

Magnetic Trapping of
Ultracold Neutrons

A thesis presented

by

Clinton Reed Brome

to

The Department of Physics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Physics

Harvard University

Cambridge, Massachusetts

May 26, 2000

©2000 Clinton Reed Brome
All Rights Reserved.

To My Parents

Magnetic trapping of neutrons has been demonstrated using a Ioffe type superconducting magnetic trap surrounding an isotopically pure ^4He bath. Cold neutrons scatter to near rest in the helium and are confined by the magnetic field until they decay. Neutron beta-decay electrons scintillate in the helium bath, producing pulses of extreme ultraviolet light, which is downconverted and detected. Results are presented in which about 500 neutrons are magnetically trapped in each loading cycle, consistent with theoretical predictions. The lifetime of the observed signal, 660^{+290}_{-170} s is consistent with the neutron beta-decay lifetime, indicating no anomalous trap losses.

Acknowledgements

I am grateful to all of my family, but particularly my parents and my grandmother for their unfailing support and belief in me throughout my life. And more recently to the members of the Cambridge Cuisine Club, who weekly reminded me that there's more to a good meal than just good food.

The experiment described in this thesis was made possible by the efforts of a number of people over a number of years. I'm particularly grateful to the excellent post-docs on the experiment, James Butterworth and Paul Huffman, who taught the second-year graduate student who thought a 1K pot was a thousand ohm variable resistor most of what he now knows about designing and running a cryogenic experiment, among other things. And that's only the beginning of the many debts I owe Paul. I am grateful to Kevin Coakley, for several illuminating discussions, clarifying my thoughts on the analysis of data. To Bob Golub not only for his contributions to this work and my physics education, but also for his unique perspectives on life. And to Roxanne Huffman for reminding us all that there were things to talk about at meals other than our experiment. Last, but by no means least, I owe an enormous debt to my advisor, John Doyle, from whom I have learned far more than I expected to learn in graduate school.

I am also grateful to everyone who read one or more draft of my thesis. Every

change was an improvement. Any remaining flaws are my own.

Finally, I thank my colleagues in the basement of Lyman for puting up with my many eccentricities during the last stages of writing this thesis, as well as for interesting discussions, physics-related and otherwise.

Contents

Abstract	iv
Acknowledgements	v
List of Figures	xi
Chapter 1 Introduction	1
1.1 Neutron Beta Decay	1
1.2 Experimental Motivation	2
1.2.1 Big Bang Nucleosynthesis	2
1.2.2 Weak Interaction Physics	4
1.3 Measurement of the Neutron Lifetime	12
1.3.1 This Work	16
Chapter 2 Experimental Concepts	20
2.1 Magnetic Trapping	20
2.2 Superthermal Production of UCN	25
2.3 Detection of decays in the trap	31
2.4 Experimental Method	34

2.5	Loss Mechanisms	35
2.5.1	^3He Absorption Losses	35
2.5.2	Marginal Trapping	36
2.6	Backgrounds	37
2.6.1	Cosmic Rays	37
2.6.2	Gamma Rays	38
2.6.3	Neutron Activation	38
2.6.4	Luminescence	39
Chapter 3 Experimental Apparatus		41
3.1	Overview	41
3.2	Beam	42
3.2.1	Cold Neutron Guide NG-6	42
3.2.2	Beamline	43
3.2.3	Collimation	46
3.2.4	Beam Flux and Spectrum	48
3.3	Magnet	48
3.4	Refrigerator	54
3.5	Dewar	56
3.5.1	T-section	59
3.5.2	IVC End Cans	61
3.5.3	Dewar End Flanges and Windows	64
3.6	Alignment	67
3.7	250 mK Section	69
3.8	Neutron Shielding	76

3.9	Detection System	80
3.9.1	Calibration	91
3.10	Background Suppression	93
3.11	Data Acquisition	96
3.11.1	NIM	96
3.11.2	CAMAC Binner	99
3.11.3	Digital Oscilloscope	101
3.12	Computer Control	101
Chapter 4	Results	103
4.1	Overview	103
4.2	Processing of the Data	106
4.3	Backgrounds	110
4.3.1	Luminescence	112
4.4	Data Analysis	115
4.5	Non-trapping Data	119
4.6	Discussion	124
Chapter 5	Conclusions and Future Directions	126
5.1	Lifetime Measurement	126
5.1.1	Prospects for Improvement	126
5.1.2	Limits to the Method	128
5.2	Other Applications	129
Appendix A	Calculation of Neutron Beta Decay Rate	131

Appendix B Miscellaneous Math	138
B.1 Magnetic Field Maxima	138
B.2 Kinematics of 1-phonon Scattering	140
B.3 Buildup of Neutrons in the Cell	141
B.4 Attenuation of Neutron Absorbers	142
Appendix C Beam Calculations	144
C.1 Beam Spectrum at end of NG-6	144
C.2 Beam Divergence	148
C.3 Transmission Losses in the Guide	149
C.4 Measurement of the Beam Spectrum	150
Appendix D Calculated UCN Production Rate	155
Appendix E Activation Testing Results	159
Appendix F Some Computer Routines	161
F.1 Detection Efficiency Calculation	161
Appendix G Trap Losses	167
G.1 Marginal Trapping	167
G.2 Majorana Transitions	170
G.3 Absorption of Trapped UCN by ^3He	173
Appendix H Superfluid Helium and the Neutron Lifetime	174
Appendix I Useful Numbers	177
Bibliography	178

List of Figures

1.1	The neutron decays into a proton, an electron, and an electron anti-neutrino.	1
1.2	Sketch of the approximate four-fermion interaction.	5
1.5	Sketch of a storage experiment.	13
1.3	Plot of the 1σ constraints on $c_v V_{ud} $ and λ from various measurements.	17
1.4	Schematic of a beam lifetime measurement.	18
1.6	Plot of the measured values of the neutron lifetime found in the experiments described in the text.	18
1.7	A sketch of the idea of the NESTOR experiment.	19
2.1	A sketch of an anti-helmholtz trap.	22
2.2	A sketch of a spherical hexapole trap.	23
2.3	A sketch of a Ioffe trap.	24
2.4	The dispersion curves cross at $q = 0$ and at $q = q^*$, which corresponds to a neutron wavelength of 0.89 nm, or an energy corresponding to 11 K.	29
2.5	The beta decay spectrum of the neutron.	31
2.6	The molecular structure of TPB.	34

3.1	NIST Center for Neutron Research Hydrogen Cold Source[1]	43
3.2	NG-6 endstation.	44
3.3	A sketch of the collimation scheme.	47
3.4	Two configurations of racetrack coils giving a cylindrical quadrupole field.	49
3.5	The design of a racetrack coil.	50
3.6	Cross-sectional view of the magnet form.	51
3.7	The field profile in the x-y plane at $z = 0$ (left), and plotted along the x-axis ($y = z = 0$) (right).	52
3.8	The field profile in the x-z plane at $y = 0$ (bottom), and plotted along the z-axis ($x = y = 0$) (top), showing the “pinch” resulting from each solenoid, closing the ends of the trap.	53
3.9	A sketch of a dilution refrigerator.	54
3.10	A sketch of the T-shaped dewar.	57
3.11	A sketch of the inner vacuum can, inside the 4 K helium bath.	58
3.12	A sketch of the 300 K Sliding Seal.	59
3.13	A sketch of the 4 K Sliding Seal.	60
3.14	A sketch of the IVC end can.	61
3.15	Details of the IVC end can design.	63
3.16	Sketch of the teflon windows.	65
3.17	Sketch of the frame alignment mechanisms.	69
3.18	Sketch of a cell end cap with teflon window attached.	70
3.19	Plot of the thermal conductance of the upper section of the heat link (through the braid) as a function of the braid temperature.	73

3.20	Plot of the minimum cell temperature as a function of heat load, given by the thermal conductivity measurement shown in figure 3.19.	74
3.21	A sketch of the cell positioning scheme.	75
3.22	A sketch of a cell positioning device, in cross-section.	76
3.23	Sketch of the windows in the dewar end flanges, on the beam entry end, showing the interlocking boron nitride tubes shielding the beam.	78
3.24	A sketch of the interlocking boron nitride tubes.	79
3.25	Sketch of the windows in the dewar end flanges, on the light exit end.	81
3.26	A sketch of detector insert. (not to scale)	82
3.27	A sketch of the Vespel pegs.	83
3.28	A sketch of the aluminum “snout”	84
3.29	A sketch of the copper snout.	85
3.30	A sketch of the thermal connection between the 3-inch lightpipe and the 77 K end flange.	86
3.31	A sketch of the seal around the 3-inch lightpipe.	87
3.32	A sketch of the optical elements connecting the 3-inch lightpipe to the PMTs.	88
3.33	A histogram of pulse areas from the anode of a Burle 8850 PMT.	90
3.34	Pulse area spectrum from 360 keV beta line spectrum.	91
3.35	A sketch of a muon paddle.	95
3.36	Sketch of the orientation of three muon paddles relative to the hori- zontal section of the dewar.	95
3.37	Data Acquisition Schematic	97

4.1	A plot of the predicted trapped neutron population over the course of a data run. Three times are indicated: the <i>start</i> of the run ($t=0$ s), the times at which the neutron beam is turned <i>on</i> ($t=100$ s) and <i>off</i> ($t=1450$ s). The run ends at $t=5050$ s.	104
4.2	A plot of the trapping signal data.	111
4.3	The counting rate with a one-photoelectron threshold. Also shown is the actual coincidence rate of this signal, and the expected rate from the accidental overlap of uncorrelated events. (from a trapping run in the first data set)	113
4.4	Plot of the difference between the averaged “trap-on” and the averaged “trap-off” data for the two reactor cycles. The solid lines indicate the result of the fits described in the text.	116
4.5	Plot of the minimum χ^2 as a function of τ	118
4.6	Plot of the minimum χ^2 of the fit of the data to equation 4.1 as a function of A_1 and A_2	119
4.7	Plots of the difference between the average pooled magnetic “trap-on” and “trap-off” in the 1.2 K “warm” data.	122
4.8	Plots of the difference between the average pooled magnetic “trap-on” and “trap-off” in the ^3He data for each reactor cycle.	123
B.1	1-phonon scattering	140
C.1	Sketch of diverging beam, for one wavelength, with $\theta_{\text{max}} = \sin^{-1}(\lambda/494 \text{ \AA})$	148
C.2	Schematic of the time-of-flight measurement of the neutron spectrum.	151
C.3	Data from the time-of-flight measurement.	152

Chapter 1

Introduction

1.1 Neutron Beta Decay

As the lightest configuration of two up quarks and a down quark, a neutron is stable under the strong and electromagnetic interactions, which conserve quark flavor. Though quarks of different flavors can couple with the charged weak-interaction gauge boson, the W , under the Standard Model the weak interaction conserves baryon number. Thus the second lightest baryon, the neutron, can decay only to the lightest baryon, the proton. The mass difference between the neutron and the proton




Figure 1.1: The neutron decays into a proton, an electron, and an electron anti-neutrino.

$(m_n(939.5656 \text{ MeV}/c^2) - m_p(938.2720 \text{ MeV}/c^2) = 1.2936 \text{ MeV}/c^2)$ is small enough (i.e. much less than the mass of a charged pion, $139 \text{ MeV}/c^2$, or a muon, $106 \text{ MeV}/c^2$), that the only possible decay products for the W^- boson are an electron and an electron anti-neutrino (figure 1.1). The mass of the electron ($m_e = 511 \text{ keV}/c^2$) leaves a net mass difference in the decay of $782.6 \text{ keV}/c^2$, which is released as kinetic energy of the decay products. Since this energy is comparable to the binding energies of protons and neutrons in nuclei, the decay energy (and thus lifetime) of bound neutrons depends strongly on the nuclear environment of the neutron. In many nuclei, the neutron is stable, while the free neutron decays with a lifetime of about 15 minutes.

1.2 Experimental Motivation

The beta-decay lifetime of the neutron is of interest in a number of contexts. It has a direct cosmological impact, as an important input into the calculation of the production of light elements during the Big Bang (Big Bang Nucleosynthesis). Beta decay of the neutron is also both the simplest nuclear decay and more generally the simplest of the charged-current weak interactions in baryons. Thus, measurements of the parameters of the charged-current weak interaction can be measured using neutron beta decay with fewer and simpler theoretical corrections than measurements using the beta decay of nuclei.

1.2.1 Big Bang Nucleosynthesis

In our current understanding of the history of the universe, the matter and energy in the universe existed in a small region at incredibly high temperatures ($T \approx 10^{12} \text{ K} \approx 100 \text{ MeV}$ at $t = 10^{-4} \text{ s}$) between ten and fifteen billion years ago. The

development of the early universe as it expanded and cooled can be thought of as a series of stages characterized by the interactions permitted at the corresponding temperature and density at that time after the Big Bang (i.e. the collision energy and frequency at that time). At temperatures $T \gg 1$ MeV ($t \ll 1$ s) neutrons and protons were in thermal equilibrium, ($N_n/N_p \approx e^{-Q/T} \approx 1$, where $Q/c^2 \approx m_n - m_p$). As the universe expanded and the collision energy fell, the weak-interaction scattering which maintained this equilibrium grew weaker and the neutron:proton ratio (N_n/N_p) began to fall. Shortly before $t = 1$ s, at $T \approx 2.2$ MeV, deuterium nuclei (formed by the reaction $n + {}^1\text{H} \rightarrow {}^2\text{H} + \gamma$) began to survive collisions with other particles. However the deuterium was still consumed (by photodissociation) as quickly as it was created until $T \approx 0.1$ MeV, at which temperature deuterium began to accumulate. As the temperature of the early universe fell further, other light nuclei, ${}^3\text{He}$, ${}^7\text{Li}$, and especially ${}^4\text{He}$ (nearly a quarter of the baryonic matter in the early universe) formed. Over the course of the next several billion years, some of these gas clouds condensed to form stars, the heavier elements, and eventually us.

Although physical interactions at the energy scales of the early universe can be studied in the lab, the Big Bang itself is obviously not reproducible in the lab. As such, the experimental investigation of the Big Bang is confined to observation of its products: properties of the universe as a whole (its total mass and expansion rate), the cosmic background radiation (the dramatically red-shifted light from the early universe), and measurements on the primordial clouds of matter (such as their isotopic abundances).

The isotopic ratios measured in extragalactic gas clouds (in which isotopic abundances have remained constant since the Big Bang, since no stars have formed) can be compared to calculations in order to place limits on the ratio of baryons to

photons in the early universe. In the calculation of the expected ${}^4\text{He}/{}^1\text{H}$ ratio, the dominant uncertainty is in the decay lifetime of the neutron[2, 3].

1.2.2 Weak Interaction Physics

In very general terms, the rate of a decay or interaction is given by a matrix element (given by the initial and final wavefunctions coupled by the interaction Hamiltonian) integrated over the available phase space. For electroweak processes, the relevant Hamiltonian can be considered an exhaustive list of all possible point interactions. The matrix element for a process can then be represented as a series of Feynman diagrams, each of which represents one set of such point interactions which connects the initial and final states. Electromagnetic and weak interactions are particularly amenable to representation in a power expansion. Each “vertex” connecting the lines which represent the particles in the interaction decreases the amplitude of the matrix element. The leading order term in a charged-current weak interaction (like the beta decay of the neutron) can be represented by a Feynman diagram like that shown on the left in figure 1.2. “Higher order” diagrams (i.e. those with additional electromagnetic or weak vertices) are successively smaller corrections to the rate. For the beta-decay of the neutron, corrections and higher order diagrams alter the decay rate by less than ten percent from the rate one can calculate using only this first-order diagram. The rate given by this diagram is calculated explicitly in appendix A, in terms of the effective weak force coupling constants g_v and g_a . The full calculation of the beta-decay rate includes the higher order terms grouped together in a corrected phase space factor ($f \rightarrow f^R(1 + \delta_R)(1 + \Delta_R^V)$, see appendix A).

These effective coupling constants can be related to the fundamental weak

force coupling constant. First, in the limit of low momentum transfer ($q^2 \ll M_W^2$), the “propagator” term representing the exchange boson in the matrix element for a charged-current weak interaction is reduced to a constant. This constant can be absorbed into the coupling constant of an effective four-fermion point interaction. The full matrix element,

$$|V_{if}| = \frac{ig}{2\sqrt{2}}(\bar{\psi}_1\gamma_\mu(1-\gamma_5)\psi_2)i\frac{g_{\mu\nu}-q_\mu q_\nu/M_W^2}{M_W^2-q^2}\frac{ig}{2\sqrt{2}}(\bar{\psi}_3\gamma_\nu(1-\gamma_5)\psi_4), \quad (1.1)$$

goes, in the limit as $q^2/M_W^2 \rightarrow 0$ to:

$$\begin{aligned} |V_{if}| &\approx \frac{ig}{2\sqrt{2}}(\bar{\psi}_1\gamma_\mu(1-\gamma_5)\psi_2)i\frac{g_{\mu\nu}}{M_W^2}\frac{ig}{2\sqrt{2}}(\bar{\psi}_3\gamma_\nu(1-\gamma_5)\psi_4) \\ &\approx -i\frac{g^2}{8M_W^2}\bar{\psi}_1\gamma_\mu(1-\gamma_5)\psi_2\bar{\psi}_3\gamma_\mu(1-\gamma_5)\psi_4 \\ &\approx -i\frac{G_F}{\sqrt{2}}\bar{\psi}_1\gamma_\mu(1-\gamma_5)\psi_2\bar{\psi}_3\gamma_\mu(1-\gamma_5)\psi_4, \end{aligned} \quad (1.2)$$

where $\frac{G_F}{\sqrt{2}} \equiv \frac{g^2}{8M_W^2}$ appears as the effective strength of the four-fermion coupling.




Figure 1.2: Sketch of the approximate four-fermion interaction.

