized to initial mean) is then fit to a line to




Figure 5.2: The pulse height spectrum of vanadium activation events in coincidence with plastic scintillator events. The blue line shows the gaussian fit within the fit range, denoted by vertical dotted lines (left). The relative peak position from the gaussian fits to the high-statistics vanadium activation data versus time, along with the linear fit (right).

estimate the size of the relative drift in peak position. Figure 5.2 shows a fit to the coincident NaI spectrum, as well as the relative peak position versus time with linear fit. We find the slope of this fit to be $(3.7 \pm 2.0) \times 10^{-4} \text{ hr}^{-1}$.

5.2.2 Coincidences

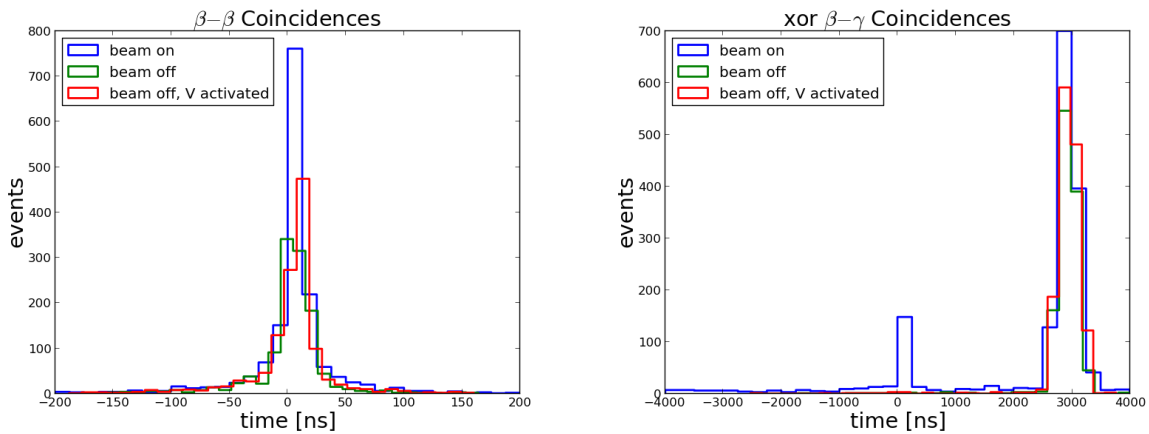


Figure 5.3: The time-to-nearest-event for the beam on period of run 655, activity measurement during that run, and beam off background from run 656, each in a 200 s window. The (left) time-to-nearest-event for the two plastic scintillators and (right) plastic scintillators and NaI detectors are shown.

An ideal candidate β event is one in which exactly one of the plastic scintillators is triggered

(referred to here as an “xor” β event). Coincidences between the two scintillator paddles must be due to cosmic ray or other environmental backgrounds, and the probability of a backscattered β triggering both scintillators is negligible. Further, the xor β event should be coincident with a NaI event with energy not above that of the photopeak. Compton-scattered γ -rays below the photopeak are included.

The time-to-nearest-event between the two plastic scintillators is shown in the upper panel of fig. 5.3 for a time window while the beam is on, beam is off and V foil unactivated, and when the beam is off and activated V foil in the detector array. In the β - β coincidences, one immediately notices an increased coincidence rate with the beam on, and little difference between the background and V activation, suggesting that very few real events are lost by the rejection of coincident events.

The lower panel of fig. 5.3 shows the same runs, but for coincidences between the plastic and NaI scintillators. There is an excess of V activation events in this region compared to background, suggesting that it can discriminate NaI events uncorrelated with events in the plastic scintillator. The peak in the time to nearest event for coincident events is offset between the β and γ signals, most probably due to the comparatively slow speed of the NaI and the use of shaping amplifiers for the NaI. Further, there is a prompt signal well timed with the plastic scintillators, which is likely due to detector noise caused during the movement of valves and actuators while the trap is being filled.

While the time-to-next-event plots offer the desired time correlation between detector events, we want to choose the coincidence and anti-coincidence windows used to select data so as to optimize the V decay signal. We accomplish this by taking a typical vanadium activation run, systematically varying the timing window, and fitting the resulting data to an exponential-plus-background:

$$S(t) = Ne^{-t/\tau_V} + B \quad (5.1)$$

with initial decay rate N and background rate B free parameters, and the ^{52}V mean lifetime fixed to $\tau_V = 324.0 \text{ s}$ [94]. The relative least squares error of N for a given fit is then determined

as a function of the time window. The offset of the time window is fixed to 7.8 ns in the case of plastic scintillator anti-coincidences, and to 3096 ns between the NaI and plastic scintillators. Figure 5.4 shows the optimized fit uncertainties for the β anti-coincidences and β - γ coincidences. In the case of the β - γ coincidences, a later run was used which had better signal to noise, in order to have sufficient statistical sensitivity to perform the optimization. This run consisted of lowering the V foil, opening the trap, and turning the beam on for 400 s, and then raising the foil to measure the activity. The rates in the NaI and β detectors from these runs is shown in fig. 5.5. From this analysis, we take a window of ± 150 ns for the xor β event window, and ± 250 ns for the β - γ coincidences.




Figure 5.4: The relative uncertainty in the activity as determined by the fit to eqn 5.1 for xor β events (left) and xor β coincident with NaI events (right).

5.2.3 Phoswich Geometry Characterization

Figure 5.6 shows example event waveforms from the plastic scintillators. Events with only a slow component, as well as both a fast and slow component. The event shown in black exhibits only a slow decay component, indicative of a charged particle (e.g. a background compton electron) that stopped entirely in the thick, outer layer of the low scintillator. The other event (blue) exhibits both fast and slow signals, and is indicative of an electron which passes through the thin fast layer and stops in the slow layer. In order to investigate the shape of these signals, events




Figure 5.5: The singles and coincidence rates in the detector for the high statistics run (see text).

were fitted to a double exponential function

$$I(t) = I_1 e^{-t/t_1} + I_2 e^{-t/t_2} \quad (5.2)$$

with I_1 , I_2 , t_1 , t_2 free parameters. The intensities $I_{1,2}$ of the fast and slow components varied substantially from event-to-event (as expected), but the time constants were found to reliably be $t_1 = 11$ ns and $t_2 = 285$ s. The slow component agrees well with the quoted decay time of the scintillator, while the fast component is substantially longer than the quoted 2 ns. This is most probably due to the light collection time and response times of the PMTs. After initial study, the time constants $t_{1,2}$ were fixed to the above values, while the intensities remained unconstrained. Figure 5.7 shows a fit to a typical waveform using the above prescription.

Alternatively, the pulse shape can be investigated by separately computing the area under the fast and slow parts of the pulse (also shown in fig. 5.7). Parallel analyses of the data show no appreciable qualitative difference between the fit and integration methods of analysis[93]. As such, the integration method is used to analyse the phoswich data as it is much less computationally intensive.

With this method, we investigate the background rejection of the phoswich signals using the




Figure 5.6: Digitized waveforms from the phoswich scintillators. An event with just a slow component (black) and an event with a fast and slow component (blue) are shown.

vanadium foil activated by UCN. A vanadium activation run with a modified guide geometry for improved signal (run 826, see section 5.4) was examined. Figure 5.8 shows the fast and slow areas for xor β and coincidence events. Events are situated most prominently in two bands in the fast/slow parameter space: those with mostly a fast component, and those with both a fast and slow component. The presence of these bands is possibly due to two scenarios: electrons which stop completely in the fast layer, perhaps due to large incident angle into the fast side of the scintillator paddle, and those which pass through both the fast and slow scintillators. As it turns out, most events that are dominantly slow are rejected simply by setting the trigger threshold sufficiently high. As such, events which originate outside of the two paddles and stop entirely within the slow layer are essentially shielded by that layer.

By comparing the xor and coincidence events in fig. 5.8, one might expect, due to the substantially different signal to background content of these two types of events, that a comparison might reveal a particular place in the parameter space where true ^{52}V events are found. However,




Figure 5.7: (left) A typical digitized waveform (solid) with the associated double exponential fit (dotted line). (right) The same waveform (solid) with vertical dotted lines representing the domain of integration for the fast part of the pulse; all times after the later dotted line form the domain of slow integration.

both branches of events appear to contain signal events for both xor and coincidence events. While in principle one could simply bin events in the two-dimensional fast/slow parameter space and directly subtract histograms of signal and background data, the events are sparse enough that reliable subtraction becomes difficult. Such an analysis was performed, and while it indicates significant signal events throughout the parameter space, it is more tractable to reduce the parameter space to that of one dimension, with a parameter that depends on the pulse shape. Because we aren't concerned with the total energy deposited in the scintillator, and fig. 5.8 suggests that events lie near two slopes in the fast-slow space, the above data are re-binned as a function of the angle of the slope (viz. $\tan^{-1}(I_{\text{fast}}/I_{\text{slow}})$). This is shown in fig. 5.9 for runs 826, as well as run 827, which is a background run in which the GV was never opened to allow UCN into the guides/trap. This shows that real ^{52}V decay events appear in both branches, though preferentially in the fast signal dominated region. Further, most of the background events in the other region are excluded by the coincidence cut, which further shows that indeed both regions are populated by real events.

Ultimately, rejecting data in certain regions of the fast/slow parameter space does not improve the statistical sensitivity of the vanadium activation measurements. While this conclusion may change if signal or background rates are much different in future run campaigns, the initially




Figure 5.8: The fast and slow integrals of events from run 826, a vanadium activation run in a modified guide geometry. The xor β events (left) and xor β - γ coincidences (right) are shown.

considered type of background event (i.e. an event with only a slow component) is rejected simply due to having lower instantaneous light output, thus being below the trigger threshold. Further, many ^{52}V decay events appear to stop in the fast scintillator, suggesting that a revision of the phoswich geometry might change the qualitative nature of events and provide better background rejection. In any case, the tagging of events with the NaI detectors provides excellent background rejection, and a substantial improvement in NaI efficiency may obviate the need for the phoswich geometry altogether.

We can estimate the gain stability of the β detectors with the same high-statistics data used to determine the NaI stability. Unlike the NaI vanadium signal, there is not a well-defined peak that can be identified. Instead, we produce a spectrum $F_r(E)$ of the total integrated pulse area E from the first of the four runs, and treat this as a reference spectrum. Each subsequent run is binned in the same way to form a trial spectrum $F(E)$, and a least-squares fit of $F(gE)$ to $F_r(E)$ is performed with g a free parameter. We then plot g versus average measurement time




Figure 5.9: The angle of the slope of fast and slow intensity for xor β and coincidence events. Run 826 (vanadium activation) is shown in black, and run 827 (background) is shown in blue. In each case, events for times after the vanadium has been raised into the detector array are included.

for the three runs (with the initial reference run at $t = 0$) and fit the data to $g = 1 + rt$ with r free. This is shown in fig. 5.10 and gives a best fit of $r = (9 \pm 6) \times 10^{-4} \text{ hr}^{-1}$ for 32 pulse area bins. However, the reference spectrum may exhibit some statistically allowable deviations in a given pulse area bin which could affect the best fit to later trial spectra. To check this, the above procedure was repeated for different numbers of pulse area bins and for different upper and lower pulse area cuts on the spectrum. This produced $1\text{-}\sigma$ deviations $|r + \delta r|$ as large as $r = 2 \times 10^{-3} \text{ hr}^{-1}$ for finer binning, though with much poorer goodness-of-fit. Thus, the value $r = 2 \times 10^{-3} \text{ hr}^{-1}$ appears to be a reasonably conservative upper bound.

5.2.4 Backgrounds

The primary advantage of using traditional UCN detectors to count surviving UCN is that they exhibit very low (in some cases negligible) background rates, even in proximity to reactor cores




Figure 5.10: The best-fit gain for trial runs compared to the initial reference run versus average measurement time, in hours.

or spallation targets. The background rates in the vanadium detector array are significant, which raises two issues: any estimate of the statistical sensitivity of this technique to a measurement of τ_n must incorporate the size of these backgrounds, and time-dependent contributions to the background can systematically effect the extraction of τ_n .

In addition to measuring the background rates, the nature of background events can be investigated by performing different types of background runs. Here, we analyze three types of background runs: beam-off, beam-on GV closed, beam-on trap closed. In all three cases, all valves/actuators were run in the same way, except for the GV never opening in the beam-on GV closed run, and the trap door remaining closed throughout the beam-on trap closed run. These runs separate a few potential sources of background: prompt background related to the spallation source, spallation-source induced radioactivity, and UCN-generated backgrounds (e.g. radiative capture of UCN on guide surfaces).

Figure 5.11 shows the background rates from these run types, as well as a foreground run with the same timing parameters, for the different types of detector events. The coincidence events

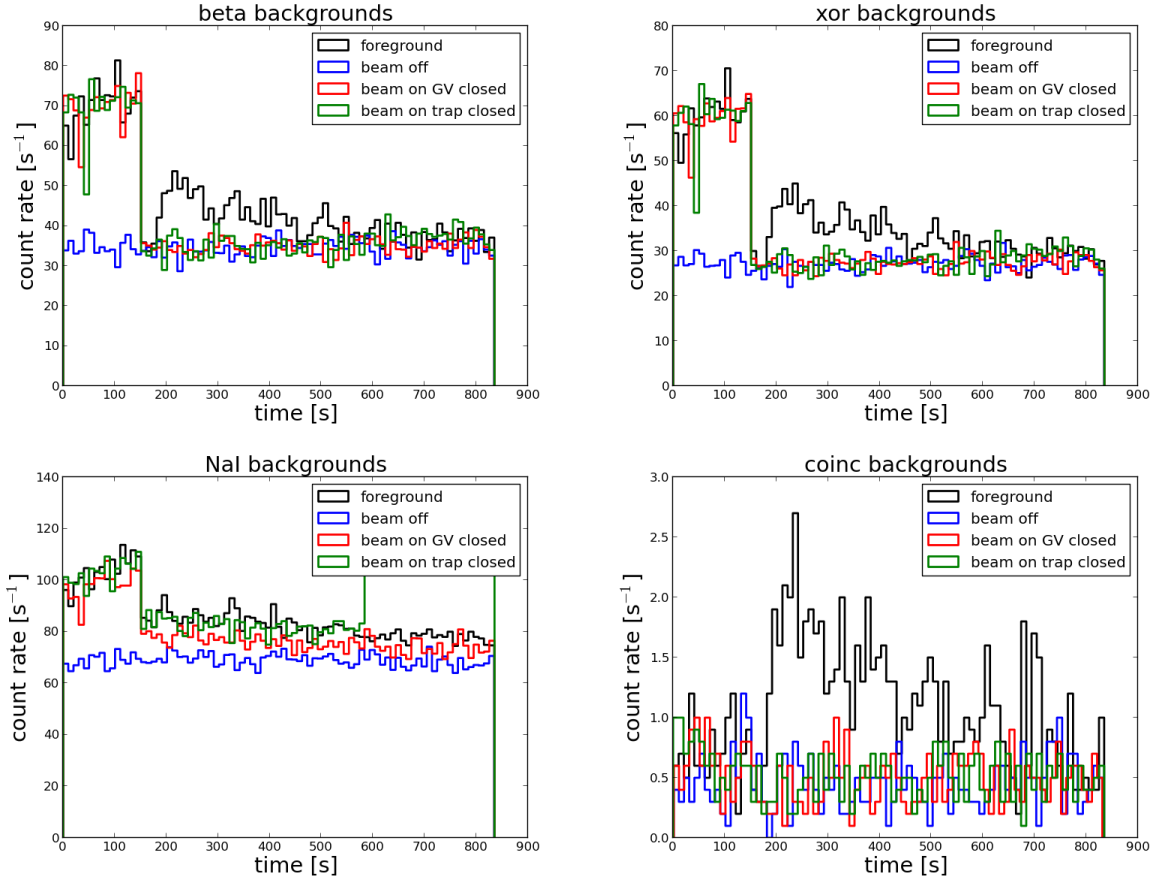


Figure 5.11: The background rates for the different run types. Rates for β singles events (upper left), xor β events (upper right), Nal singles (lower left) and xor β - γ coincidences (lower right) are shown.

and xor β events show no appreciable changes in background for the different configurations. However, the Nal detectors show an increased background while and after the beam was on. In addition, there is a small but significant increase in the Nal background rate for the beam-on trap closed compared to the beam-on GV closed run. This is possibly generated by absorption (or upscattering and subsequent absorption) of UCN within the guide system. This will be studied in more detail in subsequent run campaigns by performing more runs of this type, as well as runs in which UCN are loaded into the trap, but the vanadium foil never lowered. Of interest here, for the purpose of investigating time dependent backgrounds that could cause systematic effects for a τ_n measurement, we study the beam-on Nal detector background and the coincidence background (insofar as it can depend on the background rate in the Nal detector).

The high-statistics data used to extract the vanadium lifetime (see section 5.3) consisted of a 400 s period of beam on, and provide good sensitivity to beam related effects. Figure 5.12 shows a run in which the beam was turned on, but no UCN loaded into the trap. There is a considerable decaying background, which has been identified to be due to activity in the NaI detectors from thermal and fast neutron capture on the iodine, sodium, and aluminum casing of the detectors. The largest source of this background is due to the activation of the iodine, which has a ~ 25 m half-life[56]. Having identified this source of background, the lead shield will be fitted with ^{10}B -loaded rubber to shield the detectors and greatly reduce this background in future measurements.




Figure 5.12: The background in the NaI detector just after the beam was on for 400 s. The red curve is a fit to a constant background plus a decaying component, consistent with the ~ 25 m half-life of ^{128}I .

In spite of this considerable background in the NaI, this time-dependent background is ostensibly eliminated by forming coincidences with the β detectors. Coincident events for this same set of runs were examined to test for the presence of decaying backgrounds or slow drifts in the

background. The background are fit including an exponentially decaying term to incorporate this background, and these fits find no significant contribution. This constrains the initial rate of decay due to such a background to be less than 0.2 s^{-1} for a 25 m half-life. Likewise, we can fit these data to a constant background combined with a slow linear drift, and this constrains the size of such a drift to be less than $1.3 \times 10^{-4} \text{ s}^{-1}$ on top of the background rate of 0.7 s^{-1} . We will discuss the effect of decaying backgrounds and slow drifts in the background in section 6.10.

5.3 The ^{52}V Mean Lifetime

The high statistics runs described above were also used to measure the vanadium lifetime. There were four foreground runs, each followed by a run wherein the gate valve was never opened, which served as a background run. Ultimately, the vanadium decay signal was large enough that the 1000 s counting time was not sufficient for the vanadium activity to diminish by the start of the background run. Because the runs were timed exactly 1400 s apart, the data from the foreground and background runs were instead combined with the appropriate time offset. While the $\beta - \gamma$ coincidences give the best signal-to-background, the signal was not large enough to compensate for the $\sim 1/5$ reduction in efficiency (due to the somewhat low NaI efficiency). Thus, β -xor events were analyzed to determine the ^{52}V lifetime.

The data from a run pair is shown in fig. 5.13. Data are binned with no additional cuts on the pulse height spectrum to maximize the signal. The data are well-described by an exponential plus constant background function (with typical reduced χ^2 of 1 to 1.1). During of the first of the four pairs of the measurements, A trip of the accelerator caused a spurious fluctuation in the background, and therefore the background run could not be used in the fit. The first run of this pair was used, but contributed much less to the overall result because of the lower statistical sensitivity. The resulting measurements of the ^{52}V mean lifetime for the four pairs are shown in table 5.1, and the combined result is $\tau_V = 323.9 \pm 1.1 \text{ s}$, in agreement with ref. [94].




Figure 5.13: An example of ^{52}V lifetime run pairs. The beam-on data is cut from the runs, and the foreground and background combined together with the appropriate time offset between the runs (upper), with the red showing the best fit. The lower panel shows the fit residuals.

As a check for any time dependence in the data, the function

$$f(t) = Ne^{-t/\tau_V} (1 + rt) + B \quad (5.3)$$

with N, τ_V, r, B free parameters, was used. For all run pairs, r is found to be consistent with zero. With increased statistics in future run campaigns, this exercise could be repeated to check for time dependent effects during counting, such as gain/efficiency drifts, time dependent backgrounds, dead time, or pileup.

5.4 Improved UCN Transport

In a separate experiment, UCN transmission and polarimetry measurements were performed on various segments of the guide system as initially configured for these data[92]. These measure-

runs	τ_V	$\delta\tau_V$
831	321.0	3.46
833,834	327.0	1.7
835,836	323.9	1.7
837,838	321.6	1.7
combined	323.9	1.1

Table 5.1: The extracted vanadium lifetimes from each run pair. For the first run pair, the second run could not be used. The values are weighted by the square of their uncertainties and combined to obtain the final result.

ments demonstrated poor transmission and polarization in the guide section nearest to the trap. Moreover, the polarization was unaffected by the location of the trap door magnet plate, the strength of the holding field, or the other magnets in the experimental hall. This suggests a low conductance, high loss, and high rate of depolarization in the guide region immediately near the trap. This is readily explained by the presence of stainless steel guide components, a constricted guide geometry leading to the trap, and potential oil contamination from the actuation of the trap door magnet plate.

Motivated by these measurements, a series of alternate guide geometries were tested. The different geometries are shown in fig. 5.14. The first modification served to test losses due to the (lowered) trap door magnet plate and piston drive: these components were removed, and the bottom of the vacuum chamber replaced with a copper-lined vacuum seal (which we will refer to as the “bare” geometry). Thereafter, a Cu elbow was retrofitted into the vacuum chamber, and brass tube inserted in the upstream section of guide until the beginning of the Cu guide sections, so that no stainless steel was exposed. A short section of electro-polished Cu tube leads up to the bottom of the trap (elbow geometry). A $15.2 \times 15.2 \text{ cm}^2$ polished copper plate with a hole the size of the tube diameter was then fastened to the top to prevent UCN from leaking out of the tube and back into the SS vacuum volume (elbow & plate geometry). Finally, the tube and plate was removed, and replaced with a rectangular Cu box $\sim 15 \times 15 \text{ cm}^2$ wide, and tall enough to reach the bottom of the magnet array (box geometry).