In the hadronic sector, the charged weak current, $\frac{ig}{2\sqrt{2}}\bar{\psi}_1\gamma_\mu(1-\gamma_5)\psi_2$ must be modified in two ways. The first is that while the charged weak force couples the

leptonic isospin doublets $\begin{pmatrix} \nu_i \\ \ell_i^- \end{pmatrix}$ directly, it couples the quark doublets as $\begin{pmatrix} u_i \\ d'_i \end{pmatrix}$, where $d'_i \equiv \sum_j V_{ij} d_j$, u_i (d_i) is the mass eigenstate of the $q = +\frac{2}{3}$ ($-\frac{1}{3}$) quark of the i th generation, and V_{ij} corresponds to element ij of the Cabibbo-Kobayashi-Maskawa (CKM) quark mixing matrix. The second way in which the hadronic current differs from the leptonic current is a consequence of the fact that quarks are always found bound together in color-neutral sets (a hadron), either three quarks (a baryon) or a quark-anti-quark pair (a meson). Thus, in any real charged weak interaction the initial and final wavefunctions represent complete hadrons rather than free quarks. In the leptonic charged weak current, the parity-even vector term ($\bar{\psi}_1 \gamma_\mu \psi_2$) and parity-odd axial-vector term ($\bar{\psi}_1 \gamma_\mu \gamma_5 \psi_2$) have equal magnitude, 180° out of phase.*

Thus, the general hadronic current can be written as:

$$H_\mu = \frac{ig|V_{ij}|}{2\sqrt{2}} \bar{\psi}_1 \gamma_\mu (c_v + c_a \gamma_5) \psi_2, \quad (1.3)$$

where c_v and c_a are the strengths of the vector and axial-vector couplings for the given hadrons relative to the standard current. In keeping with the four-fermion approximation described above, these terms are absorbed into effective coupling constants $g_v \equiv c_v V_{ud} G_F$ and $g_a \equiv c_a V_{ud} G_F$.[‡] More precisely, g'_v and g'_a are measured in interactions, where the unprimed constants represent the “true” constant, and the primed constant represents the experimentally accessible value, which of course includes the effects of radiative corrections (i.e. higher order diagrams). According to the conserved vector current (CVC) hypothesis, $c_v = 1$ for all interactions. c_a however is

*In nuclear physics, the parity-even and -odd terms are called the Fermi and Gamow-Teller components of the decay.

[‡]Standard usage for V_{ij} replaces the first index with the abbreviation for the i th $q = +\frac{2}{3}$ quark and the second index with the j th $q = -\frac{1}{3}$ quark. Thus, for example, V_{12} which relates d' to s , and thus is the coefficient of the term coupling u to s in the expansion of the u to d' term, is referred to as V_{us} .

only partially conserved. In leptonic currents $c_a = -1$, while in the hadronic current for the free neutron decay $c_a = -1.2670 \pm 0.0035$ [3].

Proper normalization of the weak force eigenstates $\sum_j V_{ij} d_j$ requires that the CKM matrix be unitary, that is $V^\dagger V = I$. The first element of this matrix of equations gives:

$$(V^\dagger V)_{11} = |V_{ud}|^2 + |V_{us}|^2 + |V_{ub}|^2 = 1. \quad (1.4)$$

All three of these terms have been directly measured in semi-leptonic charged-current weak decays of hadrons. That is, they have been measured in decays represented by a matrix element in the form of equation 1.1 with one leptonic current and one hadronic current. These measurements are summarized in Table 1.1, along with their contributions to the accuracy with which equation 1.4 can be tested.

Table 1.1: Measurements of the First Row of the CKM Matrix

Term	$ V_{ij} \pm d V_{ij} $	$2 V_{ij} dV_{ij} $
$ V_{ud} $	0.9740 ± 0.0010	2.0×10^{-3}
$ V_{us} $	0.2196 ± 0.0023	1.0×10^{-3}
$ V_{ub} $	0.0033 ± 0.0008	5.3×10^{-6}

From the rates of the decays $B \rightarrow \pi \ell \nu_\ell$ and $B \rightarrow \rho \ell \nu_\ell$, $|V_{ub}|$ was determined to be $3.3 \pm 0.4 \pm 0.7 \times 10^{-3}$ [4]. Similarly, from the rate of the beta decay of the kaon, $K^+ \rightarrow \pi^0 e^+ \nu_e$, $|V_{us}|$ was determined to be 0.2196 ± 0.0023 [5]. $|V_{ud}|$ has been measured to be 0.9740 ± 0.0010 in nuclear beta decays, as described below. Using these measurements, the value of the sum in equation 1.4 is given by $(V^\dagger V)_{11} = 0.9969 \pm 0.0022$, a 1.5σ discrepancy. Since most of the uncertainty in the sum arises

from the uncertainty in $|V_{ud}|$, an improved measurement of $|V_{ud}|$ will be needed to test this discrepancy.

We can turn this argument around, and use the requirement of equation 1.4 to “predict” a value of $|V_{ud}|$:

$$|V_{ud}|^2 = 1 - |V_{us}|^2 - |V_{ub}|^2 \quad (1.5)$$

$$= 1 - (0.2196 \pm 0.0023)^2 - (0.0033 \pm 0.0008)^2$$

$$= 0.9518 \pm 0.001$$

$$|V_{ud}| = 0.9756 \pm 0.0005 \quad (1.6)$$

A measurement of $|V_{ud}|$ different from equation 1.6 (i.e. a measured value of $(V^\dagger V)_{11}$ different from 1) would imply either that the CKM matrix is incomplete and has additional terms, that the value of $|V_{us}|$ or $|V_{ub}|$ is incorrect, or that one of the assumptions connecting the measurement to $|V_{ud}|$ must be wrong (for example, CVC or the assumption that the Standard Model electroweak Lagrangian has no additional terms). The best measurements to date of $|V_{ud}|$ have been made using nuclear beta-decays, but it can also be measured using the rare decay $\pi^+ \rightarrow \pi^0 e^+ \nu_e$ or neutron beta decay (described below).

In nuclear beta decay, the vector term corresponds to initial and final nuclei with the same spin state and the axial-vector term corresponds to initial and final nuclei with the opposite spin state. In the case of the coupling between two spin-0 nuclei, there is no spin-flip component (i.e. if a spin flipped, it would have to be a coupling between a spin-0 and a spin-1 nucleus, often with a considerably higher energy) and the hadronic current for such a decay will contain no axial-vector term. In addition, since the number of nucleons and the spin are both the same, the strong-force component of the initial and final wavefunction must be the same, though the

electromagnetic component will clearly differ. This makes the calculation of the rate of such “ $0^+ \rightarrow 0^+$ ” nuclear beta decays almost as simple as that of the free neutron.

If it is true that $c_v = 1$ (CVC), then g_v is the same for all nuclear beta decays, and the matrix element for every pure vector beta decay will be the same. Thus, the half-life of every such decay, corrected for the phase space available to the decay products (i.e. the different energy of the decay) will be the same (in the jargon, Ft must be the same, where F is the calculated integral over the phase space, including radiative corrections, and t is the measured half-life). Experimental support for CVC can be found in the fact that measurements of Ft for a variety of nuclei all agree at the level of 4×10^{-4} [6]. Measurements of Ft , assuming that $c_v = 1$ and taking G_F from a pure leptonic decay ($\mu \rightarrow \nu_\mu e \bar{\nu}_e$) give a value for $|V_{ud}|$ of 0.9735 ± 0.0005 [3, 7]. However, it has been argued[8] that an additional correction would increase this value by up to 0.2%, giving a value of 0.9740 ± 0.0010 .[§] While the first value differs from equation 1.6 by nearly three standard deviations (3σ), the second differs by less than 1.5σ . In either case, the uncertainty in $|V_{ud}|$ from such measurements is dominated by the theoretical calculations of the F values, not the measurements of t . Thus, to make an improved measurement of V_{ud} , it would be advantageous to use a system in which the theoretical corrections are simpler: the beta-decay of the pion or the neutron.

Pion beta decay $\pi^+ \rightarrow \pi^0 e^+ \nu_e$ is a pure vector decay (both charged and neutral pions are spin-0 particles). The pion is a simpler system than a nucleus, and it is claimed that the radiative corrections have been calculated to 0.15%[9], accurate enough to provide a check on the measurements from nuclear $0^+ \rightarrow 0^+$ beta decays.

[§]There doesn’t appear to be any consensus on this question. However, most of the papers on the question are unpublished, appearing only on the preprint server at xxx.lanl.gov, including one of the two measurements incorporated in the average in reference [3].

The charged pion usually decays by $\pi^+ \rightarrow \mu^+ \nu_\mu$, and the beta decay has a branching ratio of only $1.026 \pm 0.039 \times 10^{-8}$ [10]. Thus, any experiment trying to measure the rate $\Gamma_{\pi^+ \rightarrow \pi^0 e^+ \nu_e}$ has to deal with 10^8 pion decays in their detector for every pion beta decay. Even with a perfect ability to distinguish beta-decay events from other pion decays, a statistics-limited 0.3% measurement of the pion beta-decay rate would require counting 10^5 beta-decay events, and thus registering (or at least vetoing) a total of 10^{13} pion decays. A collaboration at the Paul Scherrer Institute (PSI) proposes to register the occurrence of pion decays at a rate of $5 \times 10^6 \text{ s}^{-1}$ (compared to 10^4 s^{-1} in the experiment in reference [10]), at which rate it would take $2 \times 10^7 \text{ s}$, or about a month, to record this many events[9].

Finally, $|V_{ud}|$ can be measured using the beta decay of the neutron. The theoretical corrections to the neutron beta-decay rate have been calculated to 1 part in 10^4 [11], considerably better than for pion beta decay or $0^+ \rightarrow 0^+$ nuclear decays. However, unlike either of these processes, the decay of the neutron is not a pure vector process. Both the neutron and the proton are spin-1/2 particles, and so both vector (spin-unflipped) and axial-vector (spin-flipped) couplings contribute significantly to the decay. The neutron beta-decay rate is calculated in appendix A. From that appendix, equation A.21 gives the neutron beta-decay rate in terms of several measurable parameters as:

$$\frac{d\Gamma_{n \rightarrow pe\bar{\nu}_e}}{dE_e d\Omega_e d\Omega_\nu} \propto (g_v^2 + 3g_a^2)F(E_e)\left(1 + a\frac{(\vec{p}_e \cdot \vec{p}_\nu)}{E_e E_\nu} + A\frac{(\vec{\sigma}_n \cdot \vec{p}_e)}{E_e} + B\frac{(\vec{\sigma}_n \cdot \vec{p}_\nu)}{E_\nu}\right), \quad (1.7)$$

where the correlation coefficients a , A and B are given by:

$$a \equiv \frac{1 - \lambda^2}{1 + 3\lambda^2}, \quad (1.8)$$

$$A \equiv -2 \frac{\lambda^2 + \lambda}{1 + 3\lambda^2}, \quad (1.9)$$

$$B \equiv 2 \frac{\lambda^2 - \lambda}{1 + 3\lambda^2}, \quad (1.10)$$

$$\text{and } \lambda \equiv \frac{g_a}{g_v}, \quad (1.11)$$

and $F(E_e)$ (given in equation A.19) contains the dependence on the electron's energy.

A measurement of the neutron lifetime corresponds to a determination of the quantity $g_v^2 + 3g_a^2$. To extract g_v requires either an independent measurement of g_a or of the more experimentally accessible ratio $g_a/g_v \equiv \lambda$. A measurement of any one of the correlation coefficients a , A , or B is sufficient to determine λ . At present, the most precise measurements of λ comes from measurements of A . However, the four published values are not consistent with one another.[¶] Shown in figure 1.3 are the 1σ regions for selected measurements of c_a/c_v and $c_v|V_{ud}|$ (or equivalently g_v , just dividing out $G_F = 1.16639 \times 10^{-5} \text{ GeV}^{-2} = 1.4359 \times 10^{-62} \text{ Jm}^{-3}$). Although the $0^+ \rightarrow 0^+$ measurements seem theoretically uncertain and the variations between different A measurements seem to indicate poorly understood systematic effects, a discrepancy in the measured value of $c_v V_{ud}$ would imply either a violation of CVC, a larger CKM matrix, or additional terms in the electroweak Lagrangian.

[¶]The four measurements included in the average in reference [3] have values of -0.1189 ± 0.0012 [13], $-0.1160 \pm 0.0009 \pm 0.0012$ [14], -0.1135 ± 0.0014 [12] and -0.1146 ± 0.0019 [15].

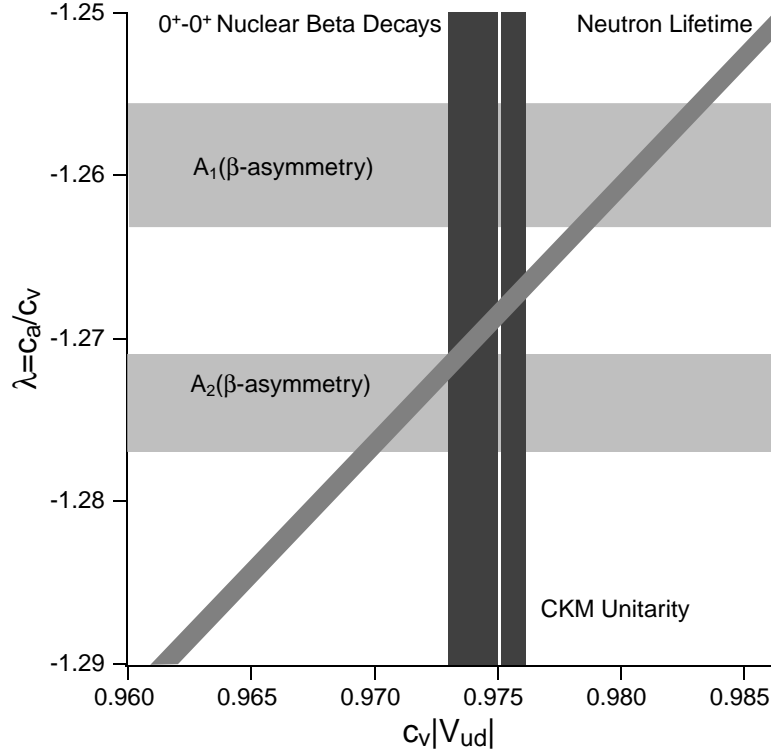


Figure 1.3: Plot of the 1σ constraints on $c_v|V_{ud}|$ and λ from various measurements. The constraint on $|V_{ud}|$ assuming CVC, CKM Unitarity and the measured values for $|V_{us}|$ and $|V_{ub}|$ is shown. Also shown are constraints from neutron lifetime measurements, $0^+ \rightarrow 0^+$ nuclear beta decays and two of the four measurements of the β -asymmetry parameter, A . The larger error bars of reference [8] are shown for $0^+ \rightarrow 0^+$ decays. A_1 and A_2 refer to the A measurements in references [12] and [13], respectively. The other two A measurements would fall between these two.

1.3 Measurement of the Neutron Lifetime

In order to measure the lifetime of a radioactive isotope, one must either independently measure both the rate of decay and the number of radioactive nuclei in the sample, or measure the change in the rate of decay in a closed sample, in which the number of radioactive nuclei is changing only by their decay. In general,

the first method is more difficult, requiring either an isotopically pure sample with known mass or else a means of accurately measuring the isotopic composition of the sample along with its mass. Measuring the change in the decay rate is generally much simpler, but can be impractical for decays with extremely long lifetimes.

Techniques analogous to each of these methods have been used to measure the neutron lifetime. With a lifetime of about 15 minutes, the neutron seems to be an ideal candidate for the closed-sample method. Unlike radioactive isotopes, however, neutrons do not form atoms, and so are harder to keep in one place. Since neutrons decay so rapidly, they are only available near facilities which either create neutrons (such as a particle accelerator) or release them from heavy nuclei (such as a nuclear reactor or a spallation source). In general, the highest average intensity sources of neutrons are nuclear reactors and the most precise neutron lifetime measurements to date have been performed at reactor facilities. Measurements of the neutron lifetime can be divided into beam measurements, analogous to the first method described above and storage measurements, analogous to closed-sample measurements.





Figure 1.4: Schematic of the apparatus used in the beam lifetime measurement in reference [16]. A ^{10}B target and α -detectors measure the beam flux, while decay protons are stored in an electromagnetic Penning trap and then counted.

In a beam measurement, a well collimated neutron beam passes continuously through a decay region of known length, and each neutron decay occurring in that volume is counted. At the end of this volume, the flux of the neutron beam is measured, weighted inversely as the velocity of the neutrons. This weighting corrects for the different amount of time spent in the decay region by neutrons of different velocities to give the total number of neutrons per unit length in the decay region. The beta-decay rate of the neutron is just given by the ratio of number of decays per unit time to the number of neutrons in the decay region. In the most precise measurement using this method ($\tau_n = 889.2 \pm 4.8$ s [17, 16], see figure 1.6), neutron decays were counted by detecting the decay protons and the beam flux was measured by detecting α -particles emitted in the $n + {}^{10}\text{B} \rightarrow {}^7\text{Li} + \alpha$ reaction (figure 1.4). Of the 4.8 s uncertainty, 3.4 s is attributed to systematic effects in the measurement of the neutron flux. An improved version of this experiment is presently running at NIST, and they hope to make a measurement with an accuracy of less than 2 seconds[18].

Storage experiments have the potential to make a substantially improved measurement of the neutron lifetime. In a storage experiment, neutrons are loaded into a containment vessel and stored for a variable length of time before being counted (figure 1.5). The neutron lifetime can be extracted from the depen-



dence of the detected neutron population on the storage time. Several storage techniques have been used to make measurements of the neutron lifetime. The most precise storage measurements (yielding values of

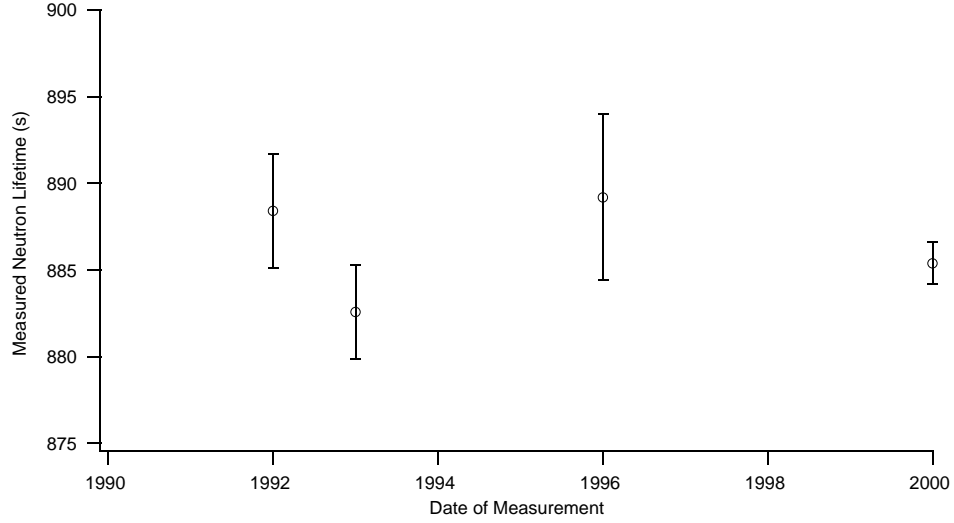


Figure 1.6: A Plot of the measured values of the neutron lifetime found in the experiments described in the text in references [20], [19], [17] and [21].

$\tau_n = 882.6 \pm 2.7$ s[19], 888.4 ± 3.3 s[20] and most recently $885.4 \pm 0.9 \pm 0.4$ s[21], see figure 1.6) have relied upon the total reflection of ultracold neutrons^{||} (UCN) from material surfaces to confine neutrons in a physical box.

An ultracold neutron with energy of 1 mK has a wavelength of 1000 Å, which is orders of magnitude larger than the interatomic spacing in a solid or liquid material. Although the interaction potential between a neutron and an atom is strongly peaked at the nuclei, the interaction of ultracold neutrons with materials can be pictured as a smeared out effective potential given by the scattering lengths of the various nuclei

^{||}UCN are traditionally defined as those with low enough kinetic energies that they can be totally internally reflected in some material. ^{58}Ni is the highest effective potential material, with a potential of 335 neV or 3.9 mK. However, Be, with a potential of 252 neV and Fomblin oil, a fluorinated oil, with a potential of 106.5 neV are more commonly used for UCN storage. A less material dependent definition would simply refer to neutrons with energies $\lesssim 1$ mK. For comparison, this energy, 86 neV, corresponds to a velocity of 4 m/s, a wavelength of 100 nm, or a rise against gravity of 85 cm or against a magnetic field change of 1.4 Tesla(appendix I).

and their density, with the potential forming a step function at the surface of the material. Thus, many materials can be used to totally reflect UCN below a certain energy at all angles of incidence, making it possible to store UCN in a material vessel.

Experiments using this storage technique to measure the neutron lifetime measure the total rate at which neutrons are lost from the storage vessel, both by beta decay and by any other loss mechanisms. For perfect storage, with no additional loss mechanisms, the total loss rate is identical to the neutron decay rate. For actual experiments, the accuracy of the lifetime measurement is limited either by unknown loss mechanisms or by the uncertainty in corrections for large loss mechanisms. For present experiments using material storage, the limiting loss mechanism is believed to arise from the interactions between the stored neutrons and the walls of the storage vessel. Most experiments have attempted to use storage data to extrapolate to the ideal case of no wall losses. In reference [19], the physical size of the storage vessel is varied, and the measurements are extrapolated to an infinite size. In reference [20], the storage time is measured as a function of UCN energy or equivalently velocity, and an extrapolation is made to zero velocity, or an infinite time between collisions with the walls. The most recent measurement, [21], used two different storage-vessel sizes with neutron detectors to measure the relative rate at which neutrons were scattered out of the storage vessel in the two configurations. UCN storage measurements using material bottles do not offer the prospect for significant improvement in τ_n .

Another storage technique uses the interaction of the magnetic moment of the neutron with magnetic field gradients. Material storage was first discussed in 1959[22]. Magnetic storage was discussed by Vladimírskii just two years later[23]. Since then, there have been several experiments to store neutrons magnetically[24, 25, 26], with varying degrees of success.




Fig. 3. Sextupole torus.

Figure 1.7: A sketch of the idea of the NESTOR experiment.

In the NESTOR (NEutron STorage Ring) experiment (figure 1.7), neutrons in a range of velocities from 6–20 m/s were stored in a sextupole magnetic storage ring analogous to the storage rings used at particle accelerators[27, 25]. Like the material surfaces in neutron beam guides (section 3.2.1), magnetic field gradients reflect neutrons at a glancing incidence, storing neutrons in certain trajectories. Unlike the material storage lifetime experiments and the experiment described in this work, however, the storage ring technique relies on the stored neutrons remaining in these trajectories, and losses from these trajectories (attributed to betatron oscillation) limited this measurement. Three-dimensional magnetic trapping was later attempted with a spherical hexapole trap[26], which also depended on neutrons to remain in certain trajectories, so as to avoid Majorana spin-flip losses of neutrons passing through the center of the trap (section 2.1). The lack of neutrons trapped in this experiment was attributed to scattering of neutrons from the trap by phonons in their helium

bath and helium atoms in their vacuum. Another experiment used a large box, with several coils along the floor and walls (the top was “closed” by gravity), but despite its size was never able to confine more than a few (of order ten) neutrons at a time, with a trap depth of order $150 \mu\text{K}$ [24, 28].

1.3.1 This Work

Another method of measuring the neutron lifetime was proposed in 1994 by Doyle and Lamoreaux in which UCN are produced and stored in a magnetic trap which has no zero-field region and is filled with superfluid helium-4[29]. In this experiment, the UCN are produced by the scattering of a cold neutron beam in superfluid helium (section 2.2) and are stored in a magnetic trap having no zero-field region (section 2.1). Unlike other storage experiments in which the number of neutrons remaining after a fixed amount of time are counted, the trapped UCN are detected storableously as they decay, using scintillations produced in the helium bath (section 2.3).

In this work, ultracold neutrons were produced, magnetically trapped, and detected as they decay, as proposed in reference [29]. The feasibility of using an improved version of this experiment to make a neutron lifetime measurement is discussed in section 5.2. In addition to the possibility of making an improved lifetime measurement, the techniques developed for this experiment have other potential applications, also discussed in section 5.2, such as in an improved measurement of the neutron’s electric dipole moment[30], detection of low energy neutrinos[31], or in other measurements using ultracold neutrons.

Chapter 2

Experimental Concepts

This thesis reports the production, magnetic trapping and detection of ultra-cold neutrons (UCN) as proposed in reference [29]. The UCN are produced using the scattering of cold neutrons in liquid helium[32] (section 2.2). Sufficiently cold UCN with the correct spin state are then trapped, using their interaction with magnetic fields (section 2.1). As the trapped neutrons decay, the decay electrons are detected through the scintillations they produce in the liquid helium bath (section 2.3). In this chapter, the main ideas of this method are described. The experimental realizations of these ideas are discussed in Chapter 3.

2.1 Magnetic Trapping

Neutrons have a magnetic moment, $\mu_n = -1.91\mu_N = -0.7\text{mK/T}$, and interact with a magnetic field, \vec{B} , by the dipole interaction:

$$H = -\vec{\mu} \cdot \vec{B} = \mu_n \vec{\sigma}_n \cdot \vec{B}, \quad (2.1)$$

where the magnetic moment of the neutron is oriented anti-parallel to the spin ($\vec{s}_n = \frac{\hbar}{2}\vec{\sigma}_n$) of the neutron. Since the neutron is a spin-1/2 particle, it has two possible quantum spin states, oriented parallel and anti-parallel to the magnetic field. When the neutron spin is aligned parallel to the magnetic field, the interaction energy is positive. Regions of high magnetic field repel neutrons in the parallel or “low-field-seeking” spin state. Similarly, neutrons with spins anti-parallel to the magnetic field have a negative interaction energy, and are attracted to regions of higher magnetic field and are called “high field seekers”.

The interaction potential given in equation 2.1 also causes the component of the magnetic moment perpendicular to the field to precess around the local direction of the field with a frequency $\omega = \frac{2\mu_n B}{\hbar}$ (or $\frac{\omega}{2\pi} \approx 29$ MHz at 1 Tesla). If this precession is considerably faster than the change in the direction of the field seen by the neutron, that is if

$$\frac{2\mu_n B}{\hbar} \gg \frac{d\vec{B}/dt}{|\vec{B}|}, \quad (2.2)$$

then $\vec{\sigma}_n \cdot \hat{B}$ will remain constant (i.e. its state will not change). This allows us to think of the magnetic potential as depending only on the magnitude, not the direction, of the magnetic field, $H = \mu_n B$.

It turns out that it is not possible to create a static magnetic field maximum in free space, but it is possible to create a static magnetic field minimum (appendix B.1). All static magnetic traps use such a minimum to confine low-field-seeking particles. In order to provide full three-dimensional confinement, the magnitude of the field must increase in every direction, giving a low-field “trapping region” completely surrounded by a higher magnetic field.

One of the simplest realizations of this uses two coaxial coils with opposite




Figure 2.1: A sketch of an anti-helmholtz trap.

current sense, as shown in figure 2.1. This “anti-Helmholtz” configuration of coils produces a field which increases in magnitude linearly in all directions from a zero-field point at the center. This zero-field point is the main drawback in using such a configuration as a magnetic trap. As a trapped particle passes near the center, where the field is very weak, equation 2.2 does not hold. Its spin thus precesses quite slowly, allowing $\vec{\mu} \cdot \vec{B}$ to flip sign, leaving the particle in the high-field-seeking or anti-trapped state. The high-field-seeking particle is quickly expelled from the low-field trapping region and lost from the trap. Such losses are known as Majorana spin-flips and can be minimized by using a field configuration without a zero-field region. Due to the geometry of the trap, no simple bias field can eliminate the zero-field “hole” in an anti-Helmholtz trap, but can only move it around.

The magnetic trap described in reference ([26]) is similar to an anti-Helmholtz trap, but uses three coils instead of two, as shown in figure 2.2. Both configurations produce traps with a spherically symmetric magnetic potential ($|B|$), but the magnetic




Figure 2.2: A sketch of a spherical hexapole trap.

potential rises quadratically with radius in the spherical hexapole. Like the anti-Helmholtz trap, the spherical hexapole has an unavoidable zero-field “hole” at its center, at which trapped particles can be lost.

A slightly more complicated trapping field has a magnetic field directed inward along one axis, outward along a second, and along the third axis, the field is directed in one direction – inward on one side of the trap and outward on the other. Such a field can be realized with four wires and two coils as shown in figure 2.3. This configuration has been used to confine ions and plasmas, and more recently to trap neutral atoms, in which context it is sometimes referred to as an Ioffe trap[33]. The four wires, with alternating current direction, provide a cylindrical quadrupole field. The two coils, coaxial with the same current sense, provide a non-zero field throughout the trapping region, suppressing Majorana losses. This minimum field can be further




Figure 2.3: A sketch of a Ioffe trap.

increased by the application of a simple bias field with the same orientation. Our magnet is of this form, without the additional bias field, and is described in section 3.3.

A magnetic trap can be characterized by its trap depth, defined to be the difference between the minimum magnetic field in the trapping region, and the lowest field at which a particle can escape from the trap. The trap depth thus corresponds to the maximum possible energy of a trapped particle. A typical trap depth of 1 T, yields a maximum energy for a trapped neutron of 60 neV, corresponding to a temperature of 0.7 mK. In practice, while a particle with energy greater than the trap depth *can* reach a point from which it can escape the trap, often trajectories exist on which it may take a long time to do so, resulting in the confinement of particles which are not actually trapped. Such marginally trapped particles are discussed further in section 2.5.2.

2.2 Superthermal Production of UCN

By definition, a particle which is magnetically trapped is surrounded by a magnetic potential greater than its total energy. Thus, to load a particle into a trap, the particle needs to either pass through the walls on entering the trap (because the trap is turned off, trap door loading, or the particle is in the untrapped spin state, RF loading), or lose energy in the trap after passing “over” the walls to enter. Many of the methods traditionally used to load atoms into traps are inappropriate for neutrons. Using light scattering (as [34]) directly to cool the neutrons is impossible, given the extremely low interaction cross section. Thermalizing the neutrons with a cold gas (as [35]) is difficult both because of the neutron’s relatively low cross section* (more than six orders of magnitude less than is typical for ions, atoms and molecules) and because the vapor pressure of anything except ^3He is negligible at 1 mK.[†] Cooling the neutrons by thermalizing them with a liquid or solid material would be difficult since the material will remain in the trapping region, continuing to scatter and absorb trapped neutrons. In principle, radio frequency fields can be used to flip the spin of the neutron from the high-field-seeking state (losing kinetic energy as it enters the trapping region) into the trapped low-field-seeking state in the trapping region[37]. Of course this still requires the neutrons to be cooled to near UCN energies elsewhere. Similarly, a pulsed neutron beam can be reflected from a rapidly moving surface (described in more detail below), however for loading a trap this can only be used for

*For example, one might imagine trying to thermalize a neutron beam with a sample of evaporatively cooled monatomic hydrogen (scattering cross section \gg absorption cross section, so the neutron can survive several collisions, and $m_H \approx m_n$, so it only takes ~ 10 collisions to thermalize from 30 K down to 1 mK) however, at typical densities ($\approx 4 \times 10^{13} \text{ cm}^{-3}$ [36]) a sample of hydrogen $4 \times 10^8 \text{ cm}$ in diameter would be required to thermalize a non-vanishing fraction of the neutrons.

[†]And ^3He has an extremely large absorption cross section for neutrons, three orders of magnitude greater than its scattering cross section, making it useless as a moderator.

a single pulse, as the moving surface will sweep neutrons back out of the trap with each pass. Likewise, a magnetic trap can be energized while a beam of UCN passes through the trapping region, confining a UCN density at most equal to that in the beam (the “trap door” method).

Ultracold neutrons have been produced in a number of ways in the last several decades. UCN were first studied using a standard cold source (like the one described in section 3.2.1). Among the many cold neutrons coming out of such a source, there will be a few UCN in the low energy tail of the approximately thermal (Maxwell-Boltzmann) distribution. The UCN can be separated out using a curved neutron guide — only neutrons which can be reflected at sharp angles (i.e. UCN) will go around a bend. This method in general produces low yields of UCN, but was used initially to confirm the prediction that neutrons could be reflected at all angles of incidence.

Later methods attained higher UCN fluences by shifting higher energy neutrons to UCN energies, such as the rotor method realized at the UCN Turbine using a cold neutron beam at the Institut Laue Langevin (ILL), in Grenoble, France[38] (or using a pulsed source [39]). A neutron reflector[‡] with a velocity half that of a neutron beam can scatter neutrons from the beam to extremely low energies. Seen in the rest frame of the reflector, the neutrons reflect off the wall with their velocity unchanged in magnitude, but opposite in direction. If the reflector’s velocity is tuned so that it is exactly half the velocity of the incident neutrons, the scattered neutrons will be left at rest. Thus neutrons from a higher energy section of the cold neutron spectrum can be turned into UCN.

Although higher fluences can be achieved with these methods, the UCN den-

[‡]either a Bragg reflecting crystal or a high potential material – ^{58}Ni at the ILL Turbine

sities achievable by this process are fundamentally limited by Liouville’s theorem, which states that the volume of phase space occupied by a group of particles will not change as the particles move reversibly under the influence of an arbitrary potential. Thus, the greatest UCN density that can be achieved by a mechanical source is equivalent to the highest phase space density of the cold neutron distribution. Further, since a Maxwell-Boltzmann distribution has constant phase space density in the low-energy tail, and a lower phase space density elsewhere, to the extent that the output of a cold source is actually Maxwellian, the phase space density of the UCN produced from the cold neutron beam by any process which can be described by a potential can be no greater than the phase space density of the UCN already in the tail of the distribution. Processes like the moderation of the neutrons in the reactor to produce a cold beam avoid this limitation by increasing the phase space volume of the particles in the moderator. The phase space volume occupied by the neutrons AND the atoms in the moderator (or equivalently, the phonons in the moderator) together is unchanged, as statistical mechanics requires.

Similarly, in order to achieve greater UCN densities, other than by starting with a higher flux cold beam, one needs a process in which the neutrons are coupled to another system whose phase space density can decrease as the phase space density of the UCN increases. The “superthermal” process (so called because it allows one to achieve phase space densities greater than that coming out of a “thermal” cold source), proposed in 1975 by Golub and Pendlebury[40, 32], does this by increasing the phase space volume occupied by phonon excitations in a material. The use of superfluid helium to produce UCN has been demonstrated in several experiments[41, 42, 43].

The superthermal production of UCN in superfluid helium depends critically on the fact that cold neutrons cannot scatter with individual helium atoms as they

would in a gas, but must scatter with excitations (phonons and rotons) in the superfluid. This is just a consequence of the fact that the wavefunction of the superfluid, unlike that of a gas, cannot be separated into independent wavefunctions associated with each atom. Since the ^4He atoms are bosons, the total wavefunction will be symmetric under the exchange of two atoms. In the “superfluid component”, the momenta of the atoms in a region are identical, giving a symmetric wavefunction $\Phi_0 \propto \exp(i\vec{K} \cdot \sum_j \vec{r}_j)$. The straightforward modification for scattering with one atom would produce an explicitly unsymmetric wavefunction, like $\Phi_1 \propto \exp(i\vec{q} \cdot \vec{r}_i)\Phi_0$. In order to symmetrize the wavefunction, the momentum transfer must be shared among the helium atoms, as in $\Phi_2 \propto \sum_j i\vec{q} \cdot \vec{r}_j \Phi_0$.

At first glance, this seems to be equivalent to the scattering of a single atom, at least in its effect on the scattered neutron. However, when the interactions between the helium atoms are taken into account, the energy of the collective excitation represented by Φ_2 is very different than the energy change in a free atom with a change in momentum \vec{q} . The dispersion curve (energy versus momentum) for excitations in superfluid helium is shown in figure 2.4 with the energy of the neutron ($\frac{p^2}{2m}$) overlayed for comparison. For the dominant “one phonon” process, momentum and energy conservation require that a scatter to the near vicinity of $q = 0$ begin either already near $q = 0$ or else near the crossing point of the curves at $q = q^*$, by creating a single 11 K phonon. Similarly, a UCN, rather than being able to scatter out of a 1 mK deep trap by absorbing energy from a 1 mK phonon, can only “upscatter” by absorbing an 11 K phonon (ignoring multi-phonon processes for the moment) (appendix B.2). This fact allows us to consider the neutron interaction with the superfluid as a two level system with an 11 K energy difference between the two levels.