For each geometry, the trap was loaded by turning on the beam, opening the GV for 200 s,




Figure 5.14: The bare (upper left), elbow (upper right), elbow & plate (lower left), and box (lower right) geometries that were tested to improve UCN transport efficiency into the trap. Cu components are shown in red, and brass components in yellow.

the V foil lowered for 30 s, and then raised to count the activity. As in section 5.5, the $\text{xor-}\beta$ signals were counted and normalized to the GV rate to determine the trap density. These runs were also performed with the spin-flipper turned off. The spin-flipper-off runs in principle cause only high-field-seeking UCN to be loaded, which cannot be stored; however, depolarization in the guide system may cause the flipper-off signal to be nonzero, and the contrast S_{on}/S_{off} of flipper on loading to flipper off loading is a relative measure of the polarization of the UCN that reach the trap.

Table 5.2 shows the signal and contrast for each geometry. While the elbow & plate geometry showed the greatest improvement in UCN density, the box geometry showed similar gains, and is a more feasible geometry for making a new UCN loading system. In addition, the data

geometry	signal	contrast
original	1.0 ± 0.06	1.3 ± 0.2
bare	1.4 ± 0.09	–
elbow	2.7 ± 0.2	5.0 ± 0.9
elbow & plate	4.9 ± 0.3	6.8 ± 1.1
box	4.6 ± 0.2	8.5 ± 0.9

Table 5.2: The GV-normalized xor- β signal and flipper-on to flipper-off contrast for each geometry, each normalized to the signal and contrast of the original geometry with the trap door and piston drive present. No flipper-off data were acquired for the bare geometry.

show increased contrast with the introduction of only Cu and brass components, which further demonstrates that the stainless steel is the primary cause of depolarization. This analysis will be used to guide the development of improved UCN guides and an improved loading system for UCN. By also improving the detector efficiencies in future run campaigns, an initial number of $\sim 10^5$ UCN should ultimately be achievable.

5.5 Studies with the Vanadium Foil

An understanding of the characteristic time to absorb the neutrons onto the vanadium foil is important for optimizing run cycles that will be used for precise determinations of the storage time, and ultimately for the measurement of τ_n . If the vanadium foil is not lowered for long enough to sufficiently absorb all of the trapped UCN, UCN phase-space time dependence might complicate the interpretation of the data. In addition, the height of the vanadium foil can be changed in order to probe the distribution of trapped UCN. Here, we analyze the time and height dependence of UCN capture on the foil, which could be repeated in further detail in future run campaigns.

Due to the somewhat low NaI detector efficiency and limited trap density, the xor β events were used to extract vanadium activation signals. This offered the best statistical sensitivity for the amount of activity observed during a typical run. For runs with the initial guide geometry (i.e. that described in chapter 4), the downstream ^{10}B -coated detector was used to normalize the number of UCN loaded into the trap. For the modified geometries studied thereafter, the

gate valve monitor was used to normalize the trap density. The downstream aluminum shutter was always closed to maximize the UCN density in the trap.




Figure 5.15: The extraction of the vanadium signal for a typical run. The time distribution of β -xor events (top) is shown along with the fit to extract the number of vanadium counts. The GV rate (bottom) is shown along with the window within which the normalization rate M is computed (red vertical lines).

In general, the analysis of these runs consists of extracting the xor β events that occur after UCN have been absorbed onto the foil in the trap, and the foil has been raised into the detector array. Events in the monitor detector are selected within a time window wherein the rate has saturated. The xor β events are then fit to eqn. 5.1, and the signal is then defined to be $N\tau_V/M$ with M the average monitor rate. In the case of studies of trap filling time, the monitor event window is varied commensurately with the duration of the filling, such that the window remains within a region of constant rate (see fig. 5.15).

A series of runs were performed wherein the gate valve and trap were opened, and beam turned on, to fill the trap. The above analysis was applied to each run with fixed time windows, and the resulting signal as a function of vanadium absorption time is shown in fig. 5.16. The

data are fit to $A(1 - \exp(-t/t_{\text{drain}}))$ with A and t_{drain} free parameters. The best fit gives $t_{\text{drain}} = 8.9 \pm 2.2$ s.




Figure 5.16: The normalized xor β signal as a function of vanadium draining time in the nominal guide configuration. The best fit for an exponential approach is also shown.

If we suppose that the time dependence of UCN draining into the foil is approximately exponential, characterized by a draining time τ_D , then we must account for the decay of both neutrons and ^{52}V nuclei before the foil has been raised. This necessitates a multiplicative correction proportional to $(1 + \tau_D/\tau_n - \tau_D/\tau_V)^{-1}$, which amounts to a 1-2% correction based on the observed t_{drain} . This is much smaller than the statistical uncertainty in the measurement performed here, though will become important in future measurements. We will examine the effect of this time-dependent draining profile with respect to systematic effects in sec. 6.11.

During the end of the run campaign, with the Cu box geometry in place (see section 5.4), a series of measurements were performed wherein the beam was turned on to fill the trap, and the vanadium foil lowered to different heights. This provides a probe for the vertical density profile of UCN while the trap is being filled.

The measurement was done for four different heights, starting with the nominal "down" position which is ~ 1 cm from the surface of the magnets to the nearest corner of the foil. The trap was filled for 200 s, and the vanadium foil lowered for 30 s. The signal for each height was extracted, and is shown in figure 5.17.




Figure 5.17: The height dependence of the detector signal after filling for 200 s.

Also shown in figure 5.17 are predictions from kinetic theory (see sec. B.2). Model 1 is a calculation of the local density near the foil for trapped UCN with a $\sqrt{E}dE$ spectrum up to 50 neV and up to 100 neV. Model 2 uses the same spectra and maximum energies, but instead plots the relative number of UCN that are energetically allowed to reach the bottom edge of the foil. One might expect the former calculation to describe the measurement if the vanadium absorption time is short enough that only the local density is sampled, and the latter would be expected if the vanadium absorption time is long enough to capture all UCN that can reach the foil. The figure shows that this measurement is likely in between these two regimes, and the measured signal above 50 cm shows that some amount of quasi-bound UCN are introduced into the trap

during filling.

5.6 Discussion and Conclusions

In this chapter, we have addressed the issues encountered in the first run campaign: we have demonstrated the detection of trapped UCN using the vanadium activation method, and assessed means of improving the trapped UCN density. The vanadium activation method offers substantially higher statistical sensitivity, and the β - γ coincidence provides a robust background rejection. While the NaI efficiency (and thus the coincidence efficiency) is somewhat low, this will be straightforwardly improved with the inclusion of more NaI detectors, and a modified detector array geometry which places the NaI detectors closer to the V foil. Simulations with GEANT4 suggest that a NaI efficiency of 50-60% is possible with only modest changes[7]. The signal can also be improved by reducing the neutron-generated NaI background using thermal neutron shielding, thus reducing potential time-dependent backgrounds such as the activated iodine. The plastic scintillator detector backgrounds could be reduced by modifying and optimizing the phoswich geometry, and possibly by changing the light collection geometry to reduce the rate of Cerenkov events in the light guides. This source of background is likely a substantial component of the background events observed in these data with only a “fast” component.

The measurement of the V absorption time constant of $t_{\text{drain}} = 8.9 \pm 2.2$ s hints that the vanadium foil counts surviving UCN much faster than the fill-and-empty technique with this apparatus and in earlier UCN bottle experiments. In the run campaign described in chapter 4, the emptying time was found to be ~ 30 s; however, UCN emptied from the trap are subject to considerable losses while being transported to the detector (or can flow back upstream in the guide system, where they are likely lost), so that 30 s is likely a lower bound.

A naive estimate of the V absorption time constant from kinetic theory is given by

$$t_{\text{drain}}^{-1} = \frac{A \langle v \rangle}{4V} \quad (5.4)$$

with A the absorption area, $\langle v \rangle$ a typical UCN velocity, and V the trap volume. Taking $A = 620 \text{ cm}^2$ to be twice the area of the foil (as it is double sided), $V = 670 \text{ l}$ the physical trap volume, and assuming that the time constant is dominated by the highest velocity component so that $\langle v \rangle \approx 3 \text{ m/s}$, we have $t_{\text{drain}} = 14.4 \text{ s}$. This estimate is somewhat higher than the measurement, possibly because UCN with $v < 3 \text{ m/s}$ occupy a smaller effective volume than the physical trap volume, and because of deviations from kinetic equilibrium (i.e. trapped UCN occupying a smaller volume of phase space than is energetically permitted). Nonetheless, this comparison suggests that the vanadium absorption probability per bounce is not much less than unity. Ultimately, a more rigorous study of this observable can help address the size of phase-space-related systematic effects (see sec. 6.11), and neutron tracking studies are better suited to provide a theoretical estimate.

In conclusion, the data in this chapter suggest that improved statistical sensitivity is achievable, and the detection methods investigated here can be used to study the dynamics of the trap and investigate certain types of systematic effects. These data, in combination with the discussion of systematic effects in chapter 6, provide the motivation for a high-precision measurement of τ_n with the magnetic bottle technique.

6 Systematic Effects

6.1 Introduction

In the previous chapters, we demonstrated that the Halbach array stores neutrons for times comparable to τ_n . In future run campaigns to measure the neutron lifetime, the storage time τ_s will be determined by storing neutrons for two times t_1 and t_2 , counting the surviving neutrons C_1 and C_2 , normalizing this measurement to the initial number of neutrons in the trap M_1 and M_2 , and computing

$$\tau_s^{-1} = \frac{\ln \left(\frac{C_1 M_2}{C_2 M_1} \right)}{t_2 - t_1}. \quad (6.1)$$

Ideally, the normalized signal goes as $C_i/M_i \propto \exp(-\tau_n^{-1}t_i)$ where $i = 1, 2$, because efficiencies and other effects cancel between the two storage measurements, and we have that the measured $\tau_s = \tau_n$. However, if some systematic shift occurs between the two measurements, this causes a shift in the measured storage time. We can fix the normalized signal for the first storage time, and separate the normalized signal for the second storage time into factors which cancel and those that do not – either due to a shift δ_C in the counting or δ_M in the normalization. In addition, the factor of τ_n^{-1} could be modified due to sources of loss in the trap, so that $\tau_n^{-1} \rightarrow \tau_n^{-1} + \tau_l^{-1}$ where τ_l^{-1} is the inverse loss lifetime which principle contains multiple terms for different sources of loss. With these considerations, we let

$$\frac{C_2}{M_2} = \frac{C_1(1 + \delta_C)}{M_1(1 + \delta_M)} e^{(\tau_n^{-1} + \tau_l^{-1})(t_2 - t_1)}. \quad (6.2)$$

Inserting into eqn. 6.1, we have that

$$\tau_n^{-1} = \tau_s^{-1} - \tau_l^{-1} + \frac{\delta_C}{t_2 - t_1} - \frac{\delta_M}{t_2 - t_1} \quad (6.3)$$

which is valid to first order in δ_C and δ_M . The first term is the measured storage time with its uncertainty, and the remaining terms are systematic corrections (with uncertainties) which we can quantify using analytic estimates, simulation, and/or by performing ancillary measurements.

The purpose of this chapter will be to discuss, provide estimates of, and propose measurements of some of the expected systematic corrections due to loss, counting of the UCN, or normalization of the UCN signal. For a given systematic effect, we quantify its effect by estimating its effect on the measured τ_n . Generally, we assume here that the storage time will be measured by pairs of short and long storage measurements, each pair being independent of subsequent pairs. Possible effects of the three types are shown in table 6.1.

loss	counting	normalizing
residual gas	t-dep. bkgd	t-dep. bkgd
material loss	gain drifts	gain drifts
heating	dead time	dead time
cleaning	phase space t-dep.	source fluctuations
depolarization	non-linearity	non-linearity

Table 6.1: Potential systematic effects related to storing, counting, and normalizing.

6.2 Residual Gas

The residual gas in the trap can absorb or upscatter the stored UCN to cold or thermal energies, thus introducing a potential source of UCN loss in the experiment. Recent experiments have corrected for this effect. Arzumanov *et al.* introduce a 0.2 s uncertainty (but no correction) in their determination of τ_n for a trap pressure of 1×10^{-6} to 5×10^{-6} Torr[66, 69]. Pichlmaier *et al.* introduce various corrections depending on the observed vacuum pressure and composition, ranging from 0.7 to 1.4 s. The larger corrections were due to the outgassing of Fomblin oil, which was observed with a mass spectrometer and fine pressure gauge. Thereafter the pressure

was stable at 3.8×10^{-6} torr, and the dominant residual gas that can cause upscattering was water vapor. All corrections quoted a 0.4 s uncertainty due to the location of the measurement devices[23].

The effect was measured by Serebrov *et al.*[27] by measuring the storage time of the trap at two pressures $P = 3.8 \times 10^{-6}$ and $P = 6 \times 10^{-4}$ Torr. The loss rate rate $\tau_g^{-1} \propto P$ is then determined from these two measurements, giving a 0.4 ± 0.024 s correction to the measurement of the neutron lifetime. We can estimate the loss rate by relating it to the correction time

$$t_{corr} = \tau_n - \tau_{store}$$

$$\tau_g^{-1} = \frac{t_{corr}}{\tau_n (\tau_n - t_{corr})} \quad (6.4)$$

from which we find that, for example, a 0.4 s correction implies $\tau_g^{-1} \approx 5 \times 10^{-7} \text{ s}^{-1}$.

Because the upscattering and absorption cross sections σ_{up} and σ_a of a UCN are proportional to $1/v$, the rate

$$\tau_g^{-1} = n(\sigma_{up} + \sigma_a)v = \tau_{up}^{-1} + \tau_a^{-1} \quad (6.5)$$

is to a good approximation independent of the UCN velocity. The time-dependent distortion of the trapped UCN energy spectrum is therefore negligible. However, the rate for different molecular species can be quite different due to the large variance of total molecular scattering/absorption cross sections for typical residual gases. Care must therefore be taken to monitor for changes in the molecular composition of the gas over time.

For a free, monatomic gas (atomic mass M), we can write the double differential scattering cross section

$$\frac{d^2\sigma}{d\Omega dE'} = \sqrt{\frac{E'}{E}} \frac{\sigma_{scat}}{4\pi} \sqrt{\frac{\beta M}{2\pi\hbar^2\kappa^2}} \exp \left[-\frac{\beta M}{2\hbar^2\kappa^2} \left(E - E' - \frac{\hbar^2\kappa^2}{2M} \right)^2 \right] \quad (6.6)$$

where $\kappa^2/2m = E + E' - 2\sqrt{EE'}x$, $x = \cos\theta$ is the cosine of the neutron scattering angle, E (E') is the initial (final) neutron energy, and $\beta = 1/kT$ is the inverse temperature. From this

expression, we compute the total upscattering cross section

$$\sigma_{up} = \frac{\sigma_{scat}}{4\pi} \sqrt{\frac{M\pi\beta}{mE}} \int_0^\infty \int_{-1}^1 \frac{dx dE' \sqrt{E'}}{\sqrt{E + E' - 2\sqrt{EE'}x}} \times \exp \left[-\frac{\beta M (E' - \sqrt{EE'}x)^2}{m (E + E' - 2\sqrt{EE'}x)} \right]. \quad (6.7)$$

As an example, the upscattering rate of argon at $P = 5 \times 10^{-6}$ Torr is $1.1 \times 10^{-9} \text{ s}^{-1}$, and this rate is constant in UCN energy to within a relative deviation of less than 10^{-4} with $0 < E_{UCN} < 50$ neV. For molecular gases, the dynamic structure factor used to compute the cross section in eqn. 6.6 must incorporate rotational and vibrational degrees of freedom in addition to translational motion, which can sometimes be calculated[135]. Regardless of the intra-molecular dynamics, the $\sigma \propto 1/v$ behavior still holds, and we can describe the scattering from a molecular species j by direct calculation or by an experimentally determined rate.

The absorption rate for a given species with cross section σ_a (quoted at thermal velocity $v_{th} = 2200 \text{ m/s}$) is given by

$$\tau_a^{-1} = \sigma_a n v_{th} = \frac{P}{kT} \sigma_a v_{th}. \quad (6.8)$$

For the most prevalent and problematic gas species (primarily water vapor and heavier hydrocarbons) the hydrogen scattering cross section is much larger. The loss rate due to absorption in water at 5×10^{-6} Torr, for example, is $4.7 \times 10^{-9} \text{ s}^{-1}$, which is two orders of magnitude smaller than its upscattering rate.

During the February 2013 experimental campaign, the vacuum pressure was initially $\sim 10^{-5}$ Torr, but continual pumping and flushing with nitrogen brought the vacuum to 7×10^{-6} Torr using a Turbo-V 300HT Macro Torr turbo pump backed by a scroll pump. A residual gas analyzer (RGA) was used to determine the composition of the vacuum, and showed that water vapor was the primary contaminant in the vacuum system, along with some N_2 and O_2 . The nitrogen and oxygen were due to small known leaks in sealing bolts and guide couplers, which can be straightforwardly remedied in subsequent experimental runs.

Further, the pumping speed can easily be increased by an order of magnitude with the addition of a cryo-pump on one of the large flanges at either end of the vacuum jacket. By increasing the pumping speed and flushing with dry gas, we expect that the residual gas upscattering rate can be decreased below 10^{-7} s, thus reducing the residual gas correction to the 10^{-4} level or below. Further, an RGA can be permanently installed on the apparatus in order to assure that the necessary vacuum conditions are not compromised during the measurement.

6.3 Depolarization

The Halbach array and holding field are designed to assure that $|\mathbf{B}| \neq 0$ everywhere in the trap volume. The data in chapter 4 show that without the holding field, UCN are depolarized at a mean rate that is of order τ_n^{-1} . Ultimately, the operating approximation is the adiabatic condition of eqn. A.30, which demands that the off-diagonal components of the spin-dependent Hamiltonian for the trapped neutron are negligible. While the physics in either case are the same, we identify two types of depolarization in the trap: 1) the presence of a low-field region in the trap, where the adiabatic approximation is badly violated, and 2) the slow rate of depolarization due to the small but non-negligible violation of the adiabatic condition. The former effect can be addressed by simply mapping the field within the trap volume and verifying that the holding field is at its nominal strength, and does not cancel anywhere with the Halbach field.

The latter effect is that of quantifying the validity of the approximation, in particular in regions near the trap surface where field gradients are large. This effect was addressed in ref. [62] for the simple 1-dimensional case of a UCN reflecting from the exponentially-varying magnetic field. The depolarization rate was calculated by assuming an asymptotic form for the spatial wavefunction, invoking the WKB approximation to compute the transition rate to the high-field-seeking spin state, and ultimately finding the depolarization rate by computing the probability current that escapes the edges of the trap. This prescription was extended to two spatial dimensions by the authors of ref. [29], which encapsulated the rotating field of the Halbach array as seen by the neutron while bouncing near the surface. The semiclassical approximation was validated in this

work by comparing results to direct numerical integration of the Schrödinger equation.

In the more robust two-dimensional model of ref. [29], the authors average the depolarization rate over an isotropic neutron distribution up to the trappable neutron speed, and find a loss rate 10^{-5} smaller than τ_n^{-1} for holding field strengths comparable to that used in the current experiment. It thus seems reasonable to suppose a $< 10^{-4}$ correction to the measured lifetime due to non-adiabatic spin-flipping, though this will ultimately need to be verified for a high-precision measurement. This can be achieved by placing thin detectors on the surface of the trap, and will only be accessible to the high-field-seeking spin state. Even such a detector with limited surface area can be used to place a bound on the size of the effect, and can be checked by de-tuning the AFP spin-flipper to assure the detector performance.

6.4 Material Losses

In order to mitigate the effect of defects in the magnets during the construction of the trap, copper tape is placed over defects of appreciable size; namely, defects with depth more than ~ 1 mm or an area of more than a few square mm. The total area covered with copper is approximately 42 ± 0.5 cm², though this is an overestimate of the area that would be accessible to trapped UCN: the tape must be anchored around the perimeter of a defect, so that the area covered by material is greater than that of the defect. In addition, a "hole" in the magnetic trap is only caused if the minimum field strength dips below ~ 0.8 T, which only happens across some fraction of the defect's area. In addition, most of the affected areas are closer to the top edges of the trap, where one might expect fewer bounces per second and lower average UCN velocity at those heights.

A naive kinetic theory estimate, assuming a loss-per-bounce $\eta = 10^{-4}$, trap volume of 600 l, 42 cm² area of tape, and neutron speed of 3 m/s, this gives a mean loss rate of 5.25×10^{-7} s⁻¹. For a nominal lifetime of 880 s, this is a 0.4 s correction to the measured storage time. If indeed the area is an upper bound, and if the loss-per-bounce is somewhat lower (Cu guides can be much lower than $\sim 10^{-4}$ per bounce), the correction could in reality be negligible. However,

the UCN properties of this copper has not been demonstrated.

A simple way of quantifying the effect experimentally is to measure the storage time with the copper in place, and with the copper removed. Any difference in the storage times will be a direct indication of material losses. Ultimately, the magnets must be repaired – molds have been obtained of all defects, so that magnets can be shaped to fill the defects, where feasible. In addition, surface profilometry of the tape would rule out the presence of any contaminants which spoil the assumptions made here. For the current experiment, and with the future experimental checks described here, a $+0.4 \pm 0.4$ correction seems reasonable.