More generally, consider the interaction of an approximately thermal source of




Figure 2.4: The dispersion curves cross at $q = 0$ and at $q = q^*$, which corresponds to a neutron wavelength of 0.89 nm, or an energy corresponding to 11 K.

neutrons, at temperature T_0 , with a colder target which can be thought of as having only two states, the ground state and an excited state with an energy Δ . This target will be able to couple neutrons from a narrow band of energies $\Delta - \Delta + 2\sqrt{\Delta E_{UCN}}$ into UCN with energies up to E_{UCN} . We can imagine in principle a large sample of neutrons coming into thermal equilibrium with the moderator. By definition, in equilibrium the scattering of cold neutrons to UCN will be exactly balanced by the scattering of UCN back to higher energies. Given that the neutrons will have a perfectly thermal (i.e. Maxwellian) distribution, we can write this “detailed balance”

explicitly:

$$\begin{aligned} & \frac{E_{UCN}}{k_B T} \exp\left(-\left(\frac{E_{UCN}}{k_B T}\right)\right) \sigma(E_{UCN} \rightarrow E_{UCN} + \Delta) = \\ & \frac{E_{UCN} + \Delta}{k_B T} \exp\left(-\left(\frac{E_{UCN} + \Delta}{k_B T}\right)\right) \sigma(E_{UCN} + \Delta \rightarrow E_{UCN}) \end{aligned} \quad (2.3)$$

which we can write as an expression for the upscattering rate:

$$\sigma(E_{UCN} \rightarrow E_{UCN} + \Delta) = \frac{E_{UCN} + \Delta}{E_{UCN}} \exp\left(-\frac{\Delta}{k_B T}\right) \sigma(E_{UCN} + \Delta \rightarrow E_{UCN}) \quad (2.4)$$

where k_B is the Boltzmann constant. For a fixed Δ (11 K in our case), the upscattering rate can be made arbitrarily small, decreasing exponentially with the temperature of the helium bath. Physically, this corresponds to decreasing the population of phonons of energy Δ (11 K).

In practice, however, as the rate of single-phonon upscattering is reduced, the rates of other (multi-phonon) scattering processes begin to dominate the upscattering rate. The second-order process is usually called two-phonon upscattering, although it applies equally to the scattering of a neutron by a single phonon, in which the phonon is not absorbed.

The rate of the two-phonon upscattering scales as $\Gamma \propto T^7$ [44][§] and has been measured down to 500 mK [41]. In a helium bath at a temperature of 1.16 K, the upscattering rate was about $2.5 \times 10^{-2} \text{ s}^{-1}$, corresponding to a storage time of about 40 s. Extrapolating to lower temperatures, at our operating temperature of 250 mK, the upscattering rate is 2000 times less than the beta-decay rate. So, in order to see trapped neutrons clearly we need to cool the helium bath well below 1 K, requiring a dilution refrigerator (section 3.4).

[§]The T^7 dependence is calculated explicitly in reference [44]. A simple explanation is that integrating the phase space of both the incoming and outgoing phonons (k^3 each) over the Maxwell-Boltzmann distribution ($\exp -k/T$) gives $\int dk k^6 \exp -k/T \propto T^7$.




Figure 2.5: The beta decay spectrum of the neutron.

2.3 Detection of decays in the trap

A free neutron decays into a proton, an electron and an electron anti-neutrino. Most of the energy of the reaction (782 keV) is released as kinetic energy of the electron and the neutrino. If there were only an electron and a proton given off, there would be exactly one electron energy (determined by energy and momentum conservation). Because the electrons share their energy with a neutrino, however, they have a range of energies (figure 2.5, equation A.19), with a peak kinetic energy of 250 keV and a maximum kinetic energy of 782 keV.

When a neutron decays in the trap, the decay electron scatters in the helium. The interaction of high energy particles (especially electrons and alphas) with helium has been well studied[45]. As the electron passes through the helium, it scatters

from atomic electrons, ionizing tens of thousands of helium atoms. The electron deposits about 30 keV/mm [45], so the range of a typical electron is no more than one centimeter.

The kinetics of the low energy electrons and the helium ions is complicated by the presence of the helium. The electrons repel the helium atoms, creating a tiny “bubble”, and the positive ions gather a cloud of polarized helium atoms in a tiny “snowball”. The ions tend to form diatomic molecular ions before recombining with electrons, leading to the creation of excited-state diatomic helium molecules. These molecules are created in both singlet and triplet states. While we observe the singlets to radiatively decay in less than 10 ns, the triplets have a radiative lifetime of about 13 s [46]. When each helium molecule decays, it emits a photon in the extreme ultraviolet ($\lambda \approx 70 - 90$ nm). The singlet decays give a relatively intense pulse of light (15 photons per keV* of beta energy, or 4000 photons for a typical neutron decay) [45] in a short time (less than 20 ns), and thus can be used to detect the neutron decays. The longer decay time for the triplets renders them less useful for this purpose. The same number of photons (4000) coming out over the course of even one second would be hard to distinguish from reasonable backgrounds (especially after allowing for detection efficiencies).[†]

Each neutron decay in the trap results on average in a pulse of a few thousand EUV photons. To detect this light directly in the trap would require detectors which work at low temperatures (250 mK) where the electronic behavior of most

*Reference [45] quotes 87 keV of EUV photons (around 16 eV/photon) for a 364 keV beta.

[†]Immediately following the beta-decay event, the triplet molecules are in a fairly small region, and they decay rather quickly by scattering, before diffusing apart. In principle, photons from these decays, in the first microsecond or so after the decay, could be used in coincidence with the initial pulse of singlet decays, to distinguish events originating in the liquid helium from those originating in the acrylic lightguide. However, so far various technical issues (high background, and PMT afterpulsing) have prevented this from working.[47]

materials is different from that at room temperature. The detectors would also need to work inside an intense magnetic field (1 Tesla) without substantially perturbing the field, making conducting materials and electric currents in general problematic. The detectors would also be exposed to an intense cold neutron beam, or at least to neutrons scattering in the helium, and would need to tolerate this radiation without creating backgrounds. Finally, there is limited space available inside the magnet bore. The thickness of any material around the trapping region directly reduces the size and depth of the trap. It would be extremely difficult to satisfy these detector requirements. Instead, we choose to detect the scintillation light indirectly.

Transporting EUV light is quite difficult. It can be passed by windows of order 100 nm thick[48], far too thin to contain the helium bath at 1 bar, and no material reflects it efficiently. Transmitting ordinary blue or violet light is easier. Standard mirrors, windows and lightpipes can be used, and with the exception of mirrors, can be made purely from elements which don't cause background problems (carbon, hydrogen and oxygen). Converting the EUV light to blue/violet light can also be accomplished with a pure hydrocarbon, 1-1-4-4 tetraphenyl 1-3 butadiene (TPB)(figure 2.6). TPB absorbs EUV light and emits light in a spectrum from about 400 nm to 500 nm [49], corresponding to the blue and violet regions of the spectrum. The detection scheme is described in section 3.9. In brief, the EUV first encounters a thin layer of TPB-doped polystyrene, on the inner surface of an acrylic tube. The TPB downconverts the EUV into blue light, some of which is trapped in the wall of the tube and piped to the end of the trapping region. A series of acrylic lightguides and windows transports the light out of the cryostat, where it is detected using a pair of photomultiplier tubes (PMTs).




Figure 2.6: The molecular structure of TPB.

2.4 Experimental Method

Each experimental run consists of a loading phase in which the neutrons are loaded into the trap and an observation phase in which the trapped neutrons are then detected as they decay. At the start of a run, the magnetic field is turned on, forming the magnetic trap which will confine the neutrons. The neutron beam is then turned on, and neutrons begin to scatter in the liquid helium, producing UCN in the trap. The UCN creation rate is proportional to the incident neutron fluence, and constant during the loading phase. UCN are lost from the trap, mainly by beta decay, in proportion to the number trapped. Thus, neutrons accumulate in the trap, with a time constant of about 15 minutes (appendix B.3). In the actual experiment, we loaded the trap for 22.5 minutes, which yielded 78% of the trapped UCN population which could be achieved with an infinite loading time. Once the beam is off, the number of neutrons begins to decrease as neutrons decay. The PMTs are turned on, and pulses of light are recorded as a function of time for one hour. Trapped neutrons would be indicated by an exponentially decaying signal consistent with the neutron beta-decay lifetime, or decaying more quickly if there are additional trap losses. Two classes of problems could prevent us from observing this signal: spurious losses (which

cause the initial trapped population to be significantly smaller due to losses during loading, as well as causing the signal to fall off more quickly) and background signals (which could give a much higher overall counting rate).

2.5 Loss Mechanisms

Possible loss mechanisms include thermal upscattering (section 2.2) which is reduced by cooling the helium target; Majorana spin-flips (section 2.1) which are reduced by the imposition of a non-zero field everywhere in the trap; and ^3He absorption (section 2.5.1). In addition, some neutrons with energies greater than the trap depth (that is, sufficient to allow them to escape) may be confined by the field for an extended period of time but “lost” from the trap other than by decay (“marginally trapped neutrons”, section 2.5.2).

2.5.1 ^3He Absorption Losses

Helium has two stable isotopes, ^3He and ^4He . ^3He makes up $1 - 2 \times 10^{-7}$ of helium gas obtained from natural gas sources (1.3×10^{-6} of atmospheric helium[50]), but has an extremely high absorption cross-section for neutrons. For UCN contained within a bath of helium including ^3He , the absorption rate is given by:

$$\Gamma_{abs} = x_3 \cdot 2.6 \times 10^7 \text{s}^{-1} \quad (2.5)$$

where x_3 is the ratio of ^3He atoms to ^4He atoms (see appendix G.3). Isotopically purified ^4He is available[‡] with $x_3 < 5 \times 10^{-13}$ [51]. The actual purity is probably better than this; the limit quoted in reference [51] was the threshold of the measuring

[‡]We purchased isotopically purified ^4He from Professor P.V.E. McClintock.

technique used. At the quoted limit of 5×10^{-13} , we would expect an absorption rate of $1.3 \times 10^{-5} \text{ s}^{-1}$, or a hundred times less than the beta-decay rate. This gives roughly a 1% shift in the trap lifetime, which is insignificant for observing the presence of trapped UCN. Reversing this, the data presented in section 4.4 puts a 3σ limit on the ^3He concentration in our helium bath of $x_3 \leq 5 \times 10^{-11}$.

2.5.2 Marginal Trapping

Though unlikely to present a problem at the present level of precision, marginally trapped neutrons could also obscure our signal. A “marginally trapped” neutron is one with an energy greater than the trap depth, but in a trajectory which remains confined in the trapping region for a significant time. The maximum energy of a marginally trapped neutron is about $2.5\times$ the trap depth (appendix G.1); the maximum number of UCN produced with such energies is four times that of the UCN produced at trappable energies. Considering the specific trajectories required, the actual ratio should be much smaller. Simulations have determined that the vast majority of these neutrons will be expelled from the trap within a few seconds[§][52, 53]. In addition, it has been proposed that by lowering the magnetic field and then raising it back up, the maximum energy of the neutrons in the trap can be reduced, eliminating the possibility of marginal trapping (discussed in more detail in appendix G.1). The final trapped sample achieved by ramping to 30% of the trap depth has no neutrons stored with energies above the trap depth, though about half of the trapped neutrons would also be thrown away in such a procedure. For simply detecting the

[§]At UCN velocities (less than 4 m/s), the 30 cm length of the trap is about a 100 ms trip, and the 5 cm width of the trap is about a 10 ms trip.

trapped neutrons, such a procedure is clearly unnecessary.[¶]

2.6 Backgrounds

Background counts are any pulses of light in the photomultiplier tubes which arise from sources other than decaying neutrons. Such a pulse could arise from the passage of a high-energy particle through the helium or through some other part of the detection system. High-energy particles passing through the apparatus could be cosmic rays; decay products from radioactive nuclei, either naturally occurring or created by neutron capture during the loading phase; or particles created by the scattering of neutral particles (gammas or fast neutrons) during the observation phase. Below, we discuss various backgrounds in general, with more detailed descriptions or quantitative analysis provided in the indicated sections.

2.6.1 Cosmic Rays

Cosmic ray muons pass through every square centimeter of the Earth's surface at a rate of $1 \text{ cm}^{-2} \text{ min}^{-1}$ [3], and penetrate any reasonable shielding. They leave an ionization trail in whatever they pass through, depositing about $2 \text{ MeV g}^{-1} \text{ cm}^2$ times the density of the material [3], so conventional scintillation detectors will easily detect these muons. A detector, placed either above or below the detection cell, can record each muon event seen in our detector. Such events can thus be easily vetoed, virtually eliminating these background events. (The muon detectors are described in section 3.10, and the vetoing in section 3.11.)

[¶]Other systematic effects have been suggested, such as a possible shift in the beta-decay lifetime of the neutron due to the presence of the helium. This and other effects, important for a lifetime measurement, but largely irrelevant to this work are discussed in section 5.1.1.

2.6.2 Gamma Rays

Backgrounds from the scattering of neutral particles (primarily gamma rays) cannot be suppressed in this way. The gamma radiation in the guide hall comes both from long-lived naturally occurring isotopes (like ^{238}U , ^{232}Th , and ^{40}K all with half-lives longer than 10^9 years) and from neutron absorption in experiments at the neutron beams there. Unlike muons, which scatter electrons nearly continuously, gammas tend to lose a significant fraction of their energy in each scattering event. Most of the gammas scattering in the cell or lightpipe will be completely absorbed before leaving the cryostat. On the other hand any detector with a sufficient density to detect most of the gammas entering it (such as a lead glass) would also serve as effective shielding. To reduce backgrounds from gamma rays, we surrounded the apparatus with 4" of lead shielding (section 3.10). Of course, backgrounds from long-lived isotopes in the materials of the apparatus itself will not be blocked by this shielding. These backgrounds can only be reduced by careful materials selection. Such long-lived isotopes will give us a flat background (i.e. constant in time on a scale of several hours).

2.6.3 Neutron Activation

Backgrounds are possible from isotopes with shorter lifetimes, primarily created by neutron absorption in materials. Such a signal can arise either directly, from charged decay products scattering in the helium or acrylic, or through gamma rays. In either case, such backgrounds will not be flat, rather decaying exponentially with their individual lifetimes. This makes them a potentially more serious problem. In order to reduce such backgrounds, the materials inside the neutron shielding sur-

rounding the detection cell were carefully chosen such that their primary constituents have small neutron absorption (or absorb neutrons to form a stable isotope, thus giving no background). These materials were then tested at the reactor using neutron activation analysis to see that in practice they did not produce any significant activation signals (appendix E). Additionally, the external portions of the apparatus were shielded to reduce, as much as possible, the scattering of neutrons into other materials (section 3.7).

2.6.4 Luminescence

When we first put neutrons into the apparatus, we observed an additional source of background which was not expected. The materials used to shield the rest of the apparatus from the neutrons (primarily hexagonal boron nitride, hBN) glow after exposure to neutrons, with a magnitude decreasing slowly on a time scale comparable to an experimental run (section 4.3.1 for more details and also reference [54, 55]). A number of methods were used to decrease this background (which was, for some time, the major obstacle to extracting our signal). These included inserting a black material (graphite) between the boron nitride and the light collection system to block the luminescence light (section 3.7), testing other neutron absorbers for their luminescence properties[55], and requiring of coincident detection in two PMTs, looking at the end of the same lightpipe. A signal was recorded when at least two photons (at least one to trigger each PMT) were detected at the same time (section 3.11.1). This, combined with a threshold on each PMT (requiring a higher signal than the typical response to a single photon) has the effect of reducing the luminescence signal by many orders of magnitude (section 4.3.1), while reducing our signal from

neutron decays to 31% of the actual number of neutron decays in the trap (section 3.9).

Despite several successful methods of reducing our backgrounds, we still had to extract a decaying signal with peak magnitude of roughly 0.1 s^{-1} from a background signal of about 6 s^{-1} initially, with a flat rate of 2 s^{-1} at longer times. We did this by measuring our background directly, taking half of our runs with the magnetic trap turned off, and subtracting this background from the trapping runs (Chapter 4).

Chapter 3

Experimental Apparatus

3.1 Overview

To magnetically trap ultracold neutrons (UCN) and detect their decays in real time, we require a source of neutrons, a magnetic trap, a bath of 250 mK helium inside the trap, and a way to detect the scintillation light produced when the trapped neutrons decay. The apparatus is described in detail in this chapter. This overview section provides a very short summary. Our neutron source is a cold neutron beam (section 3.2) from the NIST Center for Neutron Research (NCNR). The magnet is a Ioffe configuration of four racetrack-shaped coils and two solenoids (section 3.3). This produces a sausage-shaped trapping region about a foot long and an inch in diameter, with a trap depth of approximately 1 T. Inside the magnet bore is a cupronickel tube (section 3.7) which contains the helium bath and the detector insert (section 3.9). The detector insert is a thick-walled acrylic tube with a waveshifting dye doped on the inner surface, which absorbs the scintillation light produced when a neutron decays and converts it into blue light. Some of the blue light is piped down the length

of the tube through total internal reflection to a solid acrylic lightpipe. The light travels down the lightpipe and out of the cryostat through a series of windows to a pair of photomultiplier tubes (PMTs) which detect the pulse of light created by each neutron decaying in the trap (section 3.9). The signals from the PMTs are amplified and passed to a set of NIM logic modules which trigger a digital oscilloscope to record the PMT pulses (section 3.11). The data is read from the oscilloscope in groups of 100 events, and analyzed off-line (Chapter 4).

3.2 Beam

3.2.1 Cold Neutron Guide NG-6

The experiment resides at the end of cold neutron guide number six (NG-6) in the guide hall of the NCNR. The neutrons in this beam are released from fissioned uranium nuclei in the core of the NIST Research Reactor. Neutrons are moderated to lower energies by D₂O surrounding the reactor core and then by a liquid hydrogen cold source. The cold source is a hollow sphere with walls filled with liquid hydrogen at a temperature of 20 K (figure 3.1). Neutrons partially thermalize with the hydrogen after scattering a few times, then escape through a hole in the sphere, exiting towards the neutron guides.