In order to verify the UCN compatibility of the Cu for a high-precision campaign, a series of measurements can be performed. The loss-per-bounce can be determined through separate experiments in which UCN are bottled in a material volume containing the copper-coated adhesive used here. Along with this measurement of the loss-per-bounce, the tape within the trap can be replaced with a material that is known to have a high loss-per-bounce, such as polyethylene or polymethylpentene; a storage measurement in this configuration would provide an estimate of the total rate of incidence upon these regions of the trap, and the combination of this rate measurement and the loss-per-bounce measurement will measure or constrain the size of the effect.

6.5 Cleaning Quasi-Bound UCN

The data in section 4.4 suggests that many of the quasi-bound neutrons are absorbed in a mean time on the order of 8 s. However, there is likely non-exponential behavior in the draining of quasi-bound neutrons because of the distribution of UCN energies, and potentially due to non-equilibrium behavior of the trapped UCN.

In a simple model of the initial trapped population N_0 and initial quasi-bound population qN_0 (where q represents the fraction of quasi-bound UCN relative to the trapped number), we can

write the time evolution as

$$N(t) = N_0 \left[e^{-\tau_n^{-1}t} + qe^{-(\tau_n^{-1} + \tau_q^{-1})t} \right] \quad (6.9)$$

with τ_q the mean loss time for quasi-bound UCN once the cleaner is raised, and the UCN storage interval begins. The initial populations are defined to be the remaining number of trapped and quasi-bound UCN after the cleaning interval, and after the short storage time t_1 . In the following, we will set $t_1 = 0$ without loss of generality. For a precision measurement, it must be that q is small (i.e. that much of the quasi-bound UCN are removed during cleaning) and so we take the above, and compute the storage time from the log-ratio of the surviving UCN after no storage, and a long storage time t_2 . To first order in q , we have that

$$\tau_s^{-1} = \tau_n^{-1} + qt_2^{-1} (1 - e^{-t_2/\tau_q}) \quad (6.10)$$

which gives a relative pull on the lifetime of

$$\delta\tau_n/\tau_n = \frac{\tau_n q}{t_2} (e^{-t_2/\tau_q} - 1). \quad (6.11)$$

The parenthetical factor is only small if $t_2 \lesssim \tau_q$. Foregoing any knowledge of τ_q , the correction can be as large as $\tau_n q/\tau_2$, and so it is generally desirable to make the fractional population q as small as possible.

The effectiveness of a cleaner that covers the whole plane of the trap at a constant z was investigated in refs. [62, 4]. These studies showed that such a cleaner leaves $< 10^{-4}$ quasi-trapped fraction remaining after 120 s of cleaning time. In these studies, UCN were initially placed randomly throughout the volume of the trap, with their velocity determined so as to make their total energy that of the trappable energy. However, quasi-bound UCN must necessarily have a trajectory that starts from the location of the trap door, and the initial conditions chosen in that work might skew the occupied phase space of the trap, and skew the relative presence of

uncleaned quasi-bound UCN. In order to investigate this, here we perform a neutron tracking study which consists of filling the trap from the trap door for 60 s with the cleaner in the down position, and then closing the trap door and leaving the cleaner down for 250 s. In this way, only UCN which initially start at the trap door are considered. The UCN spectrum goes as $v^2 dv$ with a lower cutoff such that the UCN are energetically able to reach the top of the trap, and an upper cutoff of 4.5 m/s, which was taken from a simulation of the UCN spectrum that is loaded into the trap[91]. The cleaner in its lowered position is 42 cm from the bottom of the trap, which is similar to the lower height of the prototype cleaner and within the range of cleaning heights studied in ref. [62].

The rate of cleaning of the UCN from the trap after filling is shown in figure 6.1. The UCN are cleaned with a time constant on the order of 1 s, and less than 5×10^{-5} UCN remain after cleaning. This result, as well as the subsequent results in this section, turns out to be largely consistent whether or not the higher-order ripples in the Halbach field are included.




Figure 6.1

To investigate the ability of the current prototype cleaner to reduce the quasi-bound popula-

tion, additional neutron tracking simulations are performed. The simulation consists of loading the trap for 60 s, with trajectories beginning at the location of the trap door, and oriented isotropically upwards in 2π . The cleaner remains down for another 250 s thereafter to clean the UCN. The velocity spectrum $v^2 dv$ is cut off from below and above as with the previous simulation.




Figure 6.2: The rate of incident UCN upon the cleaner versus time. Data are fit to two-exponential functions, and the best-fit for the time constants t_1 and t_2 are shown. The fitted amplitudes of the exponential terms are (respective to the legend) 121.5 and 46.0 in relative units.

Two simulations of this type were performed. First, the dimensions of the prototype cleaner from section 3.6 were used. The trap was loaded for 60 s and cleaned for 250 s thereafter. The rate of cleaning over the 250 s time interval is shown in fig. 6.2 along with a two exponential fit. The shorter time constant of this fit is 6.2 ± 0.6 s, in reasonable agreement with the upscattered neutron measurements with the cleaner from section 4.4.

However, to quantify the effect of quasi-bound neutrons using this cleaner, the number of remaining quasi-bound neutrons after the cleaning interval must be extracted. In fact, from this simulation, 1 – 10% of the quasi-bound neutrons remain. Even correcting for the smaller relative population of quasi-bound neutrons after filling the trap, this would lead to a significant effect at the 0.1% level of measurement. Figure 6.3 shows an example trajectory that is not cleaned




Figure 6.3: A trajectory that is not cleaned by just the prototype cleaner. The black line represents the spatial trajectory, and the red box outlines the position of the cleaner.

by the prototype cleaner alone. Clearly, it reaches a height well above the cleaning height, but has a momentum primarily directed along the x direction. Even for long times, this trajectory's momentum does not shift appreciably out of the x - z plane, thus avoiding the cleaner.

In forthcoming run campaigns, an additional polyethylene sheet will be included in addition to the prototype cleaner to rapidly clean higher energy quasi-bound UCN. The sheet will occupy the side of the trap opposite that of the cleaner (i.e. the $y < 0$ side), covering nearly half of the trap from the furthest $y < 0$ edge to approximately 7 cm from the center of the trap along $y = 0$. This geometry was modeled in the same way as the other geometries, and the study shows that $q < 6 \times 10^{-4}$ after ~ 200 s. The cleaning rate is shown in fig. 6.4.

In order to translate this into a systematic uncertainty, we must take into account the fractional population of quasi-bound UCN before the cleaning interval begins (that is, the value of q when the cleaning interval begins). Figure 6.5 shows the initial spectrum, as well as the spectrum after the filling interval is complete and cleaning interval begins. Quasi-bound UCN are initially 2.19




Figure 6.4: The rate of cleaning using the prototype cleaner and additional polyethylene sheet. No UCN remain after ~ 200 s, suggesting that less than 6×10^{-4} . The fitted amplitudes of the three exponential terms are (relative to the plot legend) 681.6, 334.4, and 30.2, in relative units.

times more abundant than trappable UCN, but these UCN are cleaned during filling, reducing the population to 0.29 of that of the trappable UCN. The simulation thus suggests that, even conservatively, $q \sim \mathcal{O}(1)$ before cleaning, and less than 6×10^{-4} after cleaning, which sets a like bound on the size of the systematic shift in τ_n due to this effect.

6.6 Cleaner-Generated Effects

The use of a flat, horizontal pad of UCN absorbing or upscattering material is a conceptually and experimentally straightforward way to eliminate marginally-trapped UCN from the UCN τ storage volume. However, two potential complications arise. First, the number of trapped neutrons depends on the cleaning height. If this dependence is strong, it demands reproducibility of the movement of the absorber for each fill, clean, and storage cycle: large variations in the height can




Figure 6.5: The initial spectrum entering the trap (blue) and spectrum after the trap has been filled (red). The vertical black dotted line corresponds to the maximum trappable UCN energy.

change the number of trapped neutrons relative to the signal from the UCN monitors, thus making the normalization of each measurement unreliable. Second, the trappable neutrons absorbed by the cleaner could be removed with a much longer time constant than the quasi-bound neutrons. The number of trapped neutrons could then be sensitive to the cleaning time, and any variation or uncertainty in the time could spoil the normalization.

To quantify these effects, we simulate the cleaning of UCN at slightly varying heights, and analyze N_t with time for each height. A fractional difference between two simulations 1 and 2 of $(N_t^{(1)} - N_t^{(2)})/N_0$ (where N_0 is the initial number of UCN in the simulations) provides an indication of the size of the systematic effect on the ability to determine the initial number of UCN in the trap.

UCN are cleaned using the neutron tracking code described in chapter B. Three cleaning heights (42.0, 43.5, 44.0 cm) are used in separate simulations. For each simulation, $N_0 = 10^5$

neutrons are placed in a rectangular region near the bottom of the trap corresponding to the profile of the trap door. Their horizontal position is distributed uniformly throughout the rectangular region, and their momenta are distributed isotropically in 2π facing vertically upward with a velocity distribution $\rho(v) \propto v^2$ up to $v = 4$ m/s. Trajectories are integrated until the cleaner is reached, or until 100 s have elapsed.

From each simulation, the time of capture on the cleaner for all UCN is binned. The cumulative sum of this quantity is equal to the number $N_c(t)$ of UCN cleaned versus time. Fig. 6.6 Shows $N_0 - N_c(t)$ (the number of survivors) for each simulation.




Figure 6.6: Number of surviving UCN for the three cleaning heights.

Fig. 6.7 shows the number of survivors versus cleaning height (normalized to the 44 cm simulation), with statistical uncertainties. A linear fit to the relative survivor number versus time tells us that, for cleaning heights near ~ 43 cm, we have a fractional change in trapped UCN of 0.2% for every millimeter change in cleaning height.

We necessarily expect to eventually clean almost all of the UCN above h_{trap} . As long as the density $\rho(z)$ of UCN in the trap at a height z does not change rapidly after some time, we can




Figure 6.7: The relative number of trapped UCN after 100 seconds for the three different cleaning heights. The black line is a linear fit.

estimate the number of trapped UCN after cleaning via

$$N_t(h_{clean}) = \int_0^{h_{clean}} \rho(z) dz \quad (6.12)$$

If we assume that, in the region of 42 to 44 cm, the change in N_t is linear (which we assumed in the previous section as well), we have that the relative change m in N_t versus cleaning height is

$$m = \frac{1}{N_t(42)} \frac{d}{dh_{clean}} N_t(h_{clean})|_{42} = \frac{\rho(42)}{N_t(42)} \quad (6.13)$$

in units of centimeters.

From a simulation of trapped UCN without cleaning, we can bin the heights of the UCN (and normalize to unity) to extract $\rho(z)$. UCN are distributed randomly throughout the trap initially, with a random momentum distribution upwards in 2π and initial energy chosen such that the spectrum goes as $v^2 dv$.

Most importantly, the UCN spectrum is cut off corresponding to the height of the trap, so that




Figure 6.8: Vertical linear density (normalized to unity) extracted from simulation without cleaning, after 100 s.

there are no marginally-trapped UCN. We therefore expect the number of UCN at heights near the top of the trap to be relatively low, and therefore $\rho(z)$ to be somewhat lower at those heights than the true density profile that includes quasi-bound UCN. The distribution versus height is shown in figure 6.8. From this, we have that $m = 0.9\%$ per mm.

The second effect can be investigated by looking at how N_t varies with time for the different simulations. For all three cleaning heights, the rate of cleaning of UCN is similar after sufficiently long times. Fig. 6.9 shows the rate that UCN are absorbed versus time. That the rates are similar after only ~ 10 seconds suggests that *marginal* UCN are cleaned with a fairly short time scale, while trappable UCN above the cleaner height are cleaned with a long time constant. That is, no matter the cleaning height, there are always trappable neutrons which can reach $h_{clean} + \epsilon$ (for some small ϵ) that are cleaned with a long time constant.

However, as long as the cleaning time is repeatable, approximately the same number of “sacrificed” UCN will be removed on each fill. In order to quantify this, we consider a cleaning time of 40 s. After this time, the *rate* of absorption of UCN for all three cleaning heights is no more than $\sim 8 \text{ s}^{-1}$. If the cleaning time has an uncertainty of δt , the absolute uncertainty in the




Figure 6.9: The rate of absorption of UCN on the cleaner for 3 different cleaning heights.

number of UCN remaining in the trap is of order $(8 \text{ s}^{-1}) \cdot \delta t$. The number of remaining UCN is in the range of 3.4×10^4 to 3.7×10^4 , which gives relative uncertainties in N_t in the range of 2.2×10^{-4} to 2.4×10^{-4} per second of uncertainty. The uncertainty in the actuation of the trap door is also relevant, as it defines when the cleaning starts.

The results of this section indicate that if the timing control of the cleaner is of order 1 second, trap filling is reliable at the 10^{-4} level. Further, the above results hint that sub-millimeter control of the cleaning height is needed to fill accurately at the 10^{-4} level. Control at the millimeter level may introduce a $\sim 10^{-3}$ shift in the stored number: the current cleaner is pneumatically actuated with mechanical stops to set the position of the cleaner. It is likely that this is reliable to the level of 0.5 mm, which constrains the effect to 5×10^{-4} relative to τ_n , and the position can be monitored in future experimental campaigns to assure at least this level of stability.

There are two additional techniques that could be used to mitigate these effects. The first is to correct for discrepancies in the cleaning height using the number of upscattered UCN detected in the ^3He drift-tubes, or by turning the absorber into an active detector of UCN. If, for example, the cleaning height is discrepant for two fills, this could in principle be observed as a change in

the number of cleaned UCN that are detected. Another method is to use a separate cleaning volume, as opposed to an *in situ* cleaner. Filling cycles would be performed over a time that is long enough for equilibrium to be achieved between the cleaning and trap volumes. In this way, the absorbing surface need not be actuated.

6.7 Microphonic Heating

6.7.1 Introduction

One potential loss channel is that of small energy transfer to a trapped neutron which can subsequently render it energetically untrappable. In this way, the UCN could be slowly “evaporated” away. For both material and magnetic traps, microphonic vibrations of the trap may induce such small energy transfers.

The effect of vibrations on material traps has been discussed in the context of material storage volumes[86], though not usually in the context of precision measurements of the neutron lifetime (one exception being Mampe et al[75]). In a material storage volume, a UCN can either gain or lose velocity upon bouncing from a material wall depending on the wall's instantaneous velocity $v_{wall} \propto \omega A$ (where ω and A are the angular frequency and amplitude of the vibration, respectively). In this way, the UCN can be seen as taking a random walk in velocity with each bounce, so that the net change in velocity is proportional to $\omega A \sqrt{n}$, where n is the number of bounces. However, it is not necessarily reasonable to expect that the same behavior would also occur in magnetic traps.

In this section, the trajectories of UCN in a simplified 1-D model of the trap are investigated. The model introduces a time-dependent modulation of the potential to simulate the effect of microphonic vibrations. The size and frequency of these vibrations are then investigated using an accelerometer in the UCN experimental area to assess the potential severity of the effect.

6.7.2 Simplified Model

In the simplified model, motion is 1-D in the vertical (z) direction, the Halbach array is a horizontal plane, the holding field is neglected, and the higher order ripple terms in the Halbach array field are neglected. With no ripple, the magnitude B of the static field in a linear Halbach array of period λ is

$$B(z) = B_0 e^{-kz}, \quad (6.14)$$

where $k = 2\pi/\lambda$ and B_0 is the field at the surface of the array. When the entire Halbach array is assumed to move up and down vertically in a sinusoidal fashion with amplitude A , angular frequency ω , and phase δ , the field becomes time-dependent and takes the form

$$B(z, t) = B_0 e^{-k[z + A \sin(\omega t + \delta)]} \quad (6.15)$$

Assuming that the adiabatic condition is satisfied, the 1-D equations of motion for a field-repelled neutron in this field, with gravity, are

$$\frac{dv_z}{dt} = \frac{k|\mu|}{m} B_0 e^{-k[z + A \sin(\omega t + \delta)]} - g \quad (6.16)$$

$$\frac{dz}{dt} = v_z \quad (6.17)$$

Since we now have a time-dependent potential, the numerical integrator used for tracking in the static potential has to be modified to preserve their symplectic properties. Both modified leapfrog and fourth-order time-dependent symplectic tracking algorithms were used here, with satisfactory agreement (see for example the time-dependent prescription in ref. [117] as applied to the optimized fourth-order integrator in ref. [110]). An initial set of tracking studies was done to establish the magnitude of the effect. Three cases were considered, with $\lambda = 0.02$ m, $B_0 = 1$ T, $A = 10^{-4}$ m, $f = 100$ Hz ($\omega = 2\pi f$), with initial oscillation phases of 0, $\pi/4$, and π . The neutron was dropped from $h \equiv z(0) = 0.3$ m with zero velocity, and tracked with the fourth

order method with a 10^{-8} s time step.




Figure 6.10: Plots of z vs. t for various phases. Red: 0; blue: $\pi/4$; green: π .

Plots of z vs. t for the three phases are shown in Fig. 6.10. Note that for phases of 0 and $\pi/4$, the energy (and the maximum bounce height) is smaller than the initial energy after 1 s, but for π , the final energy is greater. Figure 6.11 is a plot of E/m_n for the trajectories of Fig. 6.10. From this we see that energy transfers of 2% can be observed after one bounce. Though the chosen amplitude of $A = 100 \mu\text{m}$ is somewhat gratuitous, it suggests that a more careful sweep of the parameter space would be beneficial to establish the size of typical energy transfers.

6.7.3 Parameter Sweep

The parameters h , f , A , and δ are individually varied (with other parameters fixed) to investigate the relationship between the energy transfer per bounce of the UCN. The salient features of the model are:

- The energy transfer is roughly linear in E_0 .
- The energy transfer obeys $\epsilon \propto \sin(\delta + \phi)$ for some ϕ with all other parameters fixed.




Figure 6.11: Plots of E vs. t for various phases. Red: 0; blue: $\pi/4$; green: π .

- The energy transfer is roughly linear in A .
- The energy transfer is maximal for some $f \sim 0 - 100$, and drops off rapidly at high and low f .

This behavior is shown in figs. 6.12, 6.13, and 6.14. For example, there is a resonant behavior as a function of frequency, perhaps unsurprisingly different from the frequency dependence of this effect for material walls. Frequencies above the few hundred Hz range induce little change in neutron energy, even for large amplitudes.

6.7.4 Spectral Random Walk

With the above behavior established, we can estimate the effect of these small UCN energy transfers by randomly walking the initial stored spectrum in a UCN trap, and observing the fraction x of the spectrum that exceeds the maximum trappable energy. In the experiment, we can remove UCN that reach 6–8 cm from the top of the trap. In this study, we conservatively take the height of 6 cm, so that the greatest UCN energy is 45.1 neV. Based on full 3-D simulations




Figure 6.12: Energy transfer after 1 bounce as a function of δ to a UCN with $h = 0.44$, $f = 40$, $A = 10^{-5}$.

of the trap, A typical number of bounces is roughly $N = 3000$ over 1000 seconds of storage.

For the random walk computation, we take an initial spectrum of 3×10^5 UCN, with energy distributed as $\rho(E) \propto \sqrt{E}$ up to $E = 45.1$ neV. Using the observed behavior from section 6.7.3, we proceed by shifting each neutron N times as per:

$$E_{i+1} = E_i \pm \epsilon \sin \eta \cdot \frac{E_i}{45.1 \text{ neV}} \quad (6.18)$$

for each energy. After N iterations, x is computed by counting the number of final energies $E_N > E_t$ where $E_t = 50.3$ neV. A value of $\epsilon = 0.06$ gives $x = 1.5 \cdot 10^{-4}$. The initial and final spectra are shown in fig. 6.15.

From the above, we find the largest value A which satisfies $\epsilon < 0.06$ for a few different frequencies. Some measured values are $A \approx 8 \cdot 10^{-6}$ at $f = 40$, $A \approx 4 \cdot 10^{-5}$ at $f = 100$, and $A \approx 3 \cdot 10^{-4}$ for $f = 150$. This, then, is an estimate of the size of vibrations at these frequencies over 1000 seconds that could evaporate more than 10^{-4} of the UCN.




Figure 6.13: Energy transfer after 1 bounce as a function of A to a UCN with $h = 0.44$, $f = 40$, $\delta = 2.0$.

6.7.5 Microphonic Vibration Measurements

Here, the amplitudes from section 6.7.4 are compared to observed vibrations in the UCN experimental area at LANSCE. The presence of, for example, vacuum pumps and compressors may induce large vibrations at low frequencies. To investigate, an STMicroelectronics accelerometer (sensitivity $50\mu\text{g}/\sqrt{\text{Hz}}$) is fastened firmly to various equipment in the area, including, for example, a vacuum jacket on the UCN beamline used for neutron Electric Dipole Moment (nEDM) research and development.

The device is calibrated to read out acceleration a versus time independently along three axes. From the acquired data $a(t)$ we can readily compute the vibrational spectral density (VSD) along a particular direction. Fig. 6.16 shows the VSD along one direction on the nEDM vacuum jacket.

The mean amplitude of oscillation at a given frequency is given by

$$\bar{A} = 2I / (2\pi f_0)^2, \quad (6.19)$$




Figure 6.14: Energy transfer after 1 bounce as a function of f to a UCN with $h = 0.44$, $\delta = 2.0$, $A = 10^{-5}$.

where I is the area under a given peak in the VSD, and f_0 is the central frequency of the peak. A few of the most intense low frequency peaks are identified; $\bar{A} = 16 \mu\text{m}$ at 30 Hz is the largest observed amplitude (measured on the vacuum manifold connecting to the jacket), with more typical values of $\bar{A} = 1$ to $5 \mu\text{m}$ measured directly on the vacuum jacket.