Some fraction of these neutrons pass into the open end of NG-6, a rectangular guide 60 m long, 15 cm tall and 6 cm wide. The inner surfaces of the guide are coated with a thin layer of ⁵⁸Ni, which has an effective neutron potential of 335 neV. Neutrons with low enough momentum perpendicular to the surface interact with the point-like potentials of the nickel nuclei which are blurred together into a single




Figure 3.1: NIST Center for Neutron Research Hydrogen Cold Source[1]

average potential. This potential corresponds to a maximum perpendicular velocity of 8 m/s. Neutrons with a perpendicular velocity below this limit will be totally reflected from the surface, and can thus be guided to an experimental station far from the reactor core, thereby reducing backgrounds. For a fixed neutron wavelength (or energy), the guide's surfaces are reflecting up to a critical angle of $\approx 2 \text{ mrad}/\text{\AA}$ (or more precisely, $\theta_c = \sin^{-1}(\lambda/\lambda_c)$, where $\lambda_c \approx 494 \text{ \AA}$).

3.2.2 Beamline

The neutrons reflect down the guide and exit at the end of NG-6 (figure 3.2). Partway along the length of the guide, the neutrons encounter a thin graphite crystal which Bragg-reflects a monochromatic slice (whose wavelength can be tuned between




Figure 3.2: NG-6 endstation.

2.3 Å and 6.1 Å by rotating the crystal) out of the beam and to a separate instrument. At the endstation the beam can either be blocked using three layers of borated aluminum attached to a 5 cm thick steel shutter or, when this shutter is raised, allowed to pass through a 6 cm diameter opening in a piece of ^6Li -loaded-glass.*

Neutrons passing through the 6 cm defining beam aperture (C_1), pass next through a filter cryostat. This aluminum vessel holds up to seven blocks of neutron wavelength filtering materials which are 7 cm square (transverse to the beamline) by 5 cm long (along the beamline). The cryostat is filled with liquid nitrogen to cool the blocks, reducing the loss of neutrons by phonon scattering in the filter blocks. During our experiment, the filter cryostat contained four blocks, two each of bismuth and

* ^6Li has a thermal neutron absorption cross-section of 941 barns. The “lithium glass” was made at the NIST glass shop, with an isotopic enrichment of about 95%, and the glass 16% LiO_2 by weight.

beryllium. The bismuth absorbs gamma rays from the reactor core or other upstream sources which cause additional backgrounds in the apparatus. The bismuth attenuates the γ flux by roughly two orders of magnitude[†] while reducing the 8.9 Å neutron flux by less than 20%. The polycrystalline blocks of beryllium Bragg-reflect neutrons with wavelengths less than 3.95 Å, twice the lattice spacing in beryllium.[‡]

The neutron guide itself, the guide exit and the filter cryostat are surrounded by steel shielding containers filled with a polyethylene and steel shot mixture to minimize background radiation. The beam exits this region through exit holes in a series of alternating steel and polyethylene plates. The combination of steel (high-Z) with polyethylene (lots of hydrogen, which absorbs neutrons) attenuates fast (MeV) neutrons. The high-Z material has a large cross-section scattering the fast neutrons, decreasing their energy slightly, but more importantly increasing the time they spend in the hydrogenous material. Neutrons lose much more energy per scattering event in the hydrogenous material and the hydrogen nuclei eventually capture the neutrons. The neutron beam exits the shielding region through a second ⁶Li-loaded-glass aperture with a diameter of 5.2 cm and passes through a hole in a 10 cm thick wall of lead bricks which blocks gammas scattered from the beam shutter.

When the beam is blocked by the shutter, the γ 's passing through the beam hole are blocked by a stack of lead bricks, moved in and out of the beamline by a translation table, the “slider”. While the neutron beam is on, the lead bricks are translated beside the beam, and the neutrons pass through a flight tube lined with

[†]Attenuation length of ≈ 2 cm for γ 's from 1-10 MeV, gives a transmission of ≈ 0.007 .

[‡]A single crystal of beryllium with lattice planes perpendicular to the beam axis would Bragg reflect neutrons at a wavelength of $\lambda_{\text{Bragg}} = 3.95$ Å. A crystal at an arbitrary angle scatters neutrons of a shorter wavelength $\lambda_{\text{Bragg}} \sin \theta$ where θ is the angle between the neutron momentum and a vector normal to the lattice planes of the crystal. A polycrystalline block is made up of many small crystals, so a neutron with a wavelength up to λ_{Bragg} is likely to see a crystal face at its Bragg angle, and thus scatter out of the beam.

boroflex[§] which absorbs neutrons scattered in the air (about 1% scatter per 10 cm of air). When the neutron beam is blocked by the shutter, the lead bricks are moved into the beamline. (The front face of the pile of lead is covered with lithium plastic[¶] as a safety measure to prevent neutrons from scattering from the lead in case the slider is accidentally closed with the neutron beam on.) The lead bricks block the thin beam of gamma rays which was observed to come from the beam guide, and which would otherwise result in a high background for this experiment.

After the slider, the neutrons pass through a hole in a second steel shielding wall, and finally into the trapping region through a series of Teflon^{||} windows[56]. At the front of the trapping region, a final aperture (C_2), made of boron carbide^{**}, collimates the beam. Unscattered neutrons passing through the trapping region are absorbed by a boron carbide beamstop.

3.2.3 Collimation

The collimators located at the end of NG-6 (C_1) and at the front of the trapping region (C_2) define the neutron beam. The collimators were chosen to be as large as possible such that any neutron passing through both C_1 and C_2 would hit the beamstop, unless it scattered in the helium. For a fixed distance between the two

[§]Boroflex is a black rubber with ^{10}B impregnated into it, purchased from Reactor Experiments, inc. www.nnc-usa.com, and used to absorb cold neutrons. ^{10}B has a thermal neutron absorption cross-section of 3840 barns.

[¶]Lithium plastic is a whitish rubber with 95% enriched ^6LiF impregnated into it, available in sheets 2 mm thick. More expensive and less flexible than boroflex, lithium plastic is superior for many applications because ^6Li doesn't emit any gamma rays when it absorbs a neutron. Lithium Plastic is produced at PNPI (The Petersburg Nuclear Physics Institute outside St. Petersburg, Russia) and distributed in the U.S. by Dora Texas, Inc.

^{||}Teflon is the trade name for a series of fluoropolymers made by DuPont.

^{**}The boron carbide aperture and beam dump were both purchased from (and machined by) Eagle Pitcher Technologies. We used boron carbide with the natural isotopic ratio of 19.9% ^{10}B to 80.1% ^{11}B .




Figure 3.3: A sketch of the collimation scheme, showing the constraints on the scrapers in the first diagram and the constraints on the collimator diameters in the second.

collimators (L_1) and between the final collimator and the beamstop (L_2), and fixed diameters of the first collimator (D_1) and beamstop (B), this requirement translates to a maximum diameter of the second collimator (D_2) given by: $D_2 = \frac{L_1 B - L_2 D_1}{L_1 + L_2}$. Between the two collimators, a series of scrapers were placed to absorb those neutrons which would not pass through C_2 . Such scrapers require less neutron-absorbing material than either a continuous tube or a large plate around C_2 . A sketch of the collimation scheme is shown in figure 3.3.

3.2.4 Beam Flux and Spectrum

From a peak flux of $4 \times 10^{14} \text{ n cm}^{-2} \text{ s}^{-1}$ in the core of the reactor[1], the neutron beam exiting the filter cryostat has a capture flux^{††} of $9 \times 10^8 \text{ n cm}^{-2} \text{ s}^{-1}$ [57]. In the cold source, the neutrons have an approximately thermal distribution at around 30 K. The walls of the guide reflect neutrons with angles up to $\theta_{max} \approx \lambda/\lambda_c$, reflecting more neutrons at longer wavelengths. The distribution at the end of guide is thus given approximately by λ^2 times the Maxwell-Boltzmann distribution. Passage through the beryllium filter removes most of the neutrons below 4 Å, leaving a spectrum very roughly given by equation C.5. A measurement of the spectrum is shown in figure C.3. A more detailed discussion of the flux at NG-6 can be found in appendix C.

3.3 Magnet

The ultracold neutrons are confined in a Ioffe-type superconducting magnetic trap[33]. In theory, the configuration of wires to create a Ioffe trap consists of four infinite wires and two coils (figure 2.3). In practice, the “infinite” wires must be joined at the ends, becoming elongated (racetrack-shaped) coils. If each coil were wound with a single turn, we would need to supply a DC current of over 10^5 A . Both to reduce this current, and for ease of manufacture, magnet coils are typically wound with many turns of narrower wire. In addition, a superconducting magnet needs to be kept cold (below 9.8 K for $\text{Nb}_{0.6}\text{Ti}_{0.4}$), which generally requires immersion in a bath of liquid helium. While the magnet itself generates no heat, the leads connecting

^{††} “capture flux” is the flux weighted by $v_{300 \text{ K}}/v$ or $\lambda/\lambda_{300 \text{ K}}$, where $v_{300 \text{ K}}$ and $\lambda_{300 \text{ K}}$ are the velocity and wavelength, respectively, of a neutron with an energy of $300 \text{ K} \times k_B$ (Boltzmann’s constant). This is the relevant number when calculating neutron activation in a “thin” material, since the chance of a neutron being absorbed varies inversely as the velocity of the neutron.




Figure 3.4: Two configurations of racetrack coils giving a cylindrical quadrupole field.

it to the power supply cannot (with present technology) be superconducting all the way to room temperature. Since a low electrical resistance implies a low thermal resistance, the design of the current leads passing from room temperature to the helium becomes a tradeoff between heating ($P = I^2R$) in the leads themselves and thermal conduction from room temperature. The obvious way to reduce the current that needs to be supplied is to decrease the diameter of the wire and wrap it in proportionately more turns, requiring less current to provide the same magnetic field.

A pair of racetrack-shaped coils can be used to provide the cylindrical quadrupole component of a Ioffe trap (figure 3.4a). To increase the depth of the trap, our magnet was designed using four racetrack coils, in two different sizes (figure 3.4b). The design of a racetrack coil is shown in figure 3.5. The coil is straight over most of its length, curving around to close the loop at each end. It is made from a single long strand of superconducting wire* wrapped around hundreds of times with the

*We used $890\mu\text{m}$ diameter kapton-wrapped 1.2:1 NbTi:Cu wire, which we purchased from Oxford Instruments.




Figure 3.5: The design of a racetrack coil.

space between the wires filled with epoxy, to hold the wires in place. The larger coils, with 777 turns, are 69.1 cm long by 9.2 cm high, with a coil cross-section ($T \times W$) of 3.2 cm by 2.0 cm. The wire in the smaller coils is wrapped around 336 times in an area of 1.5 cm by 2.0 cm. These coils are 65.2 cm long and 5.3 cm high.

When the magnet is energized, the coils exert a significant force on each other. To keep the coils from flying apart, they need to be held on a form of sufficient strength to counteract this force. In addition, the coils need to be under tension to prevent movement, which would cause them to quench. A cross-sectional view of the magnet form is shown in figure 3.6. The core section was machined out of a single piece of titanium[†]. The coils were placed onto this form with kapton and G-10 spacers.

[†]6-4 ELI grade titanium, an alloy with 6% aluminum and 4% vanadium




Figure 3.6: Cross-sectional view of the magnet form.

Aluminum plates, machined to give a cylindrical outer surface, were bolted onto the form. This cylindrical quadrupole assembly was wrapped in Kevlar^{‡§}, to uniformly compress the coils. The Kevlar was wound in 20 layers with a tension of 15 lbs., so that as the magnet is cooled and shrinks relative to the Kevlar, the Kevlar remains under tension. Though Kevlar has an extremely high tensile strength, it is vulnerable to abrasion. To protect the Kevlar, the magnet was first wrapped in a layer of Tedlar[¶] then the Kevlar, and then an outer layer of fiberglass. The fiber wrappings were embedded in epoxy to keep them from slipping. The quadrupole assembly had an outer diameter of 14.5 cm and an inner bore of 5.08 cm. The ends of the titanium form were machined into indium-sealing flanges, to mate to the inner vacuum can (section 3.5.2).

Denoting the long axis of the racetrack coils as the z-axis, with the origin at

[‡]DuPont's Kevlar para-aramid fibre, a high-tensile strength polymer fiber

[§]Thanks to Dr. Paul Pernambuco-Wise and Dennis Markiewicz of the National High Magnetic Field Laboratory in Tallahassee, Florida, for wrapping our quadrupole magnet.

[¶]Tedlar is a DuPont polyvinyl fluoride fiber.

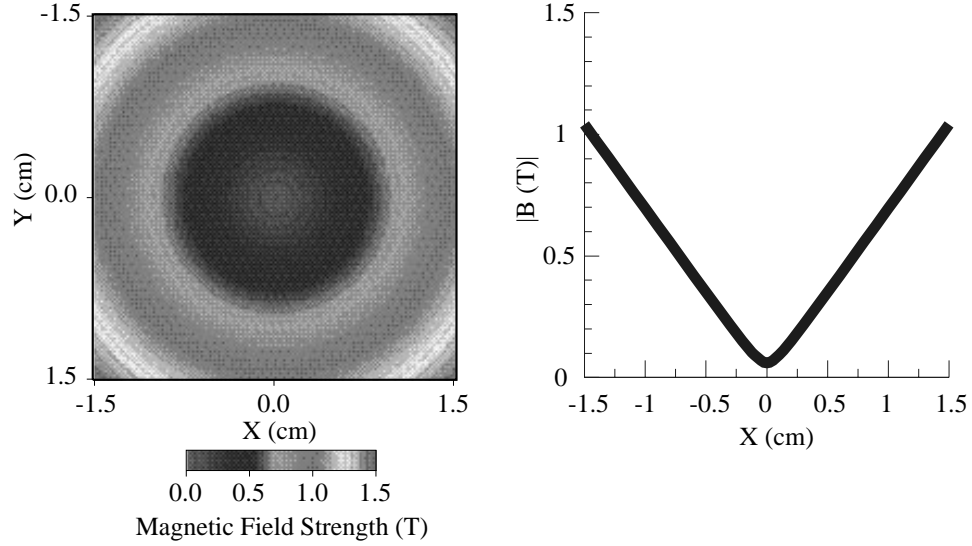


Figure 3.7: The field profile in the x-y plane at $z = 0$ (left), and plotted along the x-axis ($y = z = 0$) (right).

the geometric center of the trapping region, the field in the x-y plane is shown at $z = 0$ in figure 3.7. The magnitude of the magnetic field increases linearly with radius, with a gradient of ≈ 0.69 T/cm (at 180 A). Along the axis, the field minimum (0.11 T at 180 A) is provided by the solenoids. This non-zero minimum is crucial to minimize the loss of trapped neutrons by Majorana transitions (section 2.1).

The solenoids surrounding this quadrupole magnet were wound on forms made of 6061 aluminum. Each solenoid was flanked by a pair of smaller solenoids with opposite polarity. These “bucking” coils caused the magnetic field of the solenoids to fall off more quickly with distance along their axis. The solenoids each create a short high-field region closing off the ends of the trapping region and completing the trap (figure 3.8).

The depth of a magnetic trap is defined as the difference between the lowest




Figure 3.8: The field profile in the x-z plane at $y = 0$ (bottom), and plotted along the z-axis ($x = y = 0$) (top), showing the “pinch” resulting from each solenoid, closing the ends of the trap.

field point at which particles can be lost from the trap (B_{max}) and the lowest field point in the trap (B_{min}). In our trap, the solenoids were chosen to produce a peak field higher than the walls of the quadrupole. The trap depth was thus set by the quadrupole field at the inner edge of the detector insert, a radius of 1.59 cm corresponding to a trapping field of 1.15 T. Subtracting the minimum field of 0.11 T gives a net trap depth $B_T = 1.04$ T.

We provided current to our magnet using a Hewlett-Packard model 6681A DC power supply. The connection to the magnet was made through two long 3/0 AWG gauge welding cables, and a pair of vapor-cooled current leads with superconducting bus bars.^{||} The bus bars use superconducting wire soldered between two copper bars, allowing the current to be carried by the superconductor (without heating) for much of the length. The top sections of the leads are surrounded by hollow stainless steel tubes through which the helium boiloff can be directed. The cold helium gas cools the leads as it passes out of the dewar, greatly decreasing the heat

^{||}Model L-250 from American Magnetics, Oak Ridge, Tenn. www.americanmagnetics.com




Figure 3.9: A sketch of a dilution refrigerator.

which must be removed by the boiling of the liquid helium.

3.4 Refrigerator

Cooling the helium target to 250 mK to prevent upscattering of the trapped UCN (section 2.2) requires the use of a ^3He - ^4He dilution refrigerator (figure 3.9). Such refrigerators are commercially available with cooling powers of order 0.1 mW at 100 mK. Our refrigerator was specified to have a minimum cooling power of 0.4 mW at 100 mK. In initial tests, we measured a cooling power of 0.38 mW at 100 mK, before upgrading to a significantly larger circulation pump.

A ^3He - ^4He dilution refrigeration uses the heat of mixing of the two helium

isotopes to provide cooling. When a mixture of ^3He and ^4He is cooled, it separates into two phases, a ^3He -rich phase and a ^3He -dilute phase. The dilute phase, with a ratio $^3\text{He}/^4\text{He} \approx 6.6\%$, sits in the still and the bottom half of the mixing chamber (shown as a darker grey in figure 3.9). The ^3He -rich phase (shown as a lighter grey) floats on top of this mixture in the mixing chamber. At typical operating temperatures, the vapor pressure of the ^4He is many orders of magnitude lower than that of the ^3He in the ^4He . The ^3He can thus be pumped out of the still, and recondensed back into the mixing chamber. As ^3He is removed from the still, a concentration gradient develops between the mixing chamber and the still. As a result, ^3He flows up into the still. To maintain equilibrium in the mixing chamber, ^3He from the concentrated phase crosses the phase boundary into the dilute phase, losing kinetic energy and thus cooling the mixing chamber and whatever may be attached to it.

Pre-cooling the ^3He being condensed into the mixing chamber is crucial to the cooling of the fridge. If the helium enters the mixing chamber at too high a temperature, the heat input will reduce the cooling power of the refrigerator. The incoming helium is cooled using both counterflow and discrete heat exchangers with the outgoing helium. In addition, even before it reaches the still, the incoming helium is cooled in a “1K Pot”, in which liquid helium, drawn from the bath surrounding the refrigerator, is pumped to a low pressure, at which its boiling temperature is reduced from 4.2 K at 1 atm down to 1.3 K.