6.7.6 Discussion

While the simple model in section 6.7.2 captures some of the qualitative behavior of trapped UCN in the presence of microphonics, the 3-D UCN trajectories are more complex, and the presence of quasi-stable orbits could exacerbate the effect. It is also important to note that the escape time for a UCN with energy $E > E_t$ may not be short compared to the 1000 s storage time, and this is not considered in the above analysis.

The results of section 6.7.5 give a sense of the size and frequency content of microphonics in the experimental area. It is, however, unclear how these vibrations translate to that of the actual




Figure 6.15: The initial (upper) and final (lower) spectra after $N = 3000$ bounces.

trap walls, as this depends on geometry and the mechanical coupling of the trap to the rest of the apparatus. This suggests that *in situ* measurements should be done, so that problematic sources of vibration can be identified.

Finally, the *loss* of neutrons due to vibrational heating is only one potential systematic effect: insofar as UCN detection efficiencies are usually velocity dependent, a change in the UCN spectrum during storage folded with the efficiency of detection at different storage times may induce non-exponential behavior. The subtlety of this issue demands that it be addressed with respect to particular detection methods.

6.7.7 Conclusions

A simple model is presented to investigate the trajectories of UCN in the UCN τ trap when subjected to small microphonic vibrations. The model suggests that vibrations of amplitude ~ 10 μm at frequencies below 200 Hz may be able to heat one in ten thousand UCN to an energy above the trap potential over 1000 seconds. Vibrations of similar amplitude and frequency can be observed in the vibrational VSD data acquired from the UCN experimental area at LANSCE.




Figure 6.16: The VSD along a particular axis of the vacuum jacket on the UCN beamline at LANSCE.

By monitoring the vibrations of the trap itself, we can estimate the severity of the effect, and ultimately mitigate the effect by damping the vibrations and/or cleaning to a lower height, so that trapped UCN do not have sufficient time to “evaporate” while being stored.

6.8 Gain Drifts

The vanadium activation detector operates by counting pulses above a threshold. Because of the continuous spectra from the plastic and NaI scintillators, the observed count rate is sensitive to the detector gain relative to the (typically fixed) threshold. Therefore, gain drifts can induce a time-dependent change in the survivor counting efficiencies, systematically affecting the storage time measurement.

We can estimate the size of this effect as follows. We write the observed spectrum from a

detector as a product of pulse-height and time-dependent parts:

$$\frac{dN}{dE} = R_0(t)F(gE) \quad (6.20)$$

where R_0 represents the total event rate, and F is the unit normalized pulse height (or energy) distribution. The factor g represents the gain, which we model here as having a linear drift, so that $g = g(t) = 1 + rt$ with r the fractional gain change per unit time. The measured count rate in such a detector is given by

$$R(t) = R_0(t) \int_{E_1}^{E_2} F(gE) dE \quad (6.21)$$

with E_1 and E_2 the lower and upper thresholds, respectively. To first order in r , this expression reduces to

$$R(t) = R_0 (1 + Xrt) \quad (6.22)$$

where X is given by

$$X = E_2 F(E_2) - E_1 F(E_1) - 1 \quad (6.23)$$

The value $X \sim -1.4$ is a reasonable approximate value for the boron detector and plastic scintillators, and derived from the measured detector spectra from the previous chapters. This factor is somewhat smaller for NaI coincidences due to the greater separation of the ^{52}V photopeak from the low energy threshold. The signal S from a storage cycle is given by the counted survivors:

$$S = \int_{t_e}^{t_f} R(t) dt. \quad (6.24)$$

where, R is the rate in the survivor counting detector. To first order in the drift rate r for the survivor counter we have

$$S = S_0 + Xr \int_{t_e}^{t_f} t R_0(t) dt \quad (6.25)$$

where S_0 is the zeroth order predictions for the number of survivor and monitor counts. We

can compute the integrals assuming a survivor counting rate which takes the form $R(t) = \lambda_v N_{UCN} \exp(-\lambda_v(t - t_e))$, giving

$$S = S_0 + \frac{N_{UCN} X r}{M_0 \lambda} [1 + \lambda t_e - (1 + \lambda t_f) e^{-\lambda(t_f - t_e)}]. \quad (6.26)$$

In order to estimate the effect of gain drifts, we take the expression for S for a long and short storage and calculate the storage time to first order in r . The second storage cycle would start at a time t_w after the first, and the expected signal would be given by the above, replacing the factor in the integrand with $t \rightarrow t + t_w$. Computing the discrepancy between the neutron lifetime τ_n and the measured storage time gives

$$\tau_s^{-1} - \tau_n^{-1} = \frac{X r t_w}{t_2 - t_1} \quad (6.27)$$

or

$$\tau_s - \tau_n = \frac{X r t_w \tau_n^2}{t_1 - t_2} \quad (6.28)$$

In the above, we have taken $\exp(-\lambda(t_f - t_e)) \ll 1$, which only contributes to a few percent of the total first order shift. There is a shift in the signal S due to the gain drift over the course of the activation measurement, but this effect appears at second order, and so we are left with the dominant effect due to the waiting time t_w between measurements. Taking $X = -1.4$ (i.e. assuming that the plastic scintillators cause the largest effect), a drift of $\pm 2 \times 10^{-3} \text{ hr}^{-1}$, a long storage time of 2000 s, and a waiting time of 2300 s, gives a pull of ± 0.7 s on the measured lifetime.

6.9 Dead Time and Pileup

The dead time associated with measuring the vanadium activation can lead to a fraction of missed events which depends on detector rate. Therefore, the missed fraction of events for a short storage

time versus a long storage time could be different, leading to a systematic shift of the extracted storage time. The probability of accepting an event in the detector with event rate r is given by the area under the time-between-event distribution for all times greater than the deadtime t_d , which is $\exp(-rt_d)$. We can suppose $r = r(t) = N_{UCN} \exp(-\lambda_v t)$, and integrate over this rate weighted by the detection probability to find the measured signal versus N_{UCN} and t_d . We have that $N_{UCN} = N_0 \exp(-t_{store}/\tau_n)$, and compute the storage time $\tau_s^{-1} = \ln(N_1/N_2)/(t_2 - t_1)$ using the deadtime-diminished signal, and expand to first order in t_d , from which we find

$$\tau_s - \tau_n = \frac{\lambda_v t_d N_0 \tau_n^2}{2(t_2 - t_1)} (e^{-t_1/\tau_n} - e^{-t_2/\tau_n}). \quad (6.29)$$

Assuming an initial rate of 300 s^{-1} , a long storage time of 2000 s, and a dead time of $2 \mu\text{s}$, this gives a pull of 0.1 s on the measured lifetime. Further, this effect can be corrected for with high precision, and ultimately made negligible, as in precision studies of $0^+ \rightarrow 0^+$ nuclear decays[128].

6.10 Time Dependent Backgrounds

The analysis performed in the preceding chapters generally assumes a constant background for the monitor and vanadium detectors. The high signal-to-background and stability of the monitor detectors is likely to be sufficient for a precision measurement of τ_n . However, the backgrounds in the vanadium are significant, and the extraction of the surviving number of UCN is sensitive to changes in the background on the scale of minutes or hours. This could be due to short-lived isotopes generated by the spallation source. Slow drifts could be due to, for example, the slow diffusion of radioactive gasses through the experimental hall.

For the case of a decaying background after the trap has been loaded, we can write the rate in the vanadium detector as

$$R e^{-t/\tau_v} + I e^{-t/\tau_b} + B \quad (6.30)$$

where τ_b is the mean lifetime of the source of background. The initial time dependent background rate is I , B is the constant part of the background, and R the initial rate of true ^{52}V decays.

Computing the storage time for a pair of short and long storage times t_1 and t_2 with a vanadium counting time T_M , the first order shift in τ_n^{-1} is given by

$$\tau_s^{-1} - \tau_n^{-1} = \frac{I\tau_b}{R\tau_v(t_2 - t_1)} \left(\frac{1 - e^{-T_M/\tau_b}}{1 - e^{-T_M/\tau_v}} \right) \left(e^{(\tau_n^{-1} - \tau_b^{-1})t_1} - e^{(\tau_n^{-1} - \tau_b^{-1})t_2} \right). \quad (6.31)$$

A slow drift in the background is given by the limit of $\tau_b^{-1} \rightarrow 0$. If we express a slow (linear) drifting background as a fraction of the constant background $B(1 + rt)$, we have

$$\tau_s^{-1} - \tau_n^{-1} = \frac{BT_M^2 r}{2R\tau_v(t_2 - t_1)} \left(\frac{e^{t_1/\tau_n} - e^{t_2/\tau_n}}{1 - e^{-T_M/\tau_v}} \right). \quad (6.32)$$

From section 5.2.4, we have constrained the size of I to be $\lesssim 0.2 \text{ s}^{-1}$ for the dominant decay time of $\sim 2100 \text{ s}$ (i.e. the mean lifetime of ^{128}I). From the above, if we count 5×10^4 UCN in the vanadium detectors during a short storage interval, and with a long storage time of 2000 s and nominal neutron lifetime of 880 s , the effect is constrained to 9×10^{-3} . This is not sufficient for a 0.1% measurement of τ_n , however, the ^{10}B -loaded rubber shield will reduce the intensity of thermal neutrons by $\sim 50 - 100$, which will drive a commensurate decrease in the initial rate I of any induced activity. Further, a detailed statistical analysis shows that such a background can be accounted for in the analysis of the data, which ultimately exacerbates the statistical precision by approximately a factor of 4[15]. Thus, the shielding is likely to eliminate the effect of this background, and at worst could be accounted for with sufficient reliability. If, then, 10^{-4} statistical precision of background measurements over the run campaign can be achieved this will contribute a relative effect of approximately $(-4 \pm 4) \times 10^{-4}$ on τ_n .

As also discussed in section 5.2.4, the high-statistics data constrain slow linear drifts in the background. The results from that section only constrain the size of a systematic effect to $\sim 0.5\%$, though this can be readily suppressed by decreasing the background rate of coincident events, improving the V detector efficiency, and by performing dedicated studies of potential linear drifts to higher precision – all of these contribute linearly to the size of the effect. Further still, in the presence of a linear drift, background runs before and after a foreground run can be

averaged to partially correct for this effect. The combination of these techniques should lead to a sub-0.1% effect, and the effect must ultimately be limited to 5×10^{-4} in order to not be the leading systematic.

6.11 Phase Space Evolution and Vanadium Activation

As discussed in chapter 2, if the occupied phase space of a magnetic bottle can change with time over the course of a measurement, this can cause a systematic effect; if the detection of surviving UCN or decay products varies over the phase space of the trap, the time-evolution of the phase space will translate into a time-dependent detection efficiency. Here, we consider, in basic terms, the implications of such an effect for a τ_n measurement using the vanadium activation technique.

The activation of the vanadium versus time can be expressed as

$$\dot{V} = A(t) - \tau_v^{-1}V \quad (6.33)$$

where V is the activity of the foil, $A(t)$ is the rate of capture of UCN on the foil, and $\tau_v = 324$ s is the mean lifetime of ^{52}V . if the foil is lowered for a time t_e , then the activity during the activity measurement is given by

$$V(t) = e^{-t/\tau_v} \int_0^{t_e} e^{t'/\tau_v} A(t') dt' \quad (6.34)$$

While the foil is lowered into the trap, the number of UCN N decreases either due to β -decay or due to absorption on the foil, so that

$$N = N_0 \exp \left(-\frac{t + t_s}{\tau_n} - \int_0^t D(t) dt \right) \quad (6.35)$$

where $D(t)$ is some function of time expressing the rate of absorption of UCN on the vanadium foil, and t_s is the storage time. This rate is given by $A(t) = D(t)N$, and the signal C from vanadium activation measured from the end of the emptying time t_e until some time t_m is given

by

$$C = \int_{t_e}^{t_e+t_m} \dot{V} dt. \quad (6.36)$$

Phase space time dependence causes a systematic effect if the draining function $D(t)$ changes between two counting measurements C_1 and C_2 . To gain insight, we suppose that for the first storage measurement, the absorption of UCN onto the foil is exponential with a constant characteristic time τ_D , and thus $D(t) = \tau_D^{-1}$. The draining time for the second measurement is then given by a characteristic time $\tau_D'^{-1} = \tau_D^{-1}(1 + \eta)$ due to a time-dependent change in the occupied phase space of the trapped UCN. We then compute the lifetime from the signals C_1 and C_2 including this difference in time constants. The relative pull on the lifetime for a typical draining duration and draining time constant is shown in fig. 6.17 as a function of η and as a function of t_e .




Figure 6.17: The relative correction to the lifetime (for nominal 885 s lifetime) as a function of the vanadium absorption time t_m (left) and as a function of the change of draining η (right).

From this we see that a change in the shape of the draining curve can cause two counter-acting effects. Suppose the characteristic time increased at later time. In this case, more counts would appear at the end of the tail, and thus more would be missed by the foil, causing later times to exhibit unexpectedly lower activity, and thus τ_n would appear shorter. At the same time, a longer draining time means that most of the activated vanadium nuclei have had less time to decay before the foil is raised for counting, thus causing later times to exhibit unexpectedly

higher activity. The latter effect occurs no matter how much of the tail-end of the distribution is missed, which is why the pull on the lifetime does not vanish for long t_m as long as $\eta \neq 0$.

For typical measurement times, the extraction of τ_n is affected at the level of 0.1% if the change in mean draining time is as big as 1-10% using this simple exponential model. This sets the scale of the effect, but there is no reason to expect that the draining function is truly exponential in nature: UCN of different energies presumably have different characteristic draining times, and the possibility of non-equilibrium behavior in the trap permits non-exponential behavior.




Figure 6.18: The rate of absorption on the vanadium versus time after three different storage times. Two-exponential fits are also shown with best-fit values for time constants t_1 and t_2 .

To provide a more realistic estimate of the effect, we simulate three measurements. In each measurement, the trap is filled for 60 s, stored for either 10, 200, or 800 s, and then the vanadium absorbs the trapped neutrons for 60 s. In each case, the spectrum loaded into the trap is a $v^2 dv$ distribution initially at the trap door, up to the trappable neutron velocity. The momentum distribution is isotropic in 2π facing up into the trap. We record the time of absorption of the UCN onto the vanadium foil, in effect measuring the draining curve $D(t)$. These data are

normalized to the number of neutrons that are filled into the trap for the given simulation.

The distribution in all cases is effectively described by a two-exponential fit:

$$A(t) = A_1 e^{-t/t_1} + A_2 e^{-t/t_2}. \quad (6.37)$$

The draining profiles for each simulation are shown in fig. 6.18. Also shown are double exponential fits with the best-fit values for t_1 and t_2 . The time t_1 is in approximate agreement with the data from chapter 5. The 10 s storage time exhibits a different time constant from the longer storage times, which suggests that after ~ 200 s, the draining profile stabilizes. We can compute the observed vanadium activity at the time that the foil is raised:

$$N = \int_0^{60} e^{t/\tau_v} D(t) dt. \quad (6.38)$$

The 200 and 800 s storage time signals are discrepant at a relative level of $(5 \pm 5) \times 10^{-4}$, which also sets the size of this effect on the extraction of τ_n . However, this would be further mitigated by using a larger foil, performing more detailed studies of the draining time of UCN into the foil, and more extensive simulations. Further, new detector technology will permit the direct measurement of the capture rate $A(t)$, as will be discussed in chapter 8.

6.12 UCN Source Fluctuations

The previous UCN bottle experiments discussed in chapter 2 used a UCN detector outside of the bottle as a means of normalizing each storage measurement. All of the high-precision experiments using material bottles were performed at the reactor-based Institut Laue-Langevin UCN source. Because it is a reactor-based source with a constantly rotating turbine to downscatter UCN, the incoming rate of UCN into the apparatus is assumed to be constant. Therefore, measuring the rate in a monitor detector is sufficient to normalize the initial trap density. This assumption appears to be implicit in the precision measurements using the material bottle technique with the exception of Arzumanov et al[66] who state explicitly that the turbine-based UCN source at the

ILL is sufficiently stable for their measurements.

The proton beam delivered to the LANSCE UCN source operates by delivering pulses at a rate of 201.25 MHz over a $625 \mu\text{s}$ wide RF gate. There are 5 pulse gates at a rate of 20 Hz that direct protons to the UCN spallation target, and this 5-pulse chain occurs every 5 seconds. During each pulse chain, UCN are produced and are transported through several meters of UCN guides until the apparatus is reached. There is thus a substantial delay between this pulse and the arrival of the UCN at the apparatus and monitor detectors, and the time-of-arrival distribution can have a non-trivial structure due to the geometry of the UCN source, guides, apparatus, and the shape of the UCN spectrum.

The rate in the monitor detectors and the time dependence of the density of UCN in the trap is thus determined by the number of UCN produced during a pulse chain convolved with the time of arrival spectrum to the apparatus (call it $G(t)$). That is, The number of UCN in the trap can be expressed as

$$N(t) \propto \sum_n^N I_n G(t - n\Delta t) \quad (6.39)$$

where there are N pulse chains during a beam-on period to fill the apparatus, and $\Delta t = 5 \text{ s}$ is the separation between pulse chains. The I_n represent the UCN production during the n th pulse chain, each being a function of the proton current delivered to the target and the temperature and pressure of the neutron moderator and UCN converter. The system response function $G(t)$ takes the approximate form $G(t) \propto t \exp(-t/t_g)$ where t_g depends on the average lifetime of UCN in the source and guide system leading to the experiment. This functional form agrees with results from simulations; further, by operating the proton beam in a long-pulsed mode (i.e. one pulse chain every $\sim 200 - 400 \text{ s}$) this quantity can be measured directly, thus providing a more direct means of estimating this effect. This will be done in future run campaigns. For the purposes of estimating the size of the effect in this section, the analytical form will be used, with $t_g = 20 \text{ s}$ being a reasonable estimate based on neutron tracking studies and preliminary measurements.

Ultimately, a systematic effect can occur if the I_n are not constant – that is, some I_i is different from the rest of the I_n while the trap is being filled. This fluctuation will propagate




Figure 6.19: The rate in a monitor detector based on the convolution of the system response $G(t)$ with sequential pulse chains; the black curve shows the rate assuming a sequence of constant pulse chains, and the red shows a constant set of pulse chains, except for a single pulse chain with 1% lower UCN output during filling (top). The same rates, but zoomed in on the region of nearly-constant rate during filling, over which the rate is averaged to compute the normalizing factor (bottom).

through the guide system to become a fluctuation in the initial number of UCN in the trap by the time the trap door is closed. The effect is then manifest in the data analysis when the saturated rate in the monitor detector is determined by integrating the detector rate over the region of ostensibly constant rate (e.g. as is done in fig. 5.15). This fluctuation, while difficult to observe above the statistical fluctuations in rate, can systematically change the mean rate that is used to normalize each run. This is shown in fig. 6.19, which shows the essentially imperceptible difference between a completely constant fill versus a fill with a 1% fluctuation in UCN output during a pulse chain.

However, with knowledge of $G(t)$ and an independent measure of these fluctuations during each pulse chain, the effect can be estimated and, to an extent, corrected. The proton charge delivered to the target during each RF gate is measured by commercial (Bergoz) integrating current transformers with linearity of better than 1%. In addition, ^3He -filled drift tubes for

thermal neutron detection are embedded into the biological shielding near the UCN source, and these can be used to measure the neutron flux generated by each pulse chain. The combination of these two measurements can identify fluctuations at the 1% level (i.e. fluctuations of the size shown in fig. 6.19). Integrating over the rate in the time window of the lower panel of the figure, this leads to a 3×10^{-4} relative discrepancy in the average rate, and thus in the normalization factor for a UCN storage measurement.

This is further complicated by the fact that the time response function for the monitor detectors could be somewhat different from that for the trap itself. The response for the trap itself can be measured during a run campaign by temporarily placing a UCN detector in or near the trap, and compared to that for the separate monitor detectors. Ultimately, a detector near or above the trap could be used as the dedicated monitor detector.

It is also worth noting that the average rate during filling is the simplest quantity that can be used as a normalization factor. A more sophisticated normalizing quantity, which for example weights earlier and later times during filling differently, or includes non-statistical fluctuations in the uncertainty of the normalization factor, can further mitigate the effect.

6.13 Summary and Conclusions

We have discussed the potential systematic effects for a future measurement of the neutron lifetime using the current apparatus to a precision of 0.1% or better. In general, these effects appear to be manageable at the 0.1% level, but will require much more sophistication to quantify for a next-generation 0.01% measurement. A summary of the systematic effects, with estimated corrections and uncertainties, are shown in table 6.2.

As will be discussed in chapter 8, forthcoming advances in the UCN detector technology could further address or avoid some of the effects in this chapter. In particular, an *in situ* detector with high efficiency and low background will mitigate the effects of time dependent backgrounds and gain drifts, and will offer a means of investigating the effect of phase space evolution on UCN detection. Such a detector could be employed within the trap in order to monitor the initial trap

effect	correction	comments
residual gas	$(+1 \pm 1) \times 10^{-4}$	RGA monitoring of gas content
depolarization	$< 10^{-4}$	<i>in situ</i> verification of depolarization
material losses	$(+4 \pm 4) \times 10^{-4}$	measure η for the Cu
cleaning	$(+6 \pm 6) \times 10^{-4}$	vary cleaning heights and times
cleaner reliability	$\pm 5 \times 10^{-4}$	verify position reproducibility
microphonic heating	$\pm 1 \times 10^{-4}$	<i>in situ</i> accelerometry
dead time/pileup	$< 10^{-4}$	pileup ID and artificial dead time
gain drifts	$\pm 2 \times 10^{-4}$	regular gain monitoring
time dependent backgrounds	$\pm 5 \times 10^{-4}$	need sufficient background data
phase space evolution	$\pm 5 \times 10^{-4}$	measure V draining curve
source fluctuations	$\pm 3 \times 10^{-4}$	measure monitor/trap response
total	$(1.1 \pm 1.2) \times 10^{-3}$	

Table 6.2: Estimates of potential systematic effects in the current experiment, with corrections where applicable.

density, thus providing insight into the severity of UCN source fluctuations.

The sources of material loss and residual gas must be investigated. Recent experiments at LANSCE have measured the upscattering cross section for various gases; this, in combination with regular monitoring of the residual gas content of the trap, will mitigate this source of loss. The use of Cu to cover magnetic defects must be investigated by assuring a low loss-per-bounce and depolarization per bounce for the surface, and field mapping near these defects can be used to assure the absence of field zeros and low-field “holes” in the magnetic trap. Ultimately, these defects can be repaired if the results of these studies introduce large corrections.

The detector technology to be discussed in chapter 8 can also be used to study UCN cleaning. A larger-area cleaner appears to be more than adequate based on neutron tracking simulations performed here and by Walstrom et al. A portion of a cleaner could be instrumented as a detector, and the cleaning rate versus time measured to compare to simulation and assure the efficacy of the technique. This would greatly improve upon the technique of detecting the upscattered neutrons, which to-date exhibits low signal-to-background.

7 Solid Nitrogen as a UCN Converter

7.1 Introduction

Several planned UCN sources use solid deuterium (D_2) to downscatter cold or thermal neutrons to UCN energies[24, 71, 19, 41]. Since it was first proposed and investigated[85, 112, 25, 26], considerable recent effort has gone into understanding UCN production in D_2 [20, 104, 44, 43]. Deuterium is chosen for its low neutron absorption cross-section, low incoherent scattering cross-section (to minimize UCN elastic scattering within the source), and the presence of numerous phonon modes which can inelastically scatter neutrons down to UCN energies. However, aside from these phonon modes, the free rotation of D_2 molecules can cause incoherent upscattering of UCN in the deuterium. In order to suppress this upscattering, the D_2 sample is converted to the $J = 0$ (ortho) state using paramagnetic catalysts[20, 37].

Solid $\alpha\text{-}^{15}\text{N}_2$ is a potential alternative to deuterium: its absorption cross-section is only 5% of that of D_2 , and it has a negligible incoherent scattering cross-section. Additionally, free rotation of the N_2 molecules in the lattice is inhibited due to the anisotropy of the N_2 inter-molecular potential. This leads to dispersive modes for the rotational degrees of freedom (librons) which provide additional channels for neutron downscattering and eliminate the rotational incoherent upscattering prevalent in D_2 .

Due to the primarily classical nature of its inter-molecular potential, the anharmonicity of the angular degrees of freedom, and phonon/libron coupling, solid nitrogen has been a test-bed for *ab initio* lattice and molecular dynamics calculations[118]. Several models exist which utilize inter-molecular or atom-atom potentials of various forms[106, 45, 116, 8, 88, 131], and anharmonic effects have also been investigated experimentally and theoretically[6, 64, 5, 130].

Past interest notwithstanding, there is little data on the dynamic structure factor of α -N₂ from neutron scattering[99].

In this chapter, solid ¹⁵N₂ in the α phase ($T < 35$ K) is investigated as a new UCN converter. The dynamic structure factor is measured using a time-of-flight spectrometer, and the UCN production and upscattering cross sections are computed using the results from sec. A.5.

7.2 Experiment

The dynamic structure factor is measured with the IN4 TOF spectrometer at the Institut Laue Langevin (ILL) using incident neutrons of wavelength $\lambda = 2.2\text{\AA}$ [111]. The sample environment, cryostat, and sample control system are the same as that used for previous studies of oxygen and deuterium as UCN converters[20, 40]. This provides a means of direct comparison between the different converters. Gaseous ¹⁵N₂ (99.4% isotopic purity) is introduced into the 2 mm annular cell through a needle valve, and cooled using the ILL's orange helium-exchange cryostat (see [20, 40] for details). The detector array is calibrated using an isotropically scattering vanadium sample. The dynamic structure factor and related quantities are computed using the ILL's standard data treatment software LAMP[113].

Total scattering from the sample cell is measured while condensing nitrogen to assure that the cell is completely filled, and the sample temperature is then lowered to the solid α -phase. In order to ensure that the nitrogen sample did not consist of macro-crystals, the elastic scattering intensity is analyzed for various annealing times at the $\alpha - \beta$ phase transition. We find only slight variations in the elastic peaks for different annealing times, which suggests that the sample is polycrystalline as desired.

Data are then acquired with nitrogen at temperatures of 5.9, 11, 15, 20, and 25 K. A background measurement of the empty cell is performed at 11 K. The inelastic scattering background due to the aluminum phonons is temperature-dependent, especially on the anti-Stokes ($E < 0$) side, changing with temperature according to the Bose-Einstein occupation factor $n_{BE}(T, E) = [\exp(E/kT) - 1]^{-1}$. We thus rescale the background intensity by the relative

change in occupation number (as a function of E) prior to subtracting it from the nitrogen data. This is particularly important for computing UCN upscattering mean free path for $T < 11$ K (see below), where it is a $\sim 20\%$ effect if uncorrected.

We also perform measurements using $^{14}\text{N}_2$, and find that the dynamic structure is the same as that for $^{15}\text{N}_2$, except for a difference in overall scattering intensity commensurate with the isotopic difference in total scattering cross-section ($\sigma_{\text{scat}14}/\sigma_{\text{scat}15} \approx 2.2$). For this reason, data with $^{14}\text{N}_2$ are used for the determination of the temperature dependence of the dynamic structure factor because of better statistical sensitivity. The results are scaled by the above factor to determine the absolute $^{15}\text{N}_2$ cross-sections.

7.3 Results




Figure 7.1: The differential scattering cross-section versus E . $E > 0$ corresponds to energy loss, and $E < 0$ to energy gain. The vertical scale is set by integrating to the total scattering cross-section (see text).

The differential scattering cross-section $d\sigma/dE$ is shown in fig. 7.1. The vertical scale factor κ is set by demanding that

$$\kappa \int \left(\frac{d\sigma}{dE} \right)_{data} dE = \sigma_{\text{scat}} \quad (7.1)$$

where σ_{scat} is the total scattering cross-section per molecule. We find $\kappa = 36.6 \pm 7.3$ b, with the 20% uncertainty due primarily to the unknown contribution of scattering intensity outside the kinematic range of the instrument. The factor κ has units of barns so as to convert the $d\sigma/dE$ data (which has an arbitrary normalization) into a cross-section. This normalization factor is used below to calculate the absolute UCN production and upscattering cross-sections.

The dynamic structure factor $S(q, E)$ (averaged over orientations of \vec{q} due to the polycrystallinity of the sample) is shown in fig. 7.2. The intensity at $q = 1 \text{ \AA}^{-1}$ has peaks at $E = 4.5$, 5.7, and 8 meV, in reasonable agreement with the T_g , A_u , and T_u modes predicted by Kjems and Dolling[99]. However, it is difficult to comprehensively compare the data to various dynamic models due to the low energy resolution of the TOF method, as well as the polycrystalline nature of the sample.




Figure 7.2: The dynamic structure factor of polycrystalline $\alpha\text{-}^{15}\text{N}_2$. The color scale is set by integrating the measured differential cross-section and equating it to the total scattering cross-section. The black line corresponds to the UCN production curve given by eqn 7.2.

Fig. 7.3 shows the Generalized Density-Of-States (GDOS), which exhibits prominent peaks from one-phonon states, and broad low-lying intensity due to multi-phonon processes. Elastic

contamination is removed by fitting the GDOS to E^2 for $E \lesssim 2$ meV. The peaks in the 4 to 8 meV range broaden with increasing temperature due to phonon and libron conversion and dephasing, which has been studied in detail using infrared and Raman spectroscopy[64, 5].




Figure 7.3: The GDOS for solid nitrogen at two temperatures in the α -phase. Peak broadening at higher temperature is observed.

7.4 Discussion

The black line in fig. 7.2 corresponds to the necessary kinematic condition for UCN downscattering. As UCN energies are much less than cold and thermal neutron energies, we take the UCN energy $E_{UCN} \approx 0$ and arrive at the following relation for the wave-vector transfer $q = |\vec{k}_i - \vec{k}_f|$ and energy transfer $E = E_i - E_f$:

$$E = \pm \frac{\hbar^2 q^2}{2m_n} \quad (7.2)$$

with the plus sign for UCN production, and minus sign for UCN upscattering.

The intensity of $S(q, E)$ parameterized along the curve given in eqn. 7.2 is then proportional to the cross-section to nearly “stop” a cold neutron, thereby converting it to a UCN. The UCN production cross-section versus incident energy integrated over final UCN energy E_f from 0 to

some maximum UCN energy E_{max} is given by

$$\sigma_{UCN}(E_i) = \frac{2}{3} \kappa E_{max}^{3/2} \cdot \frac{1}{\sqrt{E_i}} S\left(\sqrt{2m_n E_i}/\hbar, E_i\right). \quad (7.3)$$

Upon escaping the nitrogen volume, UCN are boosted by the solid $^{15}\text{N}_2$ optical potential of ~ 69 neV. We choose $E_f = 181$ neV, which corresponds to a trappable energy of 250 neV, commensurate with UCN production cross-section calculations in [40].




Figure 7.4: The UCN production cross-section in $^{15}\text{N}_2$ for $E_{UCN} \leq 181$ neV versus incident neutron energy. The inelastic cross-section for $E < 2$ meV is difficult to determine due to elastic contamination, though is likely small compared to that of the energy range shown here.

The energy-dependent production cross-section is shown in fig. 7.4. Production peaks near 6 meV, most probably due to a combination of vibrational and librational downscattering channels. The production cross-section can vary with temperature due to the temperature dependence of normal mode frequencies and their widths: we find the variation in the cross-section to be no more than 18% in the range of 5 to 25 K (increasing slightly with increasing temperature).

From the above, the rate of UCN production in a source using $^{15}\text{N}_2$ can be estimated, making reasonable assumptions about the incident neutron flux. For a UCN source illuminated with the




Figure 7.5: The UCN yield of $^{15}\text{N}_2$ for a neutron flux of $10^{14} \text{ cm}^{-2} \text{ s}^{-1}$ from a cold moderator at temperature T . The yield in UCN is optimized for incident neutrons at a temperature $T = 40 \text{ K}$.

neutron flux from a cold moderator at temperature T , we define the UCN yield to be

$$Y(T) = n \int \phi(E, T) \sigma_{UCN}(E) dE \quad (7.4)$$

where $n = 2 \times 10^{-22} \text{ cm}^{-3}$ is the number density of $^{15}\text{N}_2$, and the cold neutron flux is given by

$$\phi(E, T) = \phi_0 \sqrt{\frac{4E}{\pi k^3 T^3}} \exp(-E/kT). \quad (7.5)$$

From this, we find an optimal moderator temperature of 40 K (see fig. 7.5). For a UCN source with total incident cold neutron flux $\phi_0 = 10^{14} \text{ cm}^{-2} \text{ s}^{-1}$ at $T = 40 \text{ K}$, the yield is approximately 8.9×10^4 UCN per second per cm^3 of source volume. In order to extract a UCN from a source for use in an experiment, the UCN mean free path λ must be comparable to or larger than the thickness of the source. The mean free path contains contributions from different processes:

$$\lambda^{-1} = \lambda_{up}^{-1} + \lambda_{abs}^{-1} + \lambda_{inc.el.}^{-1}. \quad (7.6)$$

where λ_{up} corresponds to UCN upscattering to non-UCN energies, λ_{abs} to nuclear absorption, and $\lambda_{inc.el.}$ to UCN elastically scattering within the nitrogen volume. Even for relatively poor isotopic purity, the incoherent scattering mean free path is much larger than the size of a potential UCN source, and can be neglected. The absorption mean free path is highly dependent on the isotopic purity, and purity $\sim 99.9\%$ or higher is necessary for upscattering to be the dominant contribution to the mean free path.




Figure 7.6: The mean free path λ_{up} for UCN to upscatter to non-UCN energies within the solid nitrogen volume. The error bars are statistical.

The (temperature dependent) upscattering cross-section of UCN in the nitrogen can be computed from the data by choosing the minus sign in eqn. 7.2 and integrating over the final state energy:

$$\sigma_{up} = \kappa E_0^{-1/2} \int_0^\infty \sqrt{E} S\left(\sqrt{2m_n E}/\hbar, -E\right) dE. \quad (7.7)$$

We compute $\lambda_{up} = 1/\sigma_{up}n$ for UCN with $E_0 = 181$ neV over a range of temperatures (see fig. 7.6). We observe an increased λ_{up} with decreasing temperature due to the smaller population of lattice modes that can induce upscattering. At $T = 5.9$ K, the data suggest an upscattering mean free path of approximately 46 cm, which is greater than the thickness of current D₂ sources

such as the Los Alamos source (~ 5 cm) or the PSI source (~ 15 cm)[24, 19]. The thickness of a UCN source using $^{15}\text{N}_2$ could therefore be made larger than a D_2 source, the latter of which is limited by its 8 cm incoherent scattering mean free path.

7.5 Conclusions

The dynamic structure factor of solid $\alpha\text{-}^{15}\text{N}_2$ has been measured over a range of temperatures. No appreciable difference in dynamics is observed between $^{14}\text{N}_2$ and $^{15}\text{N}_2$, and there is reasonable qualitative agreement between our data and previous studies of the dynamics of solid α -nitrogen. However, a higher resolution neutron scattering technique would be necessary to thoroughly investigate the rich structure of the vibrational and librational modes.

The absolute UCN production cross-section is determined by normalizing the scattered intensity to the total molecular cross-section of the sample. The cross-section peaks strongly near 6 meV, and the optimal incident cold neutron temperature is 40 K. The measured cross-section is found to be somewhat lower than that of D_2 and O_2 (see [40]). However, we observe an upscattering mean free path substantially longer than that of D_2 . Thus, for sufficiently high isotopic purity, an $\alpha\text{-}^{15}\text{N}_2$ source could be large while maintaining efficient UCN extraction to compensate for the lower production cross-section.

In addition, a nitrogen-based source may benefit from operating temperatures below those used here, if the upscattering cross-section can be further reduced at lower temperatures (~ 1 K). While deuterium is ultimately limited by its temperature independent absorption and incoherent scattering mean free paths, further improved isotopic purity of ^{15}N may make an even lower temperature nitrogen volume a promising candidate for a new UCN source. This motivates future work in performing direct measurements of slow neutron mean free paths and UCN production rates in $^{15}\text{N}_2$.

8 Conclusions

8.1 Summary and Overview

An independent, precise determination of τ_n using magnetic confinement is timely. Knowledge of τ_n , along with other n -decay observables, can discriminate potential beyond-SM processes which distort charged-current weak interactions in a way that is complimentary to, and potentially competitive with, high energy experiments. Moreover, the accurate resolution of the neutron lifetime with 10^{-3} relative precision would have a significant impact on cosmological predictions of the effective number of thermalized neutrinos in the early universe, and thus is a test of new physics manifest therein.

While ultimately measuring τ_n with 10^{-4} precision is desirable for competitive tests of new physics, current limits suffer from the ambiguity in the global data. The nature and size of systematic corrections in UCN material bottle measurements have been called into question in the literature, and magnetic confinement is a promising new technique due to small (and potentially negligible) losses. It is worth noting, however, that the current particle data group values for V_{ud} and λ predict a lifetime of 883.1 ± 0.8 s, which is in disagreement with both the most precise UCN bottle measurement and most precise beam measurement (the latter at the level of $2\text{-}\sigma$)[\[22, 31\]](#). While this doesn't point to the source of the experimental ambiguity, it does suggest that a new type of measurement is needed.

The UCN bottle experiments surveyed in [chapter 2](#) employed several experimental handles to quantify losses and other sources of systematic error, such as varying wall temperature, filtering the UCN spectrum, or measuring upscattered neutrons, though the data analysis is typically complicated by the precarious nature of UCN. For example, the UCN wall loss rate is UCN and

energy and momentum distribution dependent, and direct measurements of these distributions is often difficult or impossible. As such, model assumptions, equilibrium kinetic theory, simulation, or perturbative analyses are employed in order to extrapolate measurements to the true lifetime. These analyses for a given experiment often achieve remarkable internal consistency, though this does not preclude the presence of systemic effects which do not vanish when the endemic effects are extrapolated away.

We have presented the design of a magneto-gravitational UCN trap for the measurement of the neutron lifetime with large volume and geometry that permits the study of UCN cleaning, study of the occupied phase space of the trap, and new UCN detection techniques. We place an upper bound on the non- β -decay loss rate in the trap and demonstrate the ability to detect cleaned UCN. The potential for much greater statistical sensitivity is possible with future changes in the guide geometry and trap loading mechanism, and also through the development of *in situ* UCN detection schemes such as vanadium activation.

The vanadium activation technique has demonstrated high β detection efficiency. The γ efficiency with the NaI detectors is somewhat lower, but this is dominated by solid-angle efficiency, and forthcoming redesigns can increase the efficiency by a factor of two or three. The data presented here show that UCN are rapidly absorbed onto the V foil which is a substantial improvement over the fill-and-empty technique used in previous experiments. In addition, with modest improvements to the shielding, the backgrounds can be substantially stabilized. In addition, the scintillators and PMTs demonstrate stability sufficient for a 1 s measurement, and can be improved with source calibrations and active gain monitoring.

In spite of the performance of the apparatus detailed in this dissertation, there remains a pervasive issue for this and other UCN experiments: UCN densities are low. As mentioned in chapter 7, there are several new s-D₂ and superfluid helium UCN sources being developed which promise vast improvements, but none have thus far demonstrated substantially better UCN densities than LANSCE or the ILL in an apparatus external to the source itself. This motivates the investigation of new UCN converters and other improvements in UCN technology.

The improvement of UCN production and transport technology could transform the nature of UCN experiments. The experiment described in chapter 7 offers a simple means of measuring a UCN production cross section without the need for a dedicated source, and the technique could be used for other candidate materials as well, such as ^{208}Pb . Nitrogen offers a somewhat lower UCN production cross section, but also shows the potential for longer UCN mean-free-paths. This necessitates more work to measure mean free paths in these materials in order to demonstrate that a large source is beneficial.

8.2 Outlook

8.2.1 Magnetic Confinement

The past research on magnetic confinement discussed in chapter 2 and measurement of the storage time with the magneto-gravitational trap discussed in chapter 4 clearly demonstrate that sources of loss using magnetic confinement are small or negligible. This technology is therefore ideal for a measurement of the neutron lifetime, as it obviates the need for dimensional or spectral extrapolation to determine τ_n . The presence of the copper tape to cover defects in the apparatus described here must be investigated further by checking what fractional area of the array exhibits lower magnetic field strength due to the defects. In addition, separate measurements using the adhesive-covered copper used in the trap must be performed to estimate the loss-per-bounce, and in particular verify that it is no larger than $\sim 10^{-4}$.

The depolarization of magnetically trapped UCN remains a potential issue for next-generation measurements. Basic estimates, as well as more detailed calculations, of depolarization rates suggest that the effect should be negligible at the level of 10^{-4} or lower. However, these calculations must be verified with experiment. A thin-film detector could provide a means of measuring or placing limits on the rate of depolarization in a Halbach array. Boron-coated fibers or ZnS screens (see subsection 8.2.3) could be placed immediately on the array surface. Even if the whole array surface is not covered, a measurement using unpolarized (or high-field polarized) UCN could provide a reference rate in the detector. Thereafter, any rate measurement with polarized UCN

during storage could be interpreted as being due to depolarization.

From the neutron tracking results, it is clear that a large area cleaner will be sufficient to remove all but 10^{-3} to 10^{-4} UCN in the trap, and this can be confirmed by repeating measurements of τ_n with and without cleaning. In addition, the neutron tracking results show that the draining of UCN into the vanadium is stable for long and short storage times, suggesting that phase space time dependence is mitigated with this trap.

Still, there may exist other trap geometries which facilitate rapid cleaning and little phase space time dependence. This is the subject of ongoing theoretical and numerical research, and other asymmetric trap geometries have recently been proposed. These exhibit large fractions of chaotic orbits, and could utilize superconducting magnets to vary the trap potential. While a superconducting trap greatly increases the experimental complexity and turn-around time for measurements and R&D, it could provide additional means for characterizing the trapped population of UCN and reducing systematic effects related to magnetic confinement.

8.2.2 Vanadium Activation

The vanadium activation technique and detector array described in chapter 5 demonstrates rapid absorption of surviving UCN within the trap. While the phoswich geometry did not ultimately provide significant background rejection, the coincidence technique greatly reduced backgrounds. The NaI efficiency will be increased in future run campaigns by redesigning the portion of the vacuum vessel near the detectors, in order to move the NaI closer to the V foil. In addition, more NaI modules will be added. This will likely increase the efficiency by a factor of 2 or 3, for a total efficiency near 50%.

From the analysis in chapter 5 and estimates from chapter 6, the detector stability and backgrounds are likely to be sufficient for a 0.1% measurement of τ_n , as long as time dependent backgrounds can be controlled. This will require long-term studies of the background in future run campaigns, and adequate shielding for thermal neutron induced activation.

Implementing a larger foil will decrease the absorption time of UCN onto the foil, and this

could further reduce any potential systematic effects related to non-equilibrium of the trapped population. These effects can be checked by measuring this absorption time after short and long storage times. In addition, varying the height of the foil will provide a detailed measurement of the UCN distribution in the trap, and can be used for monte carlo and model validation.