Commercially available dilution refrigerators use gravity to define the phase separation, and thus are oriented vertically. To combine one with a long magnetic trap oriented horizontally along the neutron beamline, we required a custom-designed cryogenic dewar, consisting of a horizontal section to house the magnet and a vertical section to house the refrigerator. We chose a T-shaped dewar design which is shown

in figure 3.10. The dewar has two cylindrical sections, one oriented vertically and the other horizontally, joined at the bottom of the vertical section, forming an inverted “T” shape. There are two places in which this design becomes complicated: in the “T” section, where the horizontal and vertical sections mate (section 3.5.1) and at the ends of the horizontal section (section 3.5.2).

3.5 Dewar

A cryogenic dewar is like a Russian matryushka doll with layers nested inside one another, each completely surrounding the colder ones inside it. Each layer is separated from the next by vacuum to prevent heat conduction. In principle, the heat flow is radiative (i.e. black body), and so each layer is wrapped with many layers of “superinsulation”, thin ($6.5\ \mu\text{m}$) aluminized mylar,** which has a lower emissivity and reflects a larger amount of the thermal radiation back to the warmer outer surface. The nominal temperatures of the layers are designated by the boiling points of the cryogenic liquids used to cool them: liquid nitrogen (LN_2) at 77 K and liquid helium (LHe) at 4 K.

In the vertical section of the dewar, moving inward from the outside, there is the 300 K (i.e. room temperature) outer vacuum can (OVC), the 77 K liquid nitrogen jacket or shield, the 4 K helium can and the inner vacuum can (IVC). The nitrogen jacket (filled with LN_2 during normal operation) is surrounded on both sides by the OVC vacuum. Inside the OVC is the 4 K helium can, filled with liquid helium and containing the inner vacuum can which houses the dilution refrigerator. The IVC vacuum is separate from the OVC vacuum.

**type “NRC-2” superinsulation from MPI, Winchester, MA




Figure 3.10: A sketch of the T-shaped dewar.




Figure 3.11: A sketch of the inner vacuum can, inside the 4 K helium bath.

The horizontal section is similar. The outermost layer is the OVC. Inside this resides an aluminum shield which is thermally anchored to the liquid nitrogen jacket in the vertical section. Inside the nitrogen shield is a can filled with liquid helium – the same helium bath which continues up into the vertical section. The lower half of the IVC is contained within the helium can and is considerably more complicated (figure 3.11), because access to the experimental cell must be made at the far ends of the magnet. The inner vacuum can is composed of five separate cylindrical sections: first the long vertical cylindrical region surrounding the refrigerator, which extends down into the horizontal section through a long belows; second and third a pair of




Figure 3.12: A sketch of the 300 K Sliding Seal.

large-diameter cylinders, oriented horizontally, and attached to each end of the 4 K can (the IVC end cans); fourth, a cylinder connecting these end cans with the cylinder from the vertical section (the IVC T-section); and fifth, the inner bore of the magnet.

3.5.1 T-section

We assembled the horizontal and vertical sections independently. This required a section in which one could join the thermal link of the refrigerator and assemble each layer in turn around this. In the lower part of the vertical section, the 300 K and 4 K cans were connected using sliding sections and captive rings. The 300 K sliding section (figure 3.12) connects the horizontal section to the vertical section and can be slid up past the o-ring at the bottom of the vertical section, allowing access to the inner connection. The captured ring is sealed to both the vertical section and the sliding section using two o-rings with separate bolt circles. The sliding section connects to the horizontal section with a single o-ring seal.

The 77 K can does not require a vacuum seal, so the connection between the horizontal and vertical sections is made using two semi-cylindrical pieces, each




Figure 3.13: A sketch of the 4 K Sliding Seal.

covering 180° around the vertical axis. Each is bolted to the bottom of the nitrogen jacket and to the top of the horizontal 77 K shield.

The 4 K can has a sliding section similar to the 300 K section, except that it slides inside the vertical section (figure 3.13). The sliding section is sealed with indium o-rings to both a captured ring and the horizontal section. The captured ring makes two indium seals, one to the vertical section and one to the sliding section. Here, the ring is captured on the sliding section. The lower indium seal connects the sliding section to the horizontal section.

Inside the 4 K can, the IVC has a bellows which can be compressed during assembly, allowing access to the heat link and other connections within. These include the fill line for the helium in the trapping region, the heat link attached to the mixing chamber and electrical connections for sensors and heaters. The connections between the horizontal and vertical sections inside the 4 K sliding seal are both demountable, allowing the horizontal and vertical sections to be completely separated, and flexible,




Figure 3.14: A sketch of the IVC end can.

allowing for differences in thermal contraction and mechanical isolation. The heat link is bolted to a length of copper braid in the “T” section, the helium flows through a bellows connected with a VCR^{††} connector, the electrical connections are made with flexible wires all passing through Microtech connectors^{‡‡} here, and the IVC itself is indium sealed to a large bellows in this region.

3.5.2 IVC End Cans

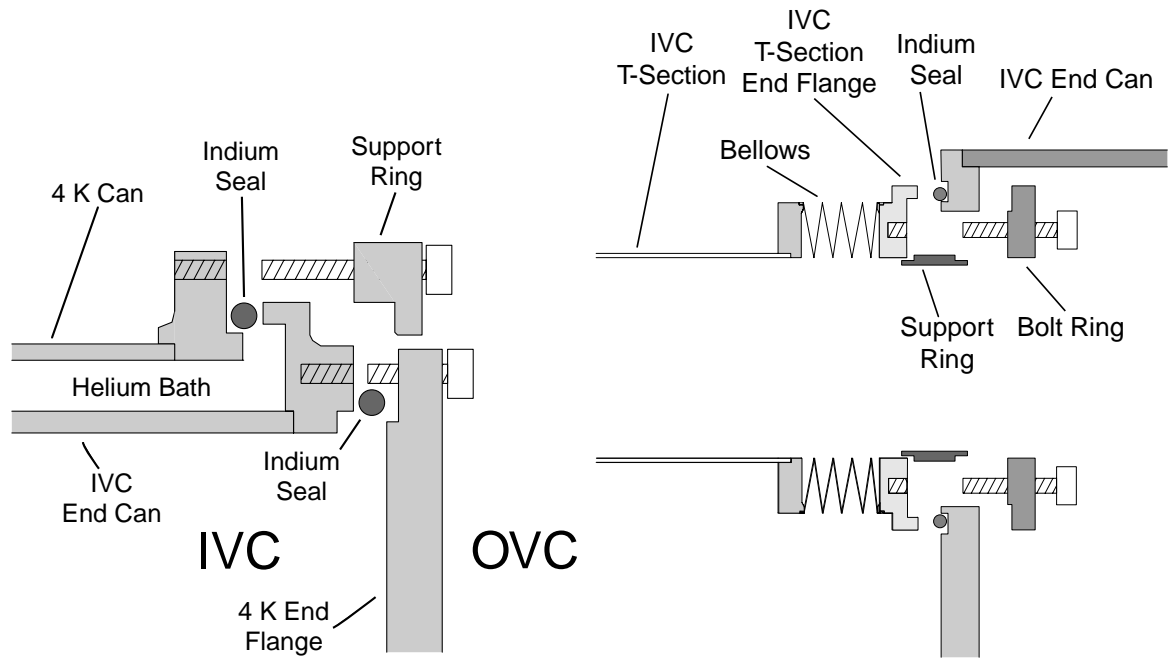
The IVC end cans are buckets inserted into each end of the 4 K can allowing a connection between the magnet bore and the IVC T-Section (figure 3.14). The

^{††}VCR connectors are face-sealed metal-gasket tube fittings available from the Swagelock companies (Cambridge Valve and Fitting). In this experiment, silver-coated nickel gaskets were used.

^{‡‡}Small electrical connectors from Microtech, Inc. Boothwyn, Penn.


helium-filled cell which resides in the magnet bore extends into each of the end cans, where the endcaps of the cell flare out to make connections with the heat link and helium fill line. The liquid helium bath surrounds the IVC end cans, as well as the IVC T-section and the magnet. The outer lip of the bucket is indium sealed to the end of the 4 K can, closing off the end of the helium bath (figure 3.15a). On the upper half of the end can, a bellows connects the horizontal section of the IVC T-section to the end can (figure 3.15b). This seal requires the use of a separate rotatable bolt ring because the two pieces are not keyed with respect to each other. As this section was initially assembled (without the “support ring” shown in figure 3.15b), the bolt ring and, more importantly, the IVC T-section end flange deformed under the force of the tightened bolts, making the indium seal to the IVC end can unreliable. Thus, the support rings were added, preventing these flanges from deforming. On the lower half of the can, a bellows connects the inner bore of the magnet to the end can. A similar bolt ring and support ring are used in this joint (figure 3.15c).

The cell mounting flange consists of two flanges mounted to a length of tube. The inner flange is thin around the heads of the bolts holding the magnet bellows flange to the magnet form and flares inward between these bolt heads to make room for bolts holding it to the magnet bellows flange (see detail of figure 3.15c). The helium-filled cell is supported from the cell mounting flange (section 3.7). Thus, the detector insert is held concentric within the magnet bore. The seal between the IVC end can and the 4 K can again requires rotational freedom in order not to overconstrain the relative positions of the pieces. A separate bolting ring is used to press together the ends of the 4 K can and the IVC end can.



(a) Detail on the IVC end can outer seal.

(b) Detail on the seal between the IVC end can and the IVC T-section.



(c) Detail on the seals between the IVC end can and the magnet form. Shown on the right as assembled.

Figure 3.15: Details of the IVC end can design.

3.5.3 Dewar End Flanges and Windows

The neutron beam must enter the horizontal section of the dewar from one end and light from the neutron decays must exit from the other end. Thus, the end flanges of the horizontal section have an assortment of windows. At the beam entrance end, there are windows to allow the neutron beam to enter the cryostat; at the opposite end, there are optical windows and lightpipes to allow transport of the light from neutron decays out of the cryostat; and above the beamline, on both sides, are optical windows, used for alignment of the cell when the apparatus is closed. All of the windows must in addition perform the functions of the end flanges themselves. At 300 K, the windows must hold vacuum; at 77 K they must block thermal radiation from 300 K; and at 4 K they must both make a cryogenic vacuum seal and block thermal radiation from 77 K.

The neutron entrance windows must transmit neutrons with minimal scattering and absorption. Scattered neutrons can give rise to unwanted backgrounds, in the form of activation, in addition to decreasing the number of neutrons reaching the trapping region. Windows typically used include thin aluminum and beryllium foils and silicon crystals. Single crystal silicon wafers can be thicker than aluminum or beryllium for a given amount of scattering, but have a tendency to crack under pressure. Aluminum could also be used, but the thickness required in our arrangement would result in excessive scattering. In addition to the scattering, the activation of the aluminum itself presents a background problem. Beryllium is ideal in terms of both activation and scattering, but is too brittle to make a vacuum window at a thickness which would not present a background issue (from neutron activation).*

*The purest beryllium foil we could find (from Brush-Wellman, Freemont, CA) was specified as 99% Be, 0.8% BeO, with the 0.2% other including .06% Fe, .02% Al, .01% Mg, .002% Mn, .007%




Figure 3.16: Sketch of the teflon windows.

Without a reasonable way to combine all of the needed properties into a single window material, we used a two layer window, one material to make a vacuum seal and the other to block black body radiation, with both having good neutron transmission properties. Specifically, we used a Teflon vacuum window with a thin beryllium foil to block thermal radiation.

We developed windows made out of clear fluoropolymers, Teflon PFA and FEP[†][56]. Clear plastics in general result in low activation, but the hydrogen in most

Cu, all of which give activation backgrounds.

[†]Teflon PFA (polyfluoroalkoxy) and FEP (fluorinated ethylene-propylene) are optically clear fluoropolymer resins produced by DuPont.

plastics scatters neutrons quite well. Fluoropolymers, on the other hand are hydrogen free. In addition they tend to retain their elasticity at low temperatures, unlike most plastics, which undergo a glass transition and become brittle. A simple compression seal (figure 3.16), where the edges of a Teflon disc are compressed between two metal flanges until the Teflon begins to deform, will provide a superfluid helium leak tight seal at low temperatures. Neutrons pass through three such windows, at 300 K, 4 K and the end of the 250 mK helium-filled cell. The windows used at 250 mK are 250 μm Teflon PFA cut 6.35 cm in diameter, leaving a 6.35 mm wide section compressed around the circumference of a window 5.08 cm in diameter. The other windows are essentially the same but with a PFA disc 7.62 cm in diameter making a window 6.35 cm in diameter. To block black body radiation a 50 μm thick beryllium disc[‡] covered the Teflon window at 4 K (figure 3.16), and a second covered the beam entry hole in the 77 K end flange. The beryllium was sandwiched between two 25 μm pieces of Teflon PFA for safety in handling. The flanges holding the beryllium-Teflon sandwich were vented to allow the space between the Teflon and the beryllium to be pumped out.

The light transmission windows, on the opposite end of the cryostat, are made from ultra-violet transmitting (UVT) acrylic (polymethyl methacrylate or PMMA).[§] The more common grade of acrylic (ultra-violet absorbing, UVA) is doped with a dye that absorbs UV light, to prevent the plastic from degrading under exposure to UV. For UVA windows, the effect of this dye on visible light exposure is quite small (less than the reflection losses), however in a long lightpipe the absorption of violet and near-UV light would noticeably degrade our signal. Using a krypton ion laser tuned

[‡]We used grade PF-60 beryllium sheet, purchased from Brush-Wellman, Freemont, CA.

[§]Our UVT acrylic was manufactured by Townsend-Glasflex and purchased from Boedeker Plastics.

to the 407-415 nm range, we measured the bulk transmission through various acrylic lightpipes, angling the laser so that the light passed from one end of the lightpipe to the other without reflecting off the walls and correcting for the expected 4% reflection at each surface (from $R = \left(\frac{n_1-n_2}{n_1+n_2}\right)^2$ for normal incidence and an index of refraction of 1.49 for acrylic). After correcting for the 4% lost to reflections at each surface, a bulk loss of 42% was observed in a 1 m lightpipe of UVA acrylic. Under identical conditions a 50 cm UVT lightpipe gave a 9% bulk loss. Though clear to visible light, acrylic blocks black body radiation[¶], and can be used directly to make a seal either with a viton o-ring (at 300 K) or an indium o-ring (at 4 K). These windows and lightpipes are described in section 3.9.

Above these windows, on both ends of the dewar, are optical windows used for aligning the cell with respect to the neutron beam. These windows were initially made of glass (BK-7), but several were replaced with acrylic over the course of the experiment. Once the alignment was complete, the windows at 300 K were covered to prevent light from entering the dewar.

3.6 Alignment

In order for the neutron beam to enter the trapping region, the apparatus must be properly oriented with respect to the beam axis. The alignment is initially performed while the dewar is at room temperature, with both endcaps removed. One sights through the open apparatus, along the beam axis, with a theodolite (a standard surveying telescope), and adjusts the position of the apparatus until the cell

[¶]Acrylic blocks IR with a cutoff between 2 μm and 3 μm [58]. For comparison, the black body spectrum at 300 K is peaked at 10 μm .

is correctly aligned. Thin wires ($75\text{ }\mu\text{m}$) stretched around bolts in the bolt circles of the cell endcaps cross at the cell axis, making an easily localized alignment marker.

As the cryostat is cooled, however, each layer is free to move independently under thermal contraction. Each layer is independently supported from the top plate, except for the helium-filled cell, which is held centered within the magnet bore, and so moves with the 4 K horizontal section. In order to realign the apparatus when cold, alignment markers were positioned on the ends of the heat links, visible through a series of optical windows on the upper half of the dewar end cans. Sighting on these markers, we observed that the cell rises approximately 3 mm when cooled from 300 K to 4 K. The dewar was then lowered until the alignment markers were in the same position as when the experiment was properly aligned at 300 K. Though these alignment markers were not on the beam axis itself, at a distance of less than 20 cm, the thermal contraction of the heat link changed their relative position less than 0.5 mm, within our alignment tolerances.

The position and orientation of the cell was adjusted by moving the dewar using adjustment mechanisms built into the support stand (figure 3.17). The top plate (a 2.5 cm thick, 1.42 m square aluminum plate) rests on four Teflon pads on top of the stand, and can be translated in the two horizontal directions and rotated around the vertical axis. The height and orientation around the remaining two axes, are adjusted using an adjustment mechanism integral with the legs. This consists of a threaded rod (1" diameter, 8 turns per inch) fixed to the upper section of each leg, onto which is threaded a nut resting on the leg beneath. By turning this nut, the top plate and thus the apparatus can be either lowered or raised. Together, the adjustment of the top plate and legs of the stand allow the position and orientation of the cell to be completely adjusted.




Figure 3.17: Sketch of the adjustment mechanisms built into the dewar support stand. Threaded rods integral to the legs of the stand permit vertical adjustment. At each corner of the stand two bolts push on a block attached to the top plate, sliding it on Teflon pads.

3.7 250 mK Section

The isotopically pure ^4He (section 2.5.1) and detection insert are contained within the inner vacuum can and held fixed with respect to the trapping magnet and neutron beam, while thermally anchored to the dilution refrigerator. The helium is contained within a tube made of cupronickel^{||} which extends through the magnet and into the IVC end cans, where a copper end cap is soft-soldered to each end. The

^{||}70% Cu and 30% Ni, a non-ferromagnetic alloy otherwise similar to stainless steel




Figure 3.18: Sketch of a cell end cap with teflon window attached.

endcaps (figure 3.18) provide a sealing surface for either the Teflon window or the light transmission window. The endcaps are also thermally anchored to the heat link which extends to the dilution refrigerator. On the neutron beam entrance end, a Teflon window contains the helium within the cell. At the opposite end, an optical window on a “snout” is indium sealed to the end cap. The snout will be described further in section 3.9.