8.2.3 New in situ Detection Techniques

While the vanadium activation detectors and cleaner upscatter detectors are promising, the direct counting of surviving and cleaned UCN within the trap would be desirable. This could further improve efficiency and signal-to-background, and provide a direct measurement of the cleaning time profile and survivor counting time profile, which would be desirable for monte carlo validation as well as evaluating the systematic effects discussed in chapter 6.

A promising technique is that of ^{10}B -coated ZnS films coupled to photo-detectors. Recently, such a detector was tested at the LANSCE UCN source which consisted of a $\sim 4\ \mu\text{m}$ thick ZnS screen evaporatively coated with 20 – 300 nm of isotopically pure ^{10}B . The screen is then coupled to a PMT to read out the scintillation light produced by the daughter ions of the boron-capture reaction. The ^{10}B layer is exposed directly to the UCN volume. There is no detector window, greatly improving the detector efficiency. Preliminary tests show excellent efficiency and signal-to-background with 7.6 cm diameter foils coupled to the UCN beamline[82]. This technology could replace the vanadium foil and cleaner, using wavelength-shifting fibers to transport light to PMTs or Silicon PMs outside the vacuum system.

In general, the development of new detection techniques, along with magnetic confinement of UCN, will make an independent determination of τ_n possible. This is necessary to address the ambiguity in the current neutron decay data. This is a necessary step towards improved BBN estimates, and along with the broader suite of neutron decay experiments, will provide a test for new physics in the charged-current sector.

Appendices

A Slow Neutrons

A.1 Introduction

Neutrons interact with magnetic fields via their magnetic moment $\mu_n \approx -60.3 \text{ neV/T}$, with nuclei via the strong nuclear force, with gravitational fields due to their mass $m_n \approx 939.565 \text{ MeV}/c^2$, and as discussed in chapter 1, they decay via the weak force.

Neutrons are often classified in terms of their energy range, and for neutrons with $E \lesssim 0.1 \text{ eV}$, in terms of their temperature equivalent $T = k_B^{-1}E$, where $k_B \approx 0.0862 \text{ meV/K}$ is the Boltzmann constant. Here, we define four types of “slow” neutrons: thermal, cold, very cold, and ultra cold. Table A.1 shows the approximate energy range of these neutrons, along with their range of de Broglie wavelengths λ , velocities v , wavenumbers k , and temperatures T , which are related by

$$E = k_B T = \frac{1}{2} m_n v^2 = \frac{\hbar^2 k^2}{2m_n} = \frac{h^2}{2m_n \lambda^2}. \quad (\text{A.1})$$

name	E	$v \text{ [m/s]}$	$\lambda \text{ [\AA]}$	$k \text{ [\AA}^{-1}\text{]}$	$T \text{ [K]}$
thermal	25 to 100 meV	2200 to 4000	1.2 to 2.5	3.5 to 7	290 to 1200
cold	1 to 25 meV	440 to 2200	2.5 to 13	0.7 to 3.5	12 to 290
very cold	300 neV to 1 meV	8 to 440	2.5 to 740	0.012 to 0.7	0.003 to 12
ultra cold	0 to 300 neV	0 to 8	>740	>0.012	0 to 0.003

Table A.1: Four classes of slow neutrons, with approximate energy, velocity, wavelength, wavenumber, and temperature ranges.

The physics of slow neutrons, and the theory of neutron scattering are rich topics, with many applications to condensed matter and nuclear physics. There are several texts on the subjects of slow neutrons[129, 119], neutron scattering[105, 125, 84], the dynamics of crystals[17], and ultracold neutrons[87, 95]. In this chapter, we limit the discussion of these subjects to the physics

of Ultracold Neutrons (UCN) relevant for chapters 2 and 3, and the theory of neutron scattering from crystalline solids (for chapter 7).

A.2 Slow Neutrons and Nuclei

Neutrons interact with nuclei via the strong force. Due to the strong but short ranged nature of the interaction, we can approximate the nuclear potential V_{nucl} seen by the neutron as a spherical well with a depth on the order of tens of MeV, and a radius on the order of a femtometer. We can write the wave function of the incoming and scattered neutron as

$$\psi = e^{ik \cdot x} + f(\theta) \frac{e^{ik \cdot x}}{r} \quad (\text{A.2})$$

where θ is the angle formed between the initial and scattered neutron wavevector. Because the slow neutrons described in table A.1 have wavelengths much longer than the size of the nuclear well, the scattering is predominantly in the s-wave, so that $f(\theta)$ is constant. We can simplify the problem by replacing the square well potential with an effective potential V_F which can be used in the Born approximation, and its strength tuned to agree with experiment:

$$V_F(r) = \frac{2\pi\hbar^2 b^2}{m_r} \delta(r - r_0). \quad (\text{A.3})$$

Here, m_r is the reduced mass of the neutron and nucleus, $\delta(x)$ is the (three dimensional) Dirac delta function, and r_0 is the location of the nucleus in the center of mass coordinates. This is referred to the *Fermi potential*. The scattering length can in general be complex, with the real part corresponding to the strength of the scattering. The imaginary part causes a non-conservation (in particular a decrease) of the probability current of the neutron wave function, and represents the nuclear capture of the neutron. Using Fermi potential in the Born approximation gives $f = -b$ and total cross section $\sigma = 4\pi|b|^2 = \sigma_s + \sigma_a$ with σ_s and σ_a the scattering and absorption cross sections.

It can be shown that, for most nuclei, thermal (and slower) neutrons are lower in energy than

the energy of excited nuclear states. The cross section for absorption is then simply proportional to the time that the neutron wave packet spends in the vicinity of the nucleus. Thus, we have that $\sigma_a \propto v^{-1}$. Slow neutron absorption cross sections are thus typically given at a reference velocity of $v^{(th)} = 2200$ m/s, so that the cross section at any velocity is expressed as

$$\sigma_a = \frac{\sigma_a^{(th)} v^{(th)}}{v}. \quad (\text{A.4})$$

A neutron beam incident upon a collection of nuclei will diminish due to absorption. The differential decrease in neutron intensity I in a material of number density n is given by the scattering probability (i.e. the effective area $n\sigma_a dV = n\sigma_a A dx$ divided by the beam area A) times the intensity:

$$dI = I\sigma_a n dx \quad (\text{A.5})$$

and thus, for an initial intensity I_0 , we have that

$$I(x) = I_0 e^{-x/\ell_a} \quad (\text{A.6})$$

where $\ell_a = 1/\sigma_a n$ is the *mean free path* for a neutron to travel through the medium before it is absorbed. We can likewise define a scattering mean free path $\ell_s = 1/\sigma_s n$ which represents the mean distance that a neutron travels before it is scattered, when in the presence of a collection of nuclei.

A.3 Neutron Scattering from Condensed Matter

In a condensed matter system, the nuclei are spatially bound, so that we let $m_r \rightarrow m_n$ for all scatterers. We can write the potential for a neutron incident on the system is given by a sum of the potentials at time t due to nuclei at positions $R_j(t)$:

$$V(r, t) = \frac{2\pi\hbar^2}{m_n} \sum_j b_j \delta(r - R_j(t)) \quad (\text{A.7})$$

where we note that the positions R_i are operators. For a neutron with incident wavevector k_i , final wavevector k_f , initial and final energy E_i and E_f , we can express the general cross-section for a neutron to scatter into a solid angle $d\Omega = 2\pi d\theta$, with energy transfer $E = E_i - E_f$ with the use of Fermi's golden rule

$$\frac{d^2\sigma}{d\Omega dE} = \frac{m^2}{4\pi^2\hbar^4} \frac{k_f}{k_i} \sum_{j'} \sum_j P_j |\langle k_f, j' | V | k_i, j \rangle|^2 \times \delta(E + E_{j'} - E_j). \quad (\text{A.8})$$

Here, j and j' are the initial and final states of the scattering system, and P_j is the probability of the system being in the j th state. This general form for the *double differential cross-section* can be expressed using plane wave solutions for the incoming and outgoing neutron, and considering the matrix element in the fourier domain

$$\frac{d^2\sigma}{d\Omega dE} = \frac{1}{2\pi\hbar} \frac{|k_f|}{|k_i|} \sum_{jj'} b_j b_{j'} \int_{-\infty}^{\infty} \langle e^{-iq \cdot R_{j'}(0)} e^{iq \cdot R_j(t)} \rangle e^{-iEt/\hbar} dt \quad (\text{A.9})$$

where the brackets represent the expectation value of the operator, and $\hbar q = \hbar(k_i - k_f)$ is the momentum transfer which satisfies $|q|^2 = |k_i|^2 + |k_f|^2 + 2|k_i||k_f| \cos \theta$.

The cross section is often re-expressed in terms of correlation functions. First, we define the *pair correlation function* or *van Hove function* for a system of N atoms:

$$G(r, t) = \frac{1}{N} \sum_{jj'} \int \langle \delta(r' - R_{j'}(0)) \delta(r' + r - R_j(t)) \rangle d^3r'. \quad (\text{A.10})$$

The function $G(r, t)$ represents the correlation between a nucleus j at position $r = 0$ and time $t = 0$, and nucleus j' at position r time t . The double sum in A.10 can be divided into the sum over like atoms ($j = j'$) and distinct atoms ($j \neq j'$), for which we have the *self* and *distinct* pair correlation functions G_s and G_d respectively. It can be shown that the double differential cross section is given by

$$\frac{d^2\sigma}{d\Omega dE} = \frac{1}{\hbar} \frac{|k_f|}{|k_i|} \int \int [\langle b \rangle^2 G(r, t) + (\langle b^2 \rangle - \langle b \rangle^2) G_s(r, t)] e^{-iEt/\hbar} e^{iq \cdot r} dt d^3r \quad (\text{A.11})$$

with scattering length averages $\langle b \rangle = \sum_j p_j b_j$ and $\langle b^2 \rangle = \sum_j p_j b_j^2$. The scattering lengths may vary due to the atomic composition of the system, as well as the spin dependence of the scattering length of a given nucleus. The first term in the integrand is referred to *coherent scattering*, and includes the interference between the scattered neutron wavefunctions of different nuclei. This leads to, for example, the phenomena of neutron diffraction and coherent phonon scattering. The second term is referred to as *incoherent scattering*, and represents the (approximately isotropic) scattering of a neutron from a single scatterer.

The cross section can be further simplified by introducing the *dynamic structure factor* $S(q, E)$:

$$S(q, E) = \frac{1}{2\pi\hbar} \int \int G(r, t) e^{iq \cdot r} e^{-iEt/\hbar} d^3r dt. \quad (\text{A.12})$$

We thus arrive at the final expression of the cross section, which can be divided into coherent and incoherent parts

$$\begin{aligned} \frac{d^2\sigma}{d\Omega dE} &= \frac{d^2\sigma_{coh.}}{d\Omega dE} + \frac{d^2\sigma_{inc.}}{d\Omega dE} \\ &= \frac{|k_f|}{|k_i|} [\langle b \rangle^2 S_{coh.}(q, E) + (\langle b^2 \rangle - \langle b \rangle^2) S_{inc.}(q, E)] \\ &= \frac{|k_f|}{|k_i|} \frac{\sigma_{tot}}{4\pi} S(q, E) \end{aligned} \quad (\text{A.13})$$

where $\sigma_{tot} = \sigma_{coh.} + \sigma_{inc.} = 4\pi \langle b^2 \rangle$ is the *total cross section* (per atom) for a neutron to scatter from the system.

For a condensed matter system in thermal equilibrium, the populations P_j of the scattering states are proportional to $\exp(-E_j/k_B T)$. This leads to the *principle of detailed balance*, which relates the processes of neutron energy *loss* and energy *gain* when scattering from a collection of nuclei in thermal equilibrium. In terms of the dynamic structure factor, it leads to the relation

$$S(-q, -E) = e^{-E/k_B T} S(q, E). \quad (\text{A.14})$$

Scattering processes in which the neutron gains energy ($E < 0$) are often referred to as *Stokes*

processes, while energy loss ($E > 0$) are *anti-Stokes*.

A.4 Crystalline Solids

The dynamic structure factor encodes the frequency content and structure (in reciprocal space) of the scattering system. It is independent of the coupling of the neutron to the system and to the neutron kinematics. Therefore, a scattering experiment (with proper knowledge of the scattering kinematics) can be used to experimentally extract $S(q, E)$ and related quantities. This makes neutrons a powerful probe for the structure and dynamics of condensed matter.

The behavior of $S(q, E)$ depends greatly on the scattering sample. There are general characteristics of $S(q, E)$ for samples such as liquids, gasses, amorphous materials, and crystals can be found in the references given in the previous section. Here, we discuss the scattering of neutrons from crystals.

For a nucleus in a crystal, its position $R_j(t)$ can be expressed in terms of its equilibrium lattice position plus the small time dependent displacement: $R_j(t) = l_j + u_j(t)$. For this collection of nuclei, the operators u_j can be expressed in terms of creation and annihilation operators of the normal modes of the lattice:

$$u_j = \sqrt{\frac{\hbar}{2MN}} \sum_{s,q'} \frac{e_s(q')}{\sqrt{\omega_s(q')}} \left(a_s(q') e^{i[q' \cdot l_j - \omega_s(q')t]} + a_s^\dagger(q') e^{-i[q' \cdot l_j - \omega_s(q')t]} \right) \quad (\text{A.15})$$

where the normal modes have dispersion relations $E = \hbar\omega_s(q)$ and the creation and annihilation operators satisfy the usual bosonic commutation relations (here, we assume the displacements u_j are small, so that the motion is harmonic). Inserting this into eqn. A.9 and combining the exponentiated operators, it can be shown that

$$\frac{d^2}{d\Omega dE} = \frac{|k_f|}{|k_i|} \frac{e^{-2W(q)}}{2\pi\hbar} \sum_{jj'} b_j b_{j'} e^{iq \cdot (l_j - l_{j'})} \int_{-\infty}^{\infty} e^{\langle AB \rangle} e^{-i\omega t} dt \quad (\text{A.16})$$

with the operators A and B given by

$$\begin{aligned} A &= -iq \cdot u_j(0) \\ B &= iq \cdot u_j(t) \end{aligned} \quad (\text{A.17})$$

and the quantity $\exp(-2W(q))$ is the *Debye-Waller factor* which is due to the non-commutativity of the operators A and B .

Now, we consider the coherent ($j \neq j'$) and incoherent ($j = j'$) contributions to the above cross section. Because the exponentiated expectation value in eqn. A.16 is small, we Taylor expand it and examine each term individually. The zeroth-order term ($\exp(\langle AB \rangle) \approx 1$) corresponds to elastic scattering, as there is no time-dependence in this term. For this case, we arrive at the coherent and incoherent elastic cross-sections

$$\begin{aligned} \left. \frac{d\sigma}{d\Omega} \right|_{inc.} &= e^{-2W(q)} (\langle b^2 \rangle - \langle b \rangle^2) \\ \left. \frac{d\sigma}{d\Omega} \right|_{coh.} &= e^{-2W(q)} \sum_{\tau} |S_{\tau}|^2 \delta(q - \tau) \end{aligned} \quad (\text{A.18})$$

where the *static structure factor* S_{τ} is given in terms of the reciprocal lattice vectors τ of atoms at position d in the unit cell (with unit cell volume v_0) by

$$S_{\tau} = \frac{8\pi^3}{v_0} \sum_d b_d e^{i\tau \cdot d} \quad (\text{A.19})$$

We see that the incoherent cross section is mostly isotropic (aside from the q -dependence of the Debye-Waller factor), while the coherent scattering cross section depends strongly on the structure of the lattice: scattering occurs only when the Bragg condition ($q = \tau$) is satisfied, and the scattering intensity is enhanced due to the constructive interference of the scattering off of the different nuclei. From this we see that the Debye-Waller factor physically corresponds to the damping and broadening of the scattering intensity at wave-vector transfer q due to the thermal motion of the atoms about their mean position.

Now, the higher-order terms in the Taylor expansion of the expectation value in A.16 include the harmonic time dependence of the atoms in the crystal, and the powers of $\langle AB \rangle$ can be calculated using the creation/annihilation operator commutation relations. Each of these terms physically corresponds to the cross section to scatter from a normal mode of the lattice, with the first order term representing 1 phonon scattering, the second term to 2 phonon scattering, and so on. For example, the one-phonon coherent scattering cross section is

$$\begin{aligned} \frac{d^2\sigma}{d\Omega dE} &= \frac{|k_f|}{|k_i|} \frac{4\pi^3}{v_0} \sum_{sq'} \frac{1}{\omega_s(q')} \left| \sum_d \frac{\langle b_d \rangle}{\sqrt{M_d}} e^{-W_d(q)} e^{iq \cdot d} (q \cdot e_{d,s}(q')) \right|^2 \\ &\times \left[(n_s(q') + 1) \delta(E - \hbar\omega_s(q')) \sum_{\tau} \delta(q - q' - \tau) \right. \\ &\quad \left. + n_s(q') \delta(E + \hbar\omega_s(q')) \sum_{\tau} \delta(q + q' - \tau) \right]. \end{aligned} \quad (\text{A.20})$$

Here M_d is the mass of the d th atom in the unit cell, and the first (second) bracketed term corresponds to Stokes (anti-Stokes) processes of phonon emission (absorption) by the neutron. From the δ -functions, we see that the sum of the neutron and phonon energy and momentum must be conserved in the scattering process, and the scattering angle also depends on the phonon polarization.

The incoherent one-phonon scattering cross section represents the process of a neutron scattering from an individual atom, thus emitting or absorbing potentially many normal mode excitations such that energy (but not lattice momentum) is satisfied. The incoherent scattering cross-section can be written in terms of the *density of states* $Z(E)$ of the crystal, which counts the total number of modes between energy E and $E + dE$:

$$\frac{d^2\sigma}{d\Omega dE} = \frac{\hbar^2}{4M} \frac{|k_f|}{|k_i|} \frac{\sigma_{inc}}{4\pi} e^{-2W(q)} \langle (q \cdot e_s(q))^2 \rangle \frac{Z(E)}{E} [\coth(E/2kT) \pm 1]. \quad (\text{A.21})$$

Here, the $+$ ($-$) corresponds to Stokes (anti-Stokes) processes.

The derivations in this section consider the translational degrees of freedom for scatterers in

the unit cell of a crystal. This analysis can be generalized to consider other degrees of freedom, such as the nuclear spin or the orientation of rigid molecules, which also exhibit normal mode excitations (namely magnons and librons, respectively), and these excitations also contribute to the scattering processes described above. Further, anharmonic effects can be included perturbatively (see for example ref. [17]). In the harmonic approximation, scattering only occurs due to a given normal mode if $E = \hbar\omega(q)$ exactly. The anharmonicity in an inter-atomic or inter-molecular potential has the effect of (potentially) causing a small shift in the frequency, as well as broadening the dispersion relation. Thus, instead of a sharp peak in the scattering intensity at E for a fixed momentum transfer $\hbar q$, the intensity follows a Lorentz distribution in E with full-width-half-maximum 2Γ , where Γ^{-1} is the lifetime of the normal mode.

In this section, we have described the principles of neutron scattering from crystals that are necessary to understand the conversion of neutrons by crystalline solids to UCN energies, as well as the interaction of UCN with matter.

A.5 Ultracold Neutron Production

The UCN defined in table A.1 are typically generated by inelastically *downscattering* cold or thermal neutrons. It is not experimentally feasible to *moderate* neutrons to the ~ 1 mK temperature of UCN, so they are instead produced by non-equilibrium processes. Because UCN energies are $\sim 10^{-5}$ smaller than moderated neutron energies, a UCN converter must virtually stop a neutron, and it must be extracted from the converter before a second scattering event can *upscatter* it to its equilibrium energy.

One method forms the basis of the UCN source at the Institut Laue Langevin: neutrons are guided vertically from a D₂O moderated, reactor based neutron source, to a neutron-reflecting turbine that downscatters the neutrons. The velocity of the turbine blades is chosen to be similar to incident neutron velocities so that (nearly) all of the neutron's energy is taken away by the recoil of the blade.

However, other current and planned UCN sources rely on cryogenic liquid or solid converters.

A converter is chosen that exhibits vibrational modes whose dispersion relation intersects the *free neutron dispersion relation*:

$$E = \pm \frac{\hbar^2 q^2}{2m_n} \quad (\text{A.22})$$

where the plus sign corresponds to downscattering, and the minus sign to upscattering. A neutron incident upon such a material will have some probability of scattering with energy transfer $E_i \approx E$ and momentum transfer $\hbar q \approx \hbar k_i$, thus converting it to a UCN. In this way, the phase space density of the initial thermal neutrons can be partially concentrated at UCN energy, with the associated decrease in entropy compensated by the increased phonon (or other normal mode) entropy. To then extract the UCN from the source and into a UCN guide, the material must have a small probability of upscattering or capturing the UCN.