The end cap and thermal links to the dilution refrigerator are machined from oxygen-free high conductivity (OFHC) grade copper. Connections between the individual pieces are made by clamping or bolting together two copper surfaces which have been cleaned and polished. The cell end caps each have a fin that attaches to an L-shaped piece which extends to the level of the IVC T-section. The top of the L-shaped piece bolts to a flat section of a 1.9 cm diameter, 52.5 cm long copper rod which extends into the center of the IVC T-section. In the T-section, the copper rods (one from each end) are bolted to a flexible connector which contains four 1.5 cm lengths of copper braid, which allows the lower sections of the heat link to be mechanically independent from the upper section. The braided pieces were constructed by brazing one end of a braid of copper wires under compression between two pieces of copper. The braids were compressed to maximize the area of contact to these copper pieces, which provided flat bolting surfaces for attachment to the lower section of the heat link. The other ends of the braids were brazed, again under compression, to a flat section at the base of a 2.5 cm diameter by 8.9 cm long cylinder which could be clamped to the 24 cm long 2.5 cm diameter rod above it, which bolted to the bottom of the buffer-cell housing. The buffer-cell housing was a copper cylinder hollowed out to house the buffer cell (described below). The top of the buffer-cell housing bolted to the bottom of the mixing chamber, completing the heat link between the dilution refrigerator and the cell.

We measured the thermal conductance of the heat link from the mixing chamber through the braid, with the lower sections of the heat link (and the cell) not attached. Thermometers recorded the temperature of the mixing chamber and the temperature of the bottom of the braid, as a heater applied power to the bottom of the braid. Shown in figure 3.19 is the measured conductance ($\frac{\text{Power}}{T_{\text{braid}} - T_{\text{MC}}}$) vs. the

measured temperature of the braid. Most importantly, this measurement confirmed that the heat link is sufficient to keep the cell at about 250 mK with a heat input of up to 100 μW . In addition, we can compare the measured conductance to what we might expect. The thermal conductance of a copper rod is just given by $\Omega = \frac{A\kappa}{\ell}$, where κ is the thermal conductivity of the copper and ℓ and A are the length and cross-sectional area of the rod. The two copper rods, each 2.5 cm in area with a total length of 32.9 cm give an $\frac{A}{\ell}$ of 0.156 cm. The housing is another 16 cm long, with two sections about 3.5 cm² each, for a total of 7 cm². Combining these, the heat link up to the braid has an $\frac{A}{\ell}$ of 0.12 cm. Taking a conductivity of roughly 0.2 W cm⁻¹ K⁻¹ at 200 mK[50], this gives an expected conductance of 2.4×10^{-2} W K⁻¹ at 200 mK, compared to a measured conductance of 2.3×10^{-4} W K⁻¹. This seems to indicate that the majority of the thermal resistance is in the braid, however the conductivity of the solid section is strongly dependent on the purity of the copper and was not separately measured. The braid was both necessary for mechanical isolation and sufficient to our cooling needs. Ignoring the resistance of the rest of the heat link (expected to be no more than that of the upper section, above the braid), these measurements correspond to a minimum temperature for our helium bath as a function of the heat load into the cell (figure 3.20). In practice, we were able to cool the helium bath to a temperature between 215 mK and 250 mK, however, the mixing chamber was considerably warmer than in the thermal conductivity measurements above (190 mK compared to 50 mK), indicating a significant heat load on the refrigerator itself. A likely cause of this additional heat load is the pumping line connecting the buffer cell to room temperature (described below).

The cell was filled with isotopically purified ⁴He by condensing it in from room temperature. The helium is stored as a room temperature gas at a pressure less than




Figure 3.19: Plot of the thermal conductance of the upper section of the heat link (through the braid) as a function of the braid temperature.

two atmospheres. The storage volume was connected to the cell through a series of stainless steel tubes and all-metal valves. A pair of 1/4-inch stainless steel tubes, a fill line and a pumping line, connected the storage dumps to the top of the dewar. The fill line connected to two internal fill lines, each of which could be closed independently outside of the dewar. The fill lines each pass into the dewar, are thermally anchored at 4 K, pass into the IVC, pass through the 1 K pot, wrap around the still, pass through the continuous heat exchangers, and are connected to the buffer cell. The pumping line enters the IVC and passes directly to the buffer cell, heat sunk at the 1 K pot and still. When this bypass line is open, the impedance between the buffer cell and storage dumps is greatly reduced. The buffer cell is positioned below the mixing chamber and provides an extra volume of helium such that when the cell is cooled from 4 K to 250 mK (over which range the helium contracts about 15%) the




Figure 3.20: Plot of the minimum cell temperature as a function of heat load, given by the thermal conductivity measurement shown in figure 3.19.

cell remains full. At the bottom of the buffer cell, two 1/8-inch stainless steel tubes extend to the two ends of the cell. The fill lines are soldered into a shallow groove on the horizontal section of the heat link for thermal anchoring. At the center of the T-section, each fill line can be disconnected using a VCR connector, allowing the horizontal and vertical sections to be completely separated. In the IVC end can, the fill lines are connected using VCR fittings to the cell. The VCR connectors are silver-soldered directly into the cell endcaps.

The heat link and the fill lines are flexible near the T-section and thus the cell is not supported directly from the dilution refrigerator. The lower half of the heat link and the cell are supported and positioned from the 4 K magnet form by a set of Kevlar threads. One supports the cell and the second positions the heat link. The cell support is mounted to a flange attached to the magnet form (figure 3.15c).

The cell mounting and positioning scheme is sketched in figure 3.21. Kevlar




Figure 3.21: A sketch of the cell positioning scheme.

is used as the support because its high tensile strength (1800 MPa at 77 K, for the braided cord, with a thermal conductivity of $2 \times 10^4 \text{ W cm}^{-2} \text{ K}^{-1}$ at 4 K[59])) makes it possible to support the cell from 4 K using narrow threads. The Kevlar thread is mounted to prevent any kinks in the Kevlar, which would limit the amount of weight that could be supported by a given thread. The thread** wraps around the cell (about 480°) and is secured on each end by a cell positioning device. The threads wrapped around the cell are composed of four twisted strands, each with a diameter of $400 \mu\text{m}$, able to support a total weight of 370 kg. Combined with the threads wrapped around the heat link, composed of two twisted strands, the kevlar threads from 4 K to the cell give a total heat load of $5 \mu\text{W}$.

The cell positioning device is shown in figure 3.22. The heart is a threaded

**The thread supporting the cell was composed of four twisted strands, each with a diameter of $400 \mu\text{m}$. The thread around the heat link was composed to two twisted strands.




Figure 3.22: A sketch of a cell positioning device, in cross-section.

brass rod (1/4"-80) with a hole in the center. The Kevlar thread passes through this hole and is epoxied (using stycast 2850^{††}) in place, as shown. The brass threaded rod is mounted in a square bracket and the length of the Kevlar thread can be adjusted by screwing the rod in or out of the mount. The cap is separate from the screw so that the Kevlar thread is not twisted and thus weakened when adjustments are made.

3.8 Neutron Shielding

Neutrons scattered from the beam are a potential source of background. In front of the cryostat, neutrons can scatter in the air (about 1% scatter per 10 cm of air); on their way into the experiment, a small fraction ($\lesssim 1\%$) scatter in each window*; and in the experimental cell the neutrons scatter in the helium, with a tiny fraction of these scatters producing trapped UCN. When the neutron beam is open to the cell, $4 \times 10^8 \text{ n s}^{-1}$ enter the helium bath. About a meter later, those which

^{††}Stycast 2850 is a black two-component filled epoxy from Emerson & Cuming.

*We measured an absorption of 3.5% and a scattering of 1.9% for 760 μm thick PFA windows (three times the thickness of our windows)

haven't scattered are absorbed by the beam dump. The scattering cross section for cold neutrons in liquid helium[60] decreases rapidly with wavelength. Although at 8.9 Å only 6% of the neutrons scatter, at 4.5 Å, nearly 40% scatter. As a result, $1-2 \times 10^8$ n s⁻¹ scatter into the walls, potentially producing backgrounds if they activate materials.

In order to prevent neutron activation of materials surrounding the trapping region, particularly near the detector insert and lightpipes, neutron absorbing materials are placed such that scattered neutrons do not reach materials that could activate. Several different shielding materials were used in this experiment, all of them containing either lithium or boron which have high cross sections for absorbing neutrons.

The precise materials used depended on the physical properties and locations of the materials. For example, boroflex and lithium plastic have the advantage of being flexible in the form of 3-6 mm rubber sheets, which can be cut to shape with a razor blade, or rolled into tubes to place around the beamline. For permanent installations, hexagonal boron nitride (hBN or just BN), which has the color and consistency of chalk, was used.[†] All of the shielding, along the beamline, inside the cryostat (except for the final collimator and the beamstop, both of which are graphite-coated boron carbide) are graphite-coated boron nitride (figure 3.23).[‡] Elsewhere in the cryostat, parts were sprayed with aerosol BN[§] to absorb stray neutrons. The graphite coating was used to minimize the luminescence we observed from the relaxation of color

[†]The solid hexagonal boron nitride purchased from Carborundum (www.carbobj.com) was grade AX05, with the exception of one piece on the 77 K end flange, which was grade A.

[‡]The graphite coating was applied using Aerodag G, an aerosol spray containing micron-sized graphite particles in isopropyl alcohol, a product of Acheson Colloids Company.

[§]in two varieties: Combat BN aerosol spray from Carborundum (www.carbobj.com) and BN aerosol lubricant from Alfa-Aesar, Ward Hill, MA




Figure 3.23: Sketch of the windows in the dewar end flanges, on the beam entry end, showing the interlocking boron nitride tubes shielding the beam.

centers produced when the cold BN was exposed to neutrons[54]. Boron oxide, was initially used as a clear beamstop, to allow light to pass, but was eventually replaced because it luminesced.[¶] Boron carbide, an extremely hard black mineral, was used instead despite the loss in signal due to its opacity, because it was found that its luminescence was not observable and many orders of magnitude below that of boron nitride.[55]

The sensitivity to activation-induced backgrounds was greatest inside the de-

[¶]The pure boron oxide also proved quite difficult to machine, being extremely hard and reasonably hygroscopic, clouding up over a period of several hours to a day when exposed to air.




Figure 3.24: A sketch of the interlocking boron nitride tubes.

tector insert itself, and in the lightpipes. Neutrons passing through the trapping region were absorbed in the beamstop. To minimize activation in the detection system the materials comprising the detector insert were carefully chosen to contain materials that do not activate. Surrounding these materials was a shield of boron nitride to absorb the remaining neutrons to prevent activation of the surrounding cupronickel tube, the magnet and the rest of the apparatus.

This shielding was made from a series of interlocking boron nitride cylinders (5 cm long, with an outer diameter of 4.15 cm and a wall thickness of 1.5 mm). The overlapping sections eliminate the path for neutrons to escape through the shielding. The BN had an attenuation length of 100 μm to 40 μm over the relevant neutron wavelengths (section B.4). When it was discovered that boron nitride was luminescing, it was necessary to either replace the BN with a different shielding material or to shield the detector from the luminescence. First, we coated the BN with a thin layer of graphite (Aerodag). Though black to the eye, this layer attenuated the luminescence by less than two orders of magnitude. We machined the boron nitride tube to a wall thickness of 1.1 mm to provide sufficient room for (non-interlocking) tubes of

graphite (15 cm long with a wall thickness of 0.2 mm) to be inserted inside the boron nitride tube (figure 3.24). This arrangement proved sufficient to allow the detection of the trapped UCN.

3.9 Detection System

At the heart of the experiment, immediately surrounding the region in which the neutrons are magnetically trapped, is the detector insert, a UVT acrylic tube whose inner surface is coated with TPB-doped polystyrene. This detector, together with a series of windows and lightpipes leading to a pair of PMTs, form a detection system allowing neutron decay events to be recorded as a function of time. Each beta-decay event in the trapping region gives rise to a pulse of EUV light in the helium (section 2.3). This light is downconverted to the visible range by the TPB and transported past the opaque beamstop in the acrylic tube to a 3.8 cm diameter lightpipe which extends to the end of the cell. At the end of the cell, the light passes through a pair of UVT acrylic windows into a 3 inch (7.6 cm) diameter lightpipe which extends out of the dewar (figure 3.25). Outside the dewar, the light is coupled into a pair of 2 inch (5 cm) PMTs which detect the downconverted scintillation light, registering the decay of a neutron in the trap.

Scintillation light from trapped neutron decays is absorbed in a thin layer of TPB-doped polystyrene coated on the inner surface of a UVT acrylic tube. Light emitted by the TPB in a direction along the tube* will be carried down the length of the tube by total internal reflection. In order to increase the amount of light reaching

*specifically, light emitted such that $n_a \sin \theta > n_h$, where n_a is the index of refraction of the acrylic (≈ 1.49), n_h is the index of refraction of liquid helium (≈ 1.024) and θ is the angle between the direction of the light and the vector normal to the surface




Figure 3.25: Sketch of the windows in the dewar end flanges, on the light exit end.

the lightpipe, the acrylic tube is wrapped in a sheet ($175\text{ }\mu\text{m}$ thick) of Tyvek[†]. The Tyvek diffusely reflects the light not captured in the walls of the tube, allowing some to reach the lightpipe. Around the Tyvek are the graphite and boron nitride tubes described in section 3.8.

At the neutron beam entrance end of the trapping region, the boron carbide

[†]A DuPont product, Tyvek is a “spunbonded olefin”, made from long, fine filaments of high density polyethylene bonded together by heat and pressure.




Figure 3.26: A sketch of detector insert. (not to scale)

collimator is positioned against the acrylic tube, with a piece of Tyvek sandwiched in between to help reflect light back towards the detectors. At the opposite end of the trapping region, the boron carbide beamstop is mounted inside the tube. The acrylic tube is glued to a solid 3.8 cm diameter UVT acrylic lightpipe using a clear epoxy (Stycast 1266[‡]). The Tyvek wrapping, graphite and boron nitride extend the entire length of the acrylic tube and lightpipe. At the end of the cell, an aluminum and acrylic “snout” (figure 3.28) is indium sealed to the copper end flange.[§] The lightpipe was epoxied (Stycast 1266) to an acrylic window at the end of the snout to prevent a gap opening up between the window and the lightpipe during thermal contraction and filling of the cell. This joint was made with the cell in place and the snout sealed to the end of the cell. (The acrylic tube and lightpipes fit very snugly into the cell giving no flexibility in the relative positioning of the snout and lightpipe, thus

[‡]Stycast 1266 is a transparent two-stage epoxy from Emerson & Cuming.

[§]Although the teflon windows were optically clear, they bowed under pressure, so that there would have been a larger gap (more than 1 cm) between the end of the lightpipe in the cell and the 4 K window.




Figure 3.27: A sketch of the Vespel pegs.

requiring the joint to be made in place.) The snout extended to the 4 K window, a 10 cm diameter 6.35 mm thick acrylic window, indium sealed to the 4 K flange. The snout was held a fixed distance from this window by a pair of pegs made from Vespel[¶], a soft black material with extremely low thermal conductivity ($2 \times 10^{-5} \text{ W cm}^{-1}\text{K}^{-1}$ at 1 K[61]). The pegs (figure 3.27) were sharpened to a point to minimize the contact area with the 4 K surface and mounted in the hollow end of a copper post attached to the compression ring on the cell end cap. Washers placed between the post and the compression ring allowed final adjustment, so that the snout window was positioned 1-2 mm from the 4 K window. The pegs pushed lightly against the window so that the distance between the end of the snout window and the 4 K window remained fixed.

The snout was machined from aluminum and acrylic (figure 3.28). The acrylic pieces were epoxied to each other and to the aluminum using Stycast 1266. The thermal contraction of acrylic is considerably greater than that of aluminum ($\Delta\ell/\ell = -1.2\%$ for acrylic and -0.4% for aluminum[50]). Thus, when the snout was cooled, the acrylic contracted considerably more than the aluminum. The acrylic tube was placed around the aluminum fin, so that it contracts onto the aluminum rather than pulling away. The aluminum fin also flexes with the contraction, putting less strain on the

[¶]Vespel SP-22, made by DuPont, is a polyimide resin filled with 40% graphite.




Figure 3.28: A sketch of the aluminum “snout”

epoxy joint[47]. Not shown in the figure is a groove inside the snout holding a teflon o-ring which was meant to contract around the lightpipe, in order to prevent it from sliding away from the window.^{||} In practice, despite the teflon o-ring, the lightpipe pulled away from the window by about 1 cm, so that some of the light coming out of the end of this lightpipe failed to reach the 3 inch diameter lightpipe. In the end, the lightpipe was simply glued directly to the window with Stycast 1266. The glue

^{||}The Tyvek, graphite and boron nitride are all relatively low friction surfaces, and acrylic contracts by 1.2% in cooling from 300 K to 4 K[50], more than 2 cm over the whole two meter length of the lightpipe and tube. In addition, a pressure differential could have pushed on the lightpipe while the helium was being condensed into the cell.




Figure 3.29: A sketch of the copper snout.

joint was observed to be optically clear over more than half the area of contact, with some scattering centers (i.e. cloudiness) observed across the rest.

Another version of the snout also managed to mount a rigid, superfluid-tight optical window several inches from an indium seal to the cell end cap (figure 3.29). Instead of acrylic and copper, this snout used glass** and OFHC copper. Since the thermal contraction of copper is greater than that of g ($\Delta\ell/\ell = -0.05\%$ for glass

**The glass used was “herasil”, a fused silica glass purchased from Rolyn Optics.




Figure 3.30: A sketch of the thermal connection between the 3-inch lightpipe and the 77 K end flange.

and -0.33% for copper[50]) this design places the window inside the copper snout. A $125_{25}^{+125} \mu\text{m}$ thick copper fin provides strain relief as the copper contracts around the glass window. In order to avoid machining away large amounts of copper, the snout was made in two pieces: one incorporating the indium seal and bolt circle, the other ending in a thin copper fin. These two pieces were soft-soldered together before the window was glued in place. As in the aluminum snout, the window was sealed in place using Stycast 1266. The aluminum snout was used in collecting all of the data described in Chapter 4.

On the opposite side of the 4 K window is a 3-inch diameter UVT acrylic lightpipe, which extends from the 4 K window to the outside of the dewar. The




Figure 3.31: A sketch of the seal around the 3-inch lightpipe.

lightpipe is thermally anchored by 64 aluminum foil ($100\ \mu\text{m}$) strips (3 cm by 8 cm) to the 77 K end flange. The strips of foil are attached to the end flange by a compression ring and to the lightpipe by four cable ties made of Teflon PFA. Teflon contracts even more than acrylic on cooling, with a $\Delta\ell/\ell = -2.1\%$ [50], tightening down as the cryostat is cooled.