The downscattering cross section can be computed by starting with eqn. A.13 and parametrizing it along the free neutron dispersion relation:

$$\frac{d^2\sigma}{d\Omega dE} = \frac{\sigma_{tot}}{4\pi} \sqrt{\frac{E_f}{E_i}} S(\sqrt{2m_n E_i}/\hbar, E_i). \quad (\text{A.23})$$

Here, we have used the fact that $|k_f|/|k_i| = \sqrt{E_f/E_i}$. Integrating over solid angle is trivial since, in the above approximation, the momentum transfer is independent of the scattering angle. We can fix E_i and integrate over E_f , which is trivial since the cross section is proportional to $\sqrt{E_f}$. We thus arrive at the general form of the *downscattering* or *UCN production* cross section:

$$\sigma_{down}(E_i) = \frac{2}{3} \sigma_{tot} \frac{E_{UCN}^{3/2}}{E_i^{1/2}} S(\sqrt{2m_n E_i}/\hbar, E_i) \quad (\text{A.24})$$

where E_{UCN} represents the maximum usable UCN energy. Typically, one chooses the maximum UCN energy E_{UCN} that can be extracted or used at a UCN source (~ 100 to 300 neV). Similarly, we can write the upscattering cross section by again starting with eqn. A.13 and choosing the minus sign in the free neutron dispersion curve. For upscattering, we have that $E \approx E_f$, $q \approx k_f$,

and $E_i = E_{UCN}$ so that

$$\sigma_{up}(E_{UCN}) = \sigma_{tot} E_{UCN}^{-1/2} \int_0^\infty S\left(\sqrt{2m_n E}/\hbar, -E\right) dE. \quad (\text{A.25})$$

From this we see that, to a good approximation, $\sigma_{up} \propto v^{-1}$ as is the case for nuclear absorption.

A.6 Ultracold Neutrons

The magnetic, gravitational, and material surface interactions all contribute at UCN energy scales for laboratory-scale magnetic field strengths, distances, and materials. UCN are therefore easily guided, confined, and polarized in suitably designed experiments. This has lead to a broad experimental program to use UCN to perform precise measurements of the properties of the neutron, including the neutron electric dipole moment[33] and the β -asymmetry parameter A [57] in addition to the neutron lifetime. This section notes some of the properties of these interactions that are necessary to understand measurements of the neutron lifetime with UCN.

A.6.1 UCN and Material Surfaces

As discussed in A.2, the Fermi potential given in eqn. A.3 reflects the fact that a slow neutron wavelength is much longer than the size of a nucleus. For neutrons at ultracold energy, the neutron wavelength is also substantially larger than the inter-atomic spacing in typical solid materials ($\sim 1 - 10\text{\AA}$), and the coherent interaction between a UCN and solid material cannot result in Bragg scattering. As a result, we can interpret the motion of a UCN through a solid material as that of a neutron which is moving through a uniform potential V_O due to forward coherent scattering. This “optical” potential is computed by performing a spatial average over the Fermi potential due to each nucleus in the material:

$$V_O = 2\pi\hbar^2 m^{-1} n \bar{b}_{coh} \quad (\text{A.26})$$

where n is the number density of nuclei in the material, and \bar{b}_{coh} the average coherent scattering length.

The interaction of a UCN with a flat material surface can be modeled using the 1-D Schrödinger equation. The component of the UCN wave vector k_{\perp} normal to the surface can be used to compute an effective kinetic energy $K_{\perp} = \hbar^2 k_{\perp}^2 / 2m_n$, and for a surface occupying, say, $x > 0$, the neutron potential can be written as $V = V_O \Theta(x)$, so that the neutron hamiltonian is $H = K_{\perp} + V$, and the usual results from basic 1-D quantum mechanics apply. From this, one finds that the specular reflection of a UCN from a material surface is in practice quite probable, so that UCN can be stored in material traps for times comparable to the neutron lifetime.

This simple 1-D model captures the basic interaction of UCN and material surfaces. However, surface properties such as thin film deposits and surface roughness can lead to, for example, diffusive reflection of the UCN from surfaces, or changes in the expected transmission and reflection probabilities. Such effects are discussed in more detail in refs. [87] and [95].

Here, we note one important effect. Upon reflection of a UCN from a material surface, there is a partial overlap of the neutron wavefunction with the bulk of the material. There is therefore some probability of nuclear absorption or upscattering of the UCN upon reflection. Part of this effect is captured phenomenologically by the fact that coherent scattering amplitudes are in general complex: inserting a complex b_{coh} into the expression for V_O gives the UCN potential an imaginary piece, which causes some of the neutron wave's probability current entering the material to disappear. This accounts for the loss of UCN due to absorption. In order to more generally describe the loss of UCN upon reflection, the size of the imaginary part of the optical potential is modified to incorporate inelastic scattering of the UCN as well.

This effect is parameterized in terms of the ratio η of the imaginary part to real part of the optical potential, which roughly speaking is a *loss-per-bounce*, and can be small for many materials. To first order in η , the reduction in reflection probability due to loss is given by

$$R = 1 - 2\eta \sqrt{\frac{E_{\perp}}{V - E_{\perp}}}. \quad (\text{A.27})$$

Expressions such as this, accounting for the (UCN energy dependent) probability of loss upon bouncing from a material, is critical for the “material bottle” experiments described in chapter 2.

A.6.2 UCN and Gravity

The gravitational potential for a neutron on the earth’s surface is given by

$$V_G = m_n g h \quad (\text{A.28})$$

where h is the height of the neutron and $g = 9.80665 \text{ m/s}^2$ is the acceleration due to gravity.

We can rewrite this equation as

$$V_G = (100.429 \text{ neV/m}) h. \quad (\text{A.29})$$

A.6.3 UCN and Magnetic Fields

A neutron interacts with magnetic fields via its magnetic moment. In general, the spin and translational degrees of freedom are described quantum mechanically via the Hamiltonian operator $\sigma \cdot B$ (with σ the spin and B the potentially space and time dependent field strength). However, for a neutron in a magnetic field, if the rate of change of the magnetic field as seen by the neutron is slow compared to the neutron’s Larmor precession frequency

$$\frac{1}{|B|} \left| \frac{dB}{dt} \right| \ll \omega_L \quad (\text{A.30})$$

then the neutron spin will adiabatically follow the local magnetic field in which it is immersed.

This is known as the adiabatic condition, and can also be written as

$$\frac{\hbar}{\mu_N |B|^2} \left| \frac{dB}{dt} \right| \ll 1. \quad (\text{A.31})$$

For an inhomogeneous magnetic field $B = B(r)$, and local neutron velocity $v = dr/dt$, we have that

$$\frac{\hbar}{\mu_N |B|^2} v \cdot \nabla B \ll 1. \quad (\text{A.32})$$

As an example, for a 1 m/s UCN passing through a 1 T field with gradient of 1 T/cm (parameters relevant to the field profile of the magnetic trap discussed later), this gives $10^{-6} \ll 1$. Thus, a UCN in either a high-field seeking (spin aligned with B) or low-field seeking (spin anti-aligned with B) spin state will remain in that state and follow the local direction of the field. The interaction with the magnetic field has only the effect of the magnetic field gradient applying a force on the neutron, and the interaction can be written as

$$V_M = \pm \mu_N |B| \quad (\text{A.33})$$

or

$$V_M = \mp (60.2 \text{ neV/T}) |B| \quad (\text{A.34})$$

with the sign chosen according to the neutron spin state.

B Experimental Modeling

B.1 Introduction

B.2 Kinetic Theory Model of the Trap

Although the dynamics of the trap are complex, and over the course of a given measurement some time dependent phenomena may be observed, the theoretical equilibrium density is useful to gain some traction in understanding the occupied phase space of the trap. This is useful for providing simple estimates of some measured quantities. Here, we compute the equilibrium phase space density $\rho(z)$ of the trap.

If no information about the initial conditions remains, we hypothesize that a UCN of energy ϵ can access with equal probability the locations in the phase space that are energetically allowed. More precisely, for a UCN with an energy between ϵ and $\epsilon + d\epsilon$, the phase space density is a constant with respect to z if $0 < H(z) - \epsilon < d\epsilon$, and vanishes otherwise. Defining $B(x, a)$ to be 1 if $0 < x < a$ and 0 otherwise, we can write

$$\rho(z, \epsilon) = B(H(z) - \epsilon, d\epsilon)/A(\epsilon). \quad (\text{B.1})$$

We can fix $A(\epsilon)$ by demanding that $\rho(z, \epsilon)$ satisfies

$$\int_{\omega} \rho(z, \epsilon) d\omega = 1, \quad (\text{B.2})$$

which we can achieve by letting $A(\epsilon)$ be equal to the volume of phase space that satisfies $0 < H(z) - \epsilon < d\epsilon$. This can be written in terms of the volume $\Omega(\epsilon)$ of the phase space for

which $H(z) < E$:

$$A(\epsilon) = \Omega(\epsilon + d\epsilon) - \Omega(\epsilon) \quad (\text{B.3})$$

We can use this to write eqn. B.1 as

$$\rho(z, \epsilon) = \left[\frac{\Omega(\epsilon + d\epsilon) - \Omega(\epsilon)}{d\epsilon} \right]^{-1} \frac{B(H(z) - \epsilon, d\epsilon)}{d\epsilon}. \quad (\text{B.4})$$

If we take the limit $d\epsilon \rightarrow 0$, we have

$$\rho(z, \epsilon) = \left[\frac{d\Omega}{d\epsilon} \right]^{-1} \delta(\epsilon - H(z)). \quad (\text{B.5})$$




Figure B.1: The quantity $\omega(\epsilon) \equiv d\Omega/d\epsilon$ for typical UCN energies. The results of the Monte Carlo integration are reasonably approximated by a polynomial.

Now, the trap is in fact filled with a distribution of UCN energies. We denote this distribution as $G(\epsilon)$ (unit normalized over energy). We can then proceed to compute the phase space density for a trap with this distribution of energies:

$$\rho(z) = \int \rho(z, \epsilon) G(\epsilon) d\epsilon \quad (\text{B.6})$$

from which we find that

$$\rho(z) = \left[\frac{d\Omega}{d\epsilon}(H(z)) \right]^{-1} G(H(z)). \quad (\text{B.7})$$

Now, we must compute $d\Omega/d\epsilon$. As stated above, Ω is given by the phase space volume that satisfies $H(z) < \epsilon$, or

$$\Omega = \int_{\omega} \Theta(\epsilon - H(z)) d\omega. \quad (\text{B.8})$$

so that

$$\frac{d\Omega}{d\epsilon} = \int_{\omega} \delta(\epsilon - H(z)) d\omega. \quad (\text{B.9})$$

For our system with $H(z) = H(q, p) = |p|^2/2 + V(q)$, we have

$$\omega(\epsilon) = \int_Q \int_P \delta(\epsilon - \frac{1}{2}|p|^2 - V(q)) d^3p d^3q. \quad (\text{B.10})$$

where P and Q are the volumes of the spaces spanned by p and q respectively, and we have defined $\omega \equiv d\Omega/d\epsilon$. We can convert to polar coordinates in p :

$$\omega(\epsilon) = 4\pi \int_Q \int |p|^2 \delta(\epsilon - \frac{1}{2}|p|^2 - V(q)) d|p| d^3q. \quad (\text{B.11})$$

We can treat the argument of the δ -function as a function of $|p|$ and invert it:

$$\omega(\epsilon) = 4\pi \int_Q \int |p|^2 \frac{\delta(p - \sqrt{2(\epsilon - V(q))})}{\sqrt{2(\epsilon - V(q))}} d|p| d^3q. \quad (\text{B.12})$$

The integration over $|p|$ can now be performed, noting that the argument of the delta function has no roots in the domain of integration unless $\epsilon - V(q) > 0$, which gives

$$\omega(\epsilon) = 4\pi \int_Q \Theta(\epsilon - V(q)) \sqrt{2(\epsilon - V(q))} d^3q. \quad (\text{B.13})$$

This can be estimated via Monte Carlo integration of over d^3q : the integrand (call it f) is

evaluated M times at uniformly random values $q_i \in Q$ and we compute

$$\omega(\epsilon) \approx \frac{4\pi Q}{M} \sum_{i=1}^M f(q_i) \quad (\text{B.14})$$

This is shown in fig. B.1 for various values of ϵ . With this estimate, we can numerically compute $\rho(z)$ and quantities derived from it.

B.3 Neutron Tracking

B.3.1 Equations of Motion

The classical Hamiltonian for a neutron in the trap can be written as

$$H = \frac{1}{2}|p|^2 + V(q) \quad (\text{B.15})$$

where $q, p \in \mathbb{R}^3$ and $V(q) = \pm\mu|B|$. The plus or minus sign is chosen depending on whether the UCN is low field seeking (+) or high field seeking (−). Units are such that the neutron mass $m_n = 1$, lengths are in meters, times are in seconds, and magnetic fields are in Tesla. For this choice, one energy unit corresponds to approximately 10.454 neV, and $\mu \approx 5.768825731$.

The corresponding equations of motion for this Hamiltonian are

$$\dot{p} = F(q) \quad (\text{B.16})$$

$$\dot{q} = p \quad (\text{B.17})$$

where the dot represents a time derivative, and the force is $F(q) = -\nabla_q V(q)$. From eqns. 3.2 and 3.6, we can compute $F(q)$. Combining the expressions for the holding and trap fields, we

have that

$$\mathbf{B} = B_{h0} \frac{r+R}{\sqrt{y^2+z^2}} \hat{\xi} + \frac{4B_r}{\pi\sqrt{2}} \sum_{n=1}^{\infty} \frac{(-1)^{n-1}}{4n-3} (1 - \exp(-k_n d)) \times \exp(-k_n \zeta) \left(\sin k_n \eta \hat{\eta} + \cos k_n \eta \hat{\zeta} \right). \quad (\text{B.18})$$

Because the unit vectors ξ , η , ζ are orthogonal, we can straightforwardly compute $|\mathbf{B}|^2 = \mathbf{B} \cdot \mathbf{B}$:

$$|\mathbf{B}|^2 = B_{h0}^2 \frac{(r+R)^2}{y^2+z^2} + A^2 \sum_{n,m=1}^{\infty} C_n C_m \exp(-(k_n + k_m \zeta)) \cos(k_n - k_m) \eta \quad (\text{B.19})$$

with $C_n \equiv (-1)^n (1 - \exp(-k_n d)) / (4n - 3)$ and $A \equiv \sqrt{8} B_r / \pi$. To compute the force, we must find the derivatives of $|\mathbf{B}|$ with respect to x , y , and z . The gradient of the magnetic field magnitude is then:

$$\frac{\partial |\mathbf{B}|}{\partial x} = \frac{A^2}{2|\mathbf{B}|} \left(\frac{x}{\sqrt{(\rho - R)^2 + x^2}} \Xi - \frac{r(\rho - R)}{(\rho - R)^2 + x^2} \text{H} \right) \quad (\text{B.20})$$

$$\frac{\partial |\mathbf{B}|}{\partial y} = \frac{A^2}{2|\mathbf{B}|} \left(-\frac{2B_{h0}^2 (r+R)^2 y}{A^2 (y^2 + z^2)^2} + \frac{(\rho - R)y}{\rho \sqrt{(\rho - R)^2 + x^2}} \Xi + \frac{rxy}{((\rho - R)^2 + x^2) \rho} \text{H} \right) \quad (\text{B.21})$$

$$\frac{\partial |\mathbf{B}|}{\partial z} = \frac{A^2}{2|\mathbf{B}|} \left(-\frac{2B_{h0}^2 (r+R)^2 z}{A^2 (y^2 + z^2)^2} + \frac{(\rho - R)z}{\rho \sqrt{(\rho - R)^2 + x^2}} \Xi + \frac{rxz}{((\rho - R)^2 + x^2) \rho} \text{H} \right) \quad (\text{B.22})$$

where Ξ and H are given by

$$\Xi \equiv \sum_{n,m=1}^{\infty} (k_n + k_m) C_n C_m \exp(-(k_n + k_m) \zeta) \cos(k_n - k_m) \eta \quad (\text{B.23})$$

$$\text{H} \equiv \sum_{n,m=1}^{\infty} (k_n - k_m) C_n C_m \exp(-(k_n + k_m) \zeta) \sin(k_n - k_m) \eta \quad (\text{B.24})$$

and the quantities ζ and η are expressed in terms of x , y , and z .

B.3.2 Neutron Tracking Method

A simulation of trajectories in the trap must be accurate for times on the order of the neutron lifetime. In addition, the integration step size h must be quite small because the potential rapidly varies near the surface. Walstrom *et al.* used step sizes on the order of 10^{-6} seconds using the second-order leap-frog integration method[62]. In addition, energy conservation is critical: the *physical* energy change of a trapped UCN may be a source of loss (see sec. 6.7), and a numerical energy drift may artificially create this effect in simulations.

For this reason, symplectic integration is used to study the dynamics of neutron trajectories in the trap. Symplectic integrators are a subclass of geometric integrators, which are integration methods constructed to preserve a given geometric quantity in the phase space of a differential equation[102]. In particular, a symplectic integrator preserves the volume of the phase space (viz $dq \wedge dp$) under time evolution (as is the case for the underlying continuous system by virtue of Liouville's theorem). There remains a concomitant integration error (typically $\mathcal{O}(h^l)$ where l is the order of the integrator), but unlike more general schemes such as Runge Kutta integrators, the calculated trajectory remains near the manifold in the phase space defined by $E = H(q, p)$, so that energy is conserved to a good approximation.

```

for (i=0; i<order; i++)
{
    f = force(q);
    p = p + c[i]*f*dt;
    q = q + d[i]*p*dt;
}

```

Figure B.2: A procedural representation of the symplectic integration method for separable, time independent Hamiltonians.

The general numerical scheme for separable, time independent Hamiltonians is given in fig. B.2. To map a point (q, p) to that for the next time $t+h$, the momentum is shifted by an amount proportional to the force F at the original point q . The position q is then drifted proportional to p to its new value. This process is repeated l times, and the constants c_i and d_i are determined

(non-uniquely) by demanding that $dq \wedge dp$ is mapped to itself. Different solutions for c and d can exhibit different error terms, and solutions with some $c_i = 0$ are computationally beneficial, as it reduces the number of force evaluations per time step.

name	c	d	order
symplectic Euler	1	1	1
leapfrog	1/2, 1/2	1, 0	2
optimal	$1/\sqrt{2}, 1 - 1/\sqrt{2}$	$1/\sqrt{2}, 1 - 1/\sqrt{2}$	2
Ruth	$2/3, -2/3, 1$	$7/24, 3/4, -1/24$	3

Table B.1: Some solutions for the coefficients c_i and d_i . The third order Ruth integrator is from ref. [117].

Table B.1 lists some examples of symplectic integrators. An optimized fourth order method is used to integrate the equations B.16 and B.17[110]. This integration method is optimized for problems where the kinetic and potential terms in the Hamiltonian (and their derivatives) contribute to the integration error on an equal footing; given the fairly hard potential in the trap, it is unclear if this integration scheme is truly optimal for the problem at hand. Nonetheless, its performance is empirically determined to be sufficient. Fig. B.3 shows a comparison of this integration scheme compared to the Runge Kutta method.




Figure B.3: The relative change in particle energy for the pendulum problem using fourth order symplectic (left) and fourth order Runge-Kutta (right) algorithms. The symplectic method demonstrates the long-term stability of the energy.

B.3.3 Simulation Method

The neutron tracking algorithm is used to model the procedures performed in the experiment, such as loading the trap or cleaning. UCN can be created with arbitrary initial conditions, but are typically started at the 15×15 cm area at the bottom where the trap door panel is located. The neutron trajectory is then calculated, and its position is checked to see if it is incident upon any of the relevant components within the trap: in particular, the cleaner (in up or down positions), V foil, and trap door opening. In addition, UCN that spill out of the edges of the trap are identified. The components within the trap can be activated or deactivated in order to simulate a given experimental procedure. As an example, UCN cleaning is simulated by activating the trap door opening and the cleaner in its down position, and introducing UCN into the trap. After the trap has been filled, the loss through the trap door is deactivated and the cleaner remains in the down position. The trajectory of each neutron is stored in memory and saved to disk if desired. Because each neutron trajectory is independent, N separate simulations are executed on N processor cores, and the results of each simulation are combined later in analysis.




Figure B.4: The spatial trajectory of a trapped UCN (left and bottom panels), and the calculated energy of the neutron versus time.




Figure B.5: The spatial trajectory of a trapped UCN with energy near the maximum trappable energy (left and bottom panels), and the calculated energy of the neutron versus time.

Two example trajectories are shown in figures B.4 and B.5 with a 10^{-4} s time step. The spatial trajectory is shown along with the computed energy versus time. The energy shows a slow drift over time, with the small changes in energy typically compounding at each reflection from the strong field near the array surface. The energy drift is approximately diffusive in nature, growing as \sqrt{t} , with the proportionality depending on the time step and the particular trajectory. The trajectory in figure B.5 has an energy near the maximum trappable energy (~ 46 neV), whereas the trajectory in figure B.4 has much lower energy. The higher energy UCN exhibits greater energy drift, which can be as large as 10^{-4} relative drift after one hour of time in the simulation. Lower energy trajectories typically exhibit energy drifts less than 10^{-7} over this time span.

Bibliography

- [1] H. Abele. *Prog. in Part. Nucl. Phys*, 60:1, 2008.
- [2] V. P. Alfimenkov. *JETP Lett*, 52(7):373–378, 1990.
- [3] [Http://www.amptek.com/coolfet.html](http://www.amptek.com/coolfet.html).
- [4] G. P. Berman, V. N. Gorshkov, and V. I. Tsifrinovich. *Nucl. Instrum. Meth. A*, 592:385–392, 2008.
- [5] R. Bini. *J. Chem. Phys*, 104:4365, 1996.
- [6] S. Califano and V. Schettino. *Int. Rev. Phys. Chem*, 7(1):19–57, 1988.
- [7] N.B. Callahan. *Private communication*, 2014.
- [8] G. Cardini and S. F. O’Shea. *Phys. Rev. B*, 32:4, 1985.
- [9] J. Chadwick. *Proc. Roy. Soc. A*, 136, 1932.
- [10] J. Chadwick. *Nature*, 129:312, 1932.
- [11] J. Chadwick. *Proc. Roy. Soc. A*, 142, 1933.
- [12] J. Chadwick. *Nature*, 134:237, 1934.
- [13] V. Cirigliano, S. Gardner, and B. Holstein. *Prog. Part. Nucl. Phys*, 71:93, 2013.
- [14] V. Cirigliano, J. Jenkins, and M. Gonzalez-Alonso. *Nucl. Phys. B*, 830:85, 2010.
- [15] S.M. Clayton. *Private communication*, 2013.

- [16] N. D'Angelo. *Phys. Rev.*, 114(1):285–292, 1959.
- [17] M. T. Dove. *Introduction to Lattice Dynamics*. Cambridge University Press, 1993.
- [18] D. Dubbers and M. G. Schmidt. *Rev. Mod. Phys.*, 83:1111, 2011.
- [19] A. Anghel et al. *Nucl. Instrum. Meth. A*, 611:272–275, 2009.
- [20] A. Frei et al. *Phys. Rev. B*, 80:064301, 2009.
- [21] A. N. Sosnovsky et al. *Nucl. Phys.*, 10:395–404, 1959.
- [22] A. P. Serebrov et al. *Phys. Rev. C*, 78:035505, 2008.
- [23] A. Pichlmaier et al. *Phys. Lett. B*, 693:221–226, 2010.
- [24] A. Saunders et al. *Rev. Sci. Instrum.*, 84:013304, 2013.
- [25] A. Serebrov et al. *JETP Lett*, 50:11, 1994.
- [26] A. Serebrov et al. *JETP Lett*, 62:10, 1995.
- [27] A. Serebrov et al. *Phys. Lett. B*, 605:72–78, 2005.
- [28] A. Steyerl et al. *Phys. Rev. C*, 85:065503, 2012.
- [29] A. Steyerl et al. *Phys. Rev. C*, 86:065501, 2012.
- [30] A. T. Holley et al. *Rev. Sci. Instrum*, 83:073505, 2012.
- [31] A. T. Yue et al. *Phys. Rev. Lett.*, *accepted*, 2013.
- [32] B. Plaster et al. *Phys. Rev. C*, 86:055501, 2012.
- [33] C. A. Baker et al. *Phys. Rev. Lett*, 97:131801, 2006.
- [34] C. J. Christensen et al. *Phys. Rev. D*, 5(7):1628–1640, 1972.
- [35] C. L. Morris et al. *Nucl. Instrum. Meth. A*, pages 248–250, 2009.

- [36] C. M. O'Shaughnessy et al. *Nucl. Instrum. Meth. A*, 611:171–175, 2009.
- [37] C.-Y. Liu et al. *Nucl. Instrum. Meth. A*, 508:257–267, 2003.
- [38] D. Pocanic et al. *Phys. Rev. Lett*, 93:181803, 2004.
- [39] D. S. McGregor et al. *Nucl. Instrum. Meth. A*, 500:272–308, 2003.
- [40] E. Gutsmedl et al. *Europhys. Lett*, 96:62001, 2011.
- [41] E. Korobkina et al. *Nucl. Instrum. Meth. A*, 579:530, 2007.
- [42] E.I. Sharapov et al. *Phys. Rev. C*, 88:064605, 2013.
- [43] F. Atchison et al. *Phys. Rev. Lett*, 99:262502, 2007.
- [44] F. Atchison et al. *Nucl. Instrum. Meth. A*, 611:252–255, 2009.
- [45] G. Filippini et al. *Chem. Phys. Lett*, 40:2, 1976.
- [46] H. Abele et al. *Phys. Rev. Lett*, 88:211801, 2002.
- [47] J. Beringer et al. *Phys. Rev. D*, 86:010001, 2012.
- [48] J. Byrne et al. *Phys. Lett*, 92B(3, 4):274–278, 1980.
- [49] J. Byrne et al. *Phys. Rev. Lett*, 65(3):289–292, 1990.
- [50] J. Byrne et al. et al. *Europhys. Lett*, 33(3):187–192, 1996.
- [51] J. Last et al. *Phys. Rev. Lett*, 60(11):995–998, 1988.
- [52] J. S. Nico et al. *Phys. Rev. C*, 71:055502, 2005.
- [53] K. J. Coakley et al. *J. Res. Natl. Instrum. Stand. Technol*, 110:367–376, 2005.
- [54] L. N. Bondarenko et al. *JETP Lett*, 28(5):303–307, 1978.
- [55] M. Antonelli et al. *Eur. Phys. J.*, C69:399, 2010.

- [56] M. Ellis et al. *IEEE Nucl. Sci. Symposium and Medical Imaging Conf.*, pages 11–17, 2012.
- [57] M. P. Mendenhall et al. *Phys. Rev. C*, 87:032501, 2013.
- [58] M. S. Dewey et al. *Phys. Rev. Lett*, 91(15):152302, 2003.
- [59] P. E. Spivak et al. *Proceedings of the International Conference on the Peaceful Uses of Atomic Energy*. Geneva, Switzerland, 1955.
- [60] P. Geltenbort et al. *Nucl. Instrum. Meth. A*, 642:168–172, 2010.
- [61] P. J. Mohr et al. *Rev. Mod. Phys*, 84:1527, 2012.
- [62] P. L. Walstrom et al. *Nuc. Inst. Meth. A*, 599:1, 2009.
- [63] P. R. Huffman et al. *Nature*, 403:62–64, 2000.
- [64] R. Ouillon et al. *J. Chem. Phys*, 93:3005, 1990.
- [65] R. Picker et al. *Nucl. Instrum. Meth. A*, 611:297–301, 2009.
- [66] S. Arzumanov et al. *Phys. Lett. B*, 483:15–22, 2000.
- [67] S. Baessler et al. *C.R. Physique*, 12:729–754, 2011.
- [68] S. Materne et al. *Nucl. Instrum. Meth. A*, 611:176–180, 2009.
- [69] S. S. Arzumanov et al. *JETP Lett*, 95(5):224–228, 2012.
- [70] T. Bhattacharya et al. *Phys. Rev. D*, 85:054512, 2012.
- [71] U. Trinks et al. *Nucl. Instrum. Meth. A*, 440:666–673, 2000.
- [72] V. F. Ezhov et al. *J. Res. Natl. Inst. Stand. Tech*, 110:345–350, 2005.
- [73] V. F. Ezhov et al. *Nucl. Instrum. Meth. A*, 611:167–170, 2009.
- [74] V. V. Nesvizhevskii et al. *JETP Lett*, 75(3):405–412, 1992.

- [75] W. Mampe et al. *Phys. Rev. Lett*, 63(6):593–596, 1989.
- [76] W. Mampe et al. *JETP Lett*, 57(2):82–87, 1993.
- [77] W. Paul et al. *Z. Phys C*, 45:25–30, 1989.
- [78] Yu. G. Abov et al. *Sov. J. Nucl. Phys*, 38:70, 1983.
- [79] Yu. G. Abov et al. *Sov. Phys. JETP Lett*, 44(8):472–475, 1986.
- [80] Yu. Yu. Kosvintsev et al. *JETP Lett*, 31(4):236–240, 1980.
- [81] Yu. Yu. Kosvintsev et al. *JETP Lett*, 44(10):571–574, 1986.
- [82] Z. Wang et al. *To be published*, 2015.
- [83] E. Fermi. *Z Phys*, 88:161, 1934.
- [84] A. Furrer, J. Mesot, and T. Strässle. *Neutron Scattering in Condensed Matter Physics*. World Scientific, 2009.
- [85] R. Golub and K. Böning. *Z. Phys. B*, 51:95–98, 1983.
- [86] R. Golub and J. M. Pendlebury. *Rep. Prog. Phys*, 42:439–501, 1979.
- [87] R. Golub, D. Richardson, and S. K. Lamoreaux, editors. *Ultracold Neutrons*. Taylor & Francis, 1991.
- [88] D. A. Goodings. *Can. J. Phys*, 55:573–577, 1976.
- [89] K. Halbach. *Nucl. Instrum. Meth. A*, 169:1, 1980.
- [90] J. C. Hardy and I. S. Towner. *Phys. Rev. C*, 79:055502, 2009.
- [91] A. T. Holley. *Private communication*.
- [92] A.T. Holley, E.R. Adamek, and S.J. Seestrom. *Private communication*, 2014.

- [93] A.T. Holley and C.L. Morris. *Private communication*, 2014.
- [94] Junde Huo, Su Huo, and Chunhui Ma. *Nuclear Data Sheets*, 108:773, 2007.
- [95] V. K. Ignatovich. *The Physics of Ultracold Neutrons*. Oxford University Press, 1990.
- [96] J. D. Jackson, S. B. Treiman, and jr H. W. Wyld. *Phys. Rev*, 106:517–521, 1957.
- [97] R. W. Pattie Jr., K. P. Hickerson, and A. R. Young. *Phys. Rev. C*, 88:048501, 2013.
- [98] N. Kaiser. *Phys. Rev. C*, 64:028201, 2001.
- [99] J. K. Kjems and G. Dolling. *Phys. Rev. B*, 11:4, 1975.
- [100] G. F. Knoll. *Radiation Detection and Measurement*. Wiley, 4 edition, 2010.
- [101] R. L. Kouzes. *Pacific Northwest National Laboratory Report PNNL*, (18388), 2009.
- [102] B. Leimkuhler and S. Reich. *Simulating Hamiltonian Dynamics*. Cambridge University Press, 2005.
- [103] K.K.H. Leung and O. Zimmer. *Nucl. Instrum. Meth. A*, 611:181–185, 2009.
- [104] C.-Y. Liu, A. R. Young, and S. K. Lamoreaux. *Phys. Rev. B*, 62:6, 2000.
- [105] S. W. Lovesey. *Neutron Scattering from Condensed Matter*. Oxford University Press, 1984.
- [106] T. Luty and G. S. Pawley. *Chem. Phys. Lett*, 28:4, 1974.
- [107] W. J. Marciano and A. Sirlin. *Phys. Rev. Lett*, 96:032002, 2006.
- [108] G. J. Mathews, T. Kajino, and T. Shima. *Phys. Rev. D*, 71:021302, 2005.
- [109] D. S. McGregor and J. K. Shultis. *Nucl. Instrum. Meth. A*, 517:180–188, 2004.
- [110] R. I. McLachlan and P. Atela. *Nonlinearity*, 5:541–562, 1992.
- [111] H. Mutka. *Nucl. Instrum. Meth. A*, 338:144, 1994.

- [112] Yu.N. Pokotilovski. *Nucl. Instrum. Meth. A*, 356:412–414, 1995.
- [113] D. Richard, M. Ferrand, and G. J. Kearley. *J. Neutron Res*, 4:33–39, 1996.
- [114] J. M. Robson. *Phys. Rev.*, 78:311, 1950.
- [115] J. M. Robson. *Phys. Rev.*, 83:349, 1951.
- [116] S. Romano, B. Jönsson, and G. Karlström. *Int. J. Quantum Chem*, XXIII:991–998, 1983.
- [117] R. D. Ruth. *IEEE Trans. Nucl. Sci*, NS-30(4), 1983.
- [118] T. A. Scott. *Phys. Rep*, 27:89–157, 1976.
- [119] V. F. Sears. *Neutron Optics: An Introduction to the Theory of Neutron Optical Phenomena and their Applications*. Oxford University Press, 1989.
- [120] A. P. Serebrov and A. K. Fomin. *Phys. Rev. C*, 82:035501, 2010.
- [121] N. Severijns, M. Beck, and O. Naviliat-Cuncic. *Rev. Mod. Phys*, 78:991, 2006.
- [122] A. H. Snell and L. Miller. *Phys. Rev.*, 74:1217, 1948.
- [123] A. H. Snell, F. Pleasonton, and R. V. McCord. *Phys. Rev.*, 78:310–311, 1950.
- [124] P. E. Spivak. *JETP Lett*, 67(9):1735–1740, 1988.
- [125] G. L. Squires. *Introduction to the Theory of Thermal Neutron Scattering*. Dover Publications, 1997.
- [126] M. Srednicki. *Quantum Field Theory*. Cambridge University Press, 2007.
- [127] [Http://www.srim.org](http://www.srim.org).
- [128] I. S. Towner and J. C. Hardy. *Rep. Prog. Phys*, 73:046301, 2010.
- [129] V. F. Turchin. *Slow Neutrons*. Coronet Books, 1965.

- [130] R. G. Della Valle, G. F. Signorini, and P. Procacci. *Phys. Rev. B*, 55:22, 1997.
- [131] T. H. M. van den Berg, M. M. G. Bongers, and A. van der Avoird. *J. J. Phys.: Cond. Mat*, 2:8015, 1990.
- [132] V. V. Vladimirkii. *Sov. Phys. JETP*, 12:740–746, 1961.
- [133] Z. Wang and C. L. Morris. *Nucl. Instrum. Meth. A*, 652:323–325, 2011.
- [134] F. E. Wietfeldt and G. L. Greene. *Rev. Mod. Phys*, 83:1173, 2011.
- [135] J. A. Young and J. U. Koppel. *Phys. Rev*, 135:3A, 1964.

Education

- Ph.D. Experimental Nuclear Physics (4.0/4.0 GPA), Indiana University, 2011 – 2015 (advisor [Chen-Yu Liu](#)).
- M.S. Physics (4.0/4.0 GPA), Indiana University, 2009 – 2011.
- B.S. Physics and B.S. Mathematics with High Distinction (3.9/4.0 GPA), Indiana University, 2004 – 2008.

Research Experience

- From 2009 to present, I have conducted research and development for a measurement of the free neutron lifetime at Los Alamos Neutron Science Center using a magneto-gravitational ultracold neutron (UCN) trap. I have performed GEANT4 and custom simulations of the experiment, constructed of the apparatus, calculated and studied potential systematic effects, and developed new UCN detection technology. In addition, I developed data acquisition routines and developed automation software and hardware. I performed the data analysis of the initial commissioning campaign for the apparatus, and have taken leadership in the planning and execution of data-taking campaigns and R&D for the past two years.
- I was the principle investigator of an experimental study of dynamics of solid α - $^{15}\text{N}_2$ using time-of-flight inelastic neutron scattering at the Institut Laue-Langevin. From the analysis, I extracted the UCN production cross-section and UCN up-scattering cross-section of α - $^{15}\text{N}_2$.
- I performed an experiment with colleagues at IU and LANL to study solid oxygen as a source of UCN by direct measurement of UCN production on a cold neutron beamline. I developed and constructed the UCN detector, UCN transport and vacuum systems, developed slow control and online analysis routines, and performed GEANT4 simulations of neutron transport in the experiment.

Peer-Reviewed Publications

- D. J. Salvat, et al. Storage of ultracold neutrons in the magneto-gravitational trap of the UCN τ experiment. *Phys. Rev. C* **89**, 052501 (2014). Spring 2014 PRC editor's highlight.
- E. I. Sharapov, et al. Upscattering of ultracold neutrons from the polymer [C₆H₁₂]_n. *Phys. Rev. C* **88**, 064605 (2013).
- E. I. Sharapov, et al. Measurements of ultracold neutron upscattering and absorption in polyethylene and vanadium. *Phys. Rev. C* **88**, 037601 (2013).
- D. J. Salvat, et al. Investigating solid α -¹⁵N₂ as a new source of ultra-cold neutrons. *Europhys. Lett.* **103** (2013) 12001.
- D. J. Salvat, et al. A boron-coated ionization chamber for ultra-cold neutron detection. *NIM A* **691** (2012) 109 – 112.
- C. M. Lavelle, et al. Ultracold-neutron production in a pulsed-neutron beam line. *Phys. Rev. C* **82**, 015502 (2010).

Other Publications

- D. J. Salvat and P. L. Walstrom. Vibration-Induced Loss of Ultra-Cold Neutrons in a Magneto-Gravitational Trap. Proceedings of the Next Generation Experiments to Measure the Neutron Lifetime Workshop (2012).

Presentations

- *The UCN τ Experiment at LANL*. APS DNP Meeting, Waikoloa, Hi, 10 June 2014.
- *The UCN τ Experiment at LANL*. National Nuclear Physics Summer School, Williamsburg, VA, 18 June 2014.

- *Measuring the Neutron Lifetime with a Magnetic Neutron Trap at LANL.* Center for Non-linear Studies Seminar, LANL, Los Alamos, NM, 8 May 2014.
- *A Magneto-Gravitational Trap for the Measurement of the Neutron Lifetime at LANL.* Nuclear Physics Seminar, University of Kentucky, Lexington, KY, 24 April 2014.
- *Neutron Transport and Systematic Studies with the UCN τ Experiment at LANL.* APS [April Meeting](#), Savannah, GA, 6 April 2014.
- *Status of the UCN τ Experiment at LANL.* P-Division Seminar, ORNL, Oak Ridge, TN, 7 Nov 2013.
- *Analysis of the first data from the UCN τ experiment.* APS [DNP Meeting](#), Newport News, VA, 25 Oct 2013.
- *Status of the UCN τ Experiment.* Physics Division P-25 Seminar, LANL, Los Alamos, NM, 1 Oct 2013.
- *Status of the UCN τ Experiment.* Kellogg Radiation Laboratory Seminar, Caltech, Pasadena, CA, 22 August 2013.
- *Status of the UCN τ Experiment.* Nuclear Physics Seminar, North Carolina State University, Raleigh, NC, 21 June 2013.
- *Overview of the UCN τ Experiment.* APS [DNP Meeting](#), Newport Beach, CA, 26 October 2012.
- *A Magneto-Gravitational Trap for the Measurement of the Free Neutron Lifetime.* APS [DNP Meeting](#), East Lansing, MI, 27 October 2011.
- *UCN Production in Oxygen: Experimental Results.* [7th International UCN Workshop](#), St. Petersburg, Russia, 12 June 2009.

- *An Apparatus to Study Ultra-Cold Neutron Production in Solid Oxygen*. APS [April Meeting](#), Denver, Colorado, 2 May 2009.

Awards and Honors

- Indiana University College of Arts and Sciences [Dagmar K. Riley Graduate Fellowship](#), 2013 – 2014.
 - One year of full financial support for dissertation writing.
- Indiana University College of Arts and Sciences [James H. Coon Sciences Prize](#), 2013 – 2014.
 - Awarded to students for the pursuit of the sciences at Indiana University.
- Department of Energy [Office of Science Graduate Fellowship](#), 2010 – 2013.
 - Three years of full financial support and research allowance for graduate research.
- [Cora B. Hennel Mathematics Scholarship](#), 2007 – 2008.
 - Partial financial support awarded to undergraduates who have demonstrated high ability in mathematics.
- Hugh Brown Memorial Scholarship, 2007 – 2008.
 - Partial financial support awarded to undergraduates for academic performance in physics.
- Inducted into [Phi Beta Kappa](#), 2007.
- Outstanding Physics Undergraduate Instructor Award, 2006 – 2007.
- [Marie S. Wilcox Mathematics Scholarship](#), 2005 – 2006.
 - Partial financial support awarded to students who maintain a record of high academic achievement in mathematics.

Teaching and Outreach

- Volunteer science demonstrations at Wood-Gormley and Chamisa Elementary Schools, 2014 – present.
 - Participating in science demonstrations in two classrooms, two hours every two weeks, using basic “cardboard and duct-tape” experiments to teach principles of light, sound, mechanics, and electricity.
- Associate Instructor, Physics P201 & P202 Laboratories, Introduction to Physics I & II, 2009 – 2010.
 - Laboratory instructor and grader for introductory mechanics, electricity, and magnetism. Approximately 50 students per semester.
- Monroe County School Corporation [Volunteer Math/Science Tutor](#), 2009.
 - Homework and test preparation tutor for middle and high school students in mathematics, statistics, chemistry, and physics once per week in group sessions.
- Indiana University [Physics Forum](#) Volunteer Tutor, 2006 – 2010.
 - Homework and test preparation tutor for undergraduate students in all lower-level physics courses, once per week in group sessions.
- Undergraduate Instructor, Physics Q202 Laboratory, Physical Science for Education Majors, 2006 – 2008.
 - Laboratory instructor and grader for approximately 24 students per semester in basic physical science for students in the department of education.

Professional Activities

- Committees:

- Student representative for the Los Alamos Neutron Science Center user group executive committee, 2014 – present.
- Indiana University Department of Physics Student-Staff-Faculty Relations Committee, 2010 – 2011.
- Indiana University Department of Physics Outreach Committee, 2009 – 2010.
- Indiana University Department of Physics Undergraduate Curriculum Committee, 2007 – 2008.
- Conferences and workshops attended:
 - Fall Meeting of the APS Division of Nuclear Physics, Waikoloa, HI, 07 – 11 October, 2014.
 - American Physical Society April Meeting, Savannah, GA, 05 – 08 April, 2014.
 - Fall Meeting of the APS Division of Nuclear Physics, Newport News, VA, 23 – 26 October, 2013.
 - Fall Meeting of the APS Division of Nuclear Physics, Newport Beach, CA, 24 – 27 October, 2012.
 - DOE Science Graduate Fellowship Meeting, Brookhaven National Laboratory, 29 July – 01 August, 2012.
 - The 62nd [Lindau Nobel Laureate Meeting](#), Lindau, Germany, 01 – 06 July, 2012.
 - Fall Meeting of the APS Division of Nuclear Physics, East Lansing, MI, 26 – 29 October, 2011.
 - DOE Science Graduate Fellowship Meeting, Oak Ridge National Laboratory, 17 – 20 July, 2011.
 - [GEANT4 Tutorial](#), Oak Ridge National Laboratory, 08 – 11 March, 2011.
 - DOE Science Graduate Fellowship Meeting, Argonne National Laboratory, 08 – 10 August, 2010.

- [Summer School on Fundamental Neutron Physics](#), National Institute of Standards and Technology, 22 – 26 June, 2009.
- 7th International UCN Workshop, St. Petersburg, Russia, 08 – 14 June, 2009.
- American Physical Society April Meeting, Denver, Colorado, 02 – 05 May, 2009.