At the 300 K end flange, a viton o-ring seals around the circumference of the acrylic lightpipe using a compression fitting* (figure 3.31). The base of this fitting seals to the 300 K flange with a viton o-ring. A Teflon o-ring fits into a groove here,

*This compression fitting is a scaled-up version of an “Ultratorr” fitting, a standard rubber o-ring shaft-seal tube fitting, available from the Swagelok companies (Cambridge Valve and Fitting).




Figure 3.32: A sketch of the optical elements connecting the 3-inch lightpipe to the PMTs.

to center the lightpipe within the compression fitting and a viton o-ring seals against the lightguide. As the nut is screwed onto the base of the fitting, the compression ring presses the viton o-ring against the tapered inner surface of the base, forcing it against the lightpipe, as shown. To prevent the lightpipe from moving due to the pressure difference when the dewar is evacuated, a clamp is tightened around the lightpipe just beyond the end of the compression fitting. Three set screws in this clamp allow the final position of the lightpipe to be adjusted externally when the dewar is cold.

The 3-inch lightpipe is coupled to a pair of 2-inch photomultiplier tubes using a series of optical elements (figure 3.32). The light is first “condensed” using a Winston cone condenser[†]. The condenser collects light travelling parallel to the axis of the lightpipe very efficiently (80% in combination with splitter) into the 2-inch diameter end of the condenser, with a greater divergence, while collecting light near the maximum angle in the acrylic much less efficiently. In our geometry this is a reasonable choice, since the light emerging from the end of a long lightpipe is more forward-peaked than the initial spectrum.[‡] The light exiting the condenser is split into two 2-inch lightpipes, using a “splitter”. The splitter is made from two 2-inch lightpipes cut at a 45°-angle and epoxied together, forming an ell. The corner is machined away leaving a flat elliptical surface with a 2-inch minor axis and a 2.8-inch major axis which is attached to the Winston cone. The Winston cone and the splitter were diamond turned, polished, and aluminized[§], to increase the light transmission. We measured an 80% transmission from the 3-inch lightpipe to the two ends of the splitter. Optical grease was used in the connections between the 3-inch lightpipe and the condenser, between the condenser and the splitter, and between the splitter and each photomultiplier tube. All of the optical elements were contained within a set of aluminum cans attached to the compression fitting. The cans were sealed with viton o-rings, both blocking light and creating a space through which dry N₂ gas could be flowed around the PMTs to prevent helium contamination of the PMTs. Behind each PMT, the aluminum cans were closed with flanges of insulating nyatron[¶], through

[†]a short length of acrylic lightpipe with the outer surface machined to the shape of a Winston cone[62]

[‡]In addition to travelling a longer total distance through the acrylic than axial light, the high-angle light bounces off the surface of the acrylic many more times, each time with a small chance of scattering out of the lightpipe.

[§]These pieces were manufactured by Contour Metrological Manufacturing.

[¶]polyethylene loaded with a black molybdenum sulfide dye, available from McMaster Carr




Figure 3.33: A histogram of pulse areas from the anode of a Burle 8850, using a charge-integrating preamplifier, showing distinct peaks for one-, two- and three-photoelectron events. The PMT is being illuminated by an LED operating at low current.

which the high-voltage and signal connections to the PMTs were made. Three springs on each flange pressed the PMTs against the splitter to assure optical contact.

The photomultipliers are Burle model 8850. A photomultiplier amplifies a weak light signal by successively converting the light into an ionized electron in the “photocathode”, and then amplifying the signal in discrete steps. The signal is read out at the anode as a discrete current pulse. The ’8850’s distinguishing feature is that the first dynode is gallium phosphide, which has a very low ionization threshold, and produces about an order of magnitude more electrons than a typical dynode. In most PMTs, statistical fluctuations in the number of electrons ionized in the first




Figure 3.34: Pulse area spectrum from 360 keV beta line spectrum.

dynode wash out the distinction between a shower triggered by one photoelectron from the photocathode and a shower triggered by two photoelectrons. The higher amplification in the first dynode of the '8850 makes it possible to distinguish one-, two- and three-photoelectron events. To see this, we took data with an '8850 looking at a light-emitting diode at low current (figure 3.33).

3.9.1 Calibration

A series of calibrations were performed, to determine the efficiency for detecting neutron decay events of the detector insert and light collection systems as a whole. A variety of downconverters were tested, and TPB was chosen because it was found to have the highest quantum efficiency[63]. Detectors were tested, using alpha and beta sources which create scintillations in liquid helium. These tests were used

to optimize the detection insert's parameters[64]. The assembled detector insert and light collection system (without the Y) were tested, in the cryostat at 250 mK using a ^{113}Sn radioactive source. This source has a beta line spectrum at 360 keV which is near the median energy of neutron decay betas. This measurement, performed with a single PMT at the end of the Winston cone, gave a pulse height spectrum that corresponded to between 5 and 5.5 photoelectrons (figure 3.34). A series of tests using the higher signal from an alpha source gave the relative detection efficiency for EUV light produced at different longitudinal positions along the cell. The efficiency varied from roughly 75% to 125% (relative to the efficiency at the center of the cell) along the length, and was approximated as a linear dependence. A simple Monte Carlo (source code in appendix F.1) was used to calculate our efficiency for detecting trapped neutrons. Each Monte Carlo event was characterized by a random position along the length of the cell and a random beta energy, according to the known beta spectrum (equation A.19). The results of the ^{113}Sn calibration gives an output of 14.5 ± 1.5 photoelectrons per MeV of electron kinetic energy. The expected signal is corrected for the measured 80% efficiency of the Y splitter (i.e. 40% of the photons into each PMT, with 20% lost or reflected back towards the cell) and for the dependence of the signal on the position of the event in the cell. This expected number of photoelectrons (in general, not an integer) corresponds to a Poisson distribution of probabilities of giving each (integer) number of photoelectrons.^{||} Finally, given a total number of photoelectrons, the probability of the decay being detected is given by the probability of these photoelectrons being distributed with at least two in each

^{||}For example, an expected signal of 5.2 photoelectrons corresponds to a 13% probability of 3 photoelectrons, a 17% probability of 4 photoelectrons, and so on.

photomultiplier.** Requiring coincidence between two photoelectrons in each photomultiplier results in detection of $31 \pm 4\%$ of neutron beta-decays in the trapping region (where the uncertainty is primarily due to the uncertainty in the ^{113}Sn calibration). In principle, the imperfect gain dispersion of the PMTs means that requiring an anode current pulse corresponding to at least two photoelectrons will sometimes ($\approx 2\%$) pass an event in which there was actually only one photoelectron or cut an event in which there actually were two photoelectrons. With single-photon backgrounds (such as luminescence), where the actual rate of two simultaneous photoelectrons is orders of magnitude lower than the rate of single photoelectrons, this rate can be quite significant. In the case of the neutron signal, however, the expected rate of two simultaneous photoelectrons in a PMT is of the same order of magnitude as the rate of single photoelectrons, making this a small correction.

3.10 Background Suppression

There are several different sources of “background” events recorded alongside events triggered by neutrons decaying inside the detector insert. Such events could arise from interactions which create multiple photons in a short time or from uncorrelated events which randomly coincide. In the first category, we can conceptually distinguish between events from radioactive materials in the apparatus, either naturally occurring or short-lived isotopes created by neutron absorption and events caused by particles coming from outside the experiment, either cosmic rays (primarily

**For example, 4 photoelectrons have a $6/16$ chance to be distributed 2 and 2, 5 photoelectrons have a $10/32$ chance to be 2 and 3, and a $10/32$ chance to be 3 and 2, and 6 photoelectrons have a $15/64$ chance to be 2 and 4, a $20/64$ chance to be 3 and 3 and a $15/64$ chance to be 4 and 2, giving 37.5%, 62.5%, and 78.125% chances of detection for 4, 5, and 6 photoelectrons, respectively.

muons), or fast neutrons and gamma rays, both from the nuclear reactor and from materials around the experiment (especially ^{40}K in concrete). The neutron shielding materials used to prevent activation backgrounds, and to block luminescence light are described in section 3.7. In this section, minimization of backgrounds arising external to the experiment is described.

Cosmic ray muons, scintillating either in the liquid helium or in the acrylic give rise to about 8 s^{-1} of triggers in our detector. The muons giving rise to these events can be detected outside the apparatus allowing the active removal of these events. The muons were detected using plastic scintillators surrounding the bottom of the dewar. These muon detectors, or muon “paddles”, so called because of their shape, are standard fare in experimental high energy physics and we were fortunate enough to borrow some from high energy physics groups at Harvard and the University of Maryland.* Each paddle consists of a $2.5\text{ cm} \times 32\text{ cm} \times 1.42\text{ m}$ slab of scintillator, an acrylic condensing region (40 cm long), and a photomultiplier tube (figure 3.35). The paddle was wrapped with a reflector to increase the light collection efficiency and in black rubber to block room light from entering the detector.

When a muon passes through the scintillator it deposits about 2 MeV of energy per centimeter [3] in the form of Čerenkov radiation and scattered ionized electrons. The Čerenkov light is converted by the scintillator into blue or green light, which is transmitted to the PMT and detected.

Three muon paddles reside below the horizontal section of the dewar (figure 3.36) and a fourth over the “Y” section of the lightpipe and the PMTs. The active vetoing scheme is described further in section 3.11.1. The muon paddle directly

*Thanks to Dr. George Brandenburg of the Harvard University High Energy Physics Laboratory and to Professor Elizabeth Beise of the University of Maryland for the loan of several muon paddles.




Figure 3.35: A sketch of a muon paddle, showing the rectangular slab of scintillator which acts as the detection region and the tapered section which connects to the photomultiplier tube.




Figure 3.36: Sketch of the orientation of three muon paddles relative to the horizontal section of the dewar.

beneath the cryostat vetoed about 4.8 s^{-1} , each of the side paddles vetoed about 1.7 s^{-1} , and the fourth paddle over the PMTs vetoed another 1.1 s^{-1} . The total veto rate was about 8 s^{-1} , less than the sum of these rates, as the paddles overlap slightly.

Gamma rays originating from the reactor, neighboring experiments and the environment in general were attenuated using lead bricks surrounding the dewar.

The lead bricks were supported by two walls of cement blocks stacked parallel to the beamline, with two 1.3 cm thick aluminum plates resting on top of the walls. Lead bricks (each measuring 2" x 4" x 8"), were placed around the horizontal section of the dewar and muon paddles to form a 4" thick shield. At the front, a 4"-square hole lined with boroflex allowed the neutron beam to enter the apparatus. One muon paddle rested on the lead "floor" and the other two leaned against the side walls. The top of the lead enclosure was constructed with a large square gap around the vertical section. At the back, the ends of the muon paddles and the "Y"-shaped lightpipes were covered on the top and sides, but open at the back.

Unfortunately, the varying of multiple parameters at once makes it impossible to determine exactly what rate of background was eliminated by the lead. However, the lead shielding seems to have removed of order 10 s^{-1} . A number of benchtop studies using two PMTs attached to each end of an acrylic lightpipe showed that even without the helium-filled cell and TPB-coated tube, there should have been several counts per second coming from the lightpipe itself which could be drastically reduced by stacking lead around this assembly. Cosmic ray events were similarly detected in such a setup, without helium, in coincidence with an overlain muon paddle.

3.11 Data Acquisition

3.11.1 NIM

The schematic picture of the data acquisition system is shown in figure 3.37. The anodes of the PMTs looking at the cell and the muon paddles are connected to the inputs of a NIM LeCroy Research Systems (LRS) Model 612 Amplifier, which




Figure 3.37: Data Acquisition Schematic

amplifies the voltage of the signal by a factor of ten, and fans the signal out into two outputs. The amplified signals from the PMTs looking at the cell (PMT I and PMT J) are connected both to the digital oscilloscope and to a linear fan-out (NIM LRS Model 128).

The outputs of the fan-out pass into a NIM octal Discriminator (Phillips Scientific Model 710), which has independently tuneable thresholds and output widths, and three separate output channels for each input. A discriminator takes an analog input and produces a digital pulse each time the input rises above the preset threshold. The amplified signal from PMT I is fanned out to channels 1 and 7 of the discriminator, and the amplified signal from PMT J to channels 2 and 8. The

thresholds for channels 1 and 2 were set to trigger on nearly every pulse which would give a charge-integrated signal of 2 or more photoelectrons on the scope (called $I_{1.5+}$ and $J_{1.5+}$), as determined empirically. The thresholds for channels 7 and 8 were set to correspond to a voltage that was in the valley between the noise and the 1 photoelectron peak in a pulse height spectrum (I_{1+} and J_{1+}) (figure 3.33). The remaining four channels of the discriminator were used for the muon paddles. These thresholds were set such that each muon paddle produced about 1000 counts per second, which gives a negligible false coincidence rate. Varying this threshold so as to change this counting rate by more than a factor of two had no effect on the muon coincidence rate with the detector insert, indicating that all true coincidences were being included at these threshold levels.

The four muon paddle discriminator outputs were fanned together using an LRS Model 127 FL, which sums its inputs. The analog output of the fan-in was passed to another discriminator (Phillips Scientific Model 708), giving a digital pulse whenever any of the muon paddles triggered its discriminator. This output was passed to a gate and delay (built in-house for a past experiment at NIST) to shape the signal, producing a pulse $1.15 \mu\text{s}$ wide. This signal was used to prevent triggering of the scope by muon events (any PMT pulses in a $1.15 \mu\text{s}$ window starting 8 ns before a muon trigger).

The logic portion of the NIM electronics was handled by three LeCroy 622 Quad Coincidence Modules. The outputs of the first two discriminator channels, $I_{1.5+}$ and $J_{1.5+}$, were ANDed to produce a coincident $(I_{1.5+} \cdot J_{1.5+})^\dagger$ signal. The time window of this coincidence was set by the width of the discriminator outputs, which

[†]using a notation in which ‘ \cdot ’ indicates the logical operation “AND”, ‘ $+$ ’ indicates “OR”, and ‘ \bar{A} ’ indicates “NOT(A)”

were 23 ns each. The width of the PMT pulses was about 10 ns, and the lengths of the BNC cables were chosen such that the pulses arrived at the same time. Most of the coincidences occurred in the first half of this time window, indicating that the coincidence window was sufficiently long to catch all true coincidences. Similarly, the outputs of the last two discriminator channels were ANDed to produce a coincident ($I_{1+} \cdot J_{1+}$) signal. The widths for I_{1+} and J_{1+} were 22 ns and 21 ns, respectively.

The inverted output of the first ($1.5+$) coincidence signal was delayed using a long piece of BNC cable and ORed with the delayed combined muon signal, from the gate and delay, to produce a signal at the inverted output with the form $\overline{((I_{1.5+} \cdot J_{1.5+}) + \mu)} = (I_{1.5+} \cdot J_{1.5+}) \cdot \bar{\mu}$, or a coincidence between $I_{1.5+}$ and $J_{1.5+}$, where the muon paddle hasn't recorded a coincident muon. This signal corresponds to a multi-photon pulse in the cell or lightpipe which was not caused by a cosmic ray (i.e. potentially a “real” event), and was used to trigger the digital oscilloscope. The second coincidence channel, ($I_{1+} \cdot J_{1+}$) was similarly delayed and combined with the delayed muon signal to produce $(I_{1+} \cdot J_{1+}) \cdot \bar{\mu}$. Four more coincidence channels were used to coincide ($I_{1.5+} \cdot J_{1.5+}$) with each of the muon paddle outputs individually.

3.11.2 CAMAC Binner

The “Binner” consisted of a set of three Kinetic Systems Model 3610 CAMAC 50 MHz hex counters used to record the rates of various signals. The signals recorded by each of the eighteen counters, the eighteen “channels” of the binner, are listed in Table 3.1. The first 15 channels were used to count the outputs of various NIM discriminators and coincidence modules, described above. Channel 16 was used to record the discriminated output of a sodium iodide detector (a crystal of sodium iodide

Table 3.1: Signals recorded in each Binner Channel					
1	$I_{1.5+}$	7	not used	13	I_{1+}
2	$J_{1.5+}$	8	μ_4	14	J_{1+}
3	$(I_{1.5+} \cdot J_{1.5+}) \cdot \overline{\mu}$	9	$(I_{1.5+} \cdot J_{1.5+}) \cdot \mu_1$	15	$(I_{1+} \cdot J_{1.5+}) \cdot \overline{\mu}$
4	μ_1	10	$(I_{1.5+} \cdot J_{1.5+}) \cdot \mu_2$	16	Gammas (NaI)
5	μ_2	11	$(I_{1.5+} \cdot J_{1.5+}) \cdot \mu_3$	17	Scope Triggers
6	μ_3	12	$(I_{1.5+} \cdot J_{1.5+}) \cdot \mu_4$	18	1 MHz Oscillator

mounted on the front of a photomultiplier tube) to monitor the gamma background in the guide hall during the experiment. The seventeenth channel was used to record an output from the oscilloscope each time it was triggered. This channel was used in combination with channel three, which counts the number times that the NIM tried to trigger the oscilloscope, to correct for the deadtime of the oscilloscope and also to assign time information to each scope event (section 4.2). The last channel in the binner was used to count the output of a 1 MHz oscillator (a Stanford Research Systems Model DS345 30 MHz function generator), to continuously monitor the bin size. The bin size was set to be 1 second, and in practice the counters were active for 980 ms, after which they were dead for 20 ms during which the total number of counts in each bin was read out by a Macintosh computer. In principle, the data from these counters could be used as the final data for the experiment. However, digitizing the pulses with an oscilloscope allowed us to integrate the area of each PMT pulse. The total area of the pulse corresponds to the total number of electrons arriving at the PMT anode, while the maximum height of the pulse corresponds to the maximum arrival rate during the pulse. Using the integrated charge gives a sharper distinction between one- and two-photoelectron events, increasing the signal-to-background efficiency of

our coincidence requirement. In principle, one could attain the same advantage using a charge-integrating preamplifier, but without the possibility of later varying the parameters of the event selection, or performing more complicated analyses.

3.11.3 Digital Oscilloscope

The final component of the data acquisition was a Tektronix TDS 744A digital oscilloscope, controlled over GPIB by a Macintosh computer (a different one from the one reading out the binner). The first two input channels of the oscilloscope were connected to the amplified outputs of the two photomultiplier tubes. The third, used to trigger the oscilloscope, was connected to the NIM logic signal $(I_{1.5+} \cdot J_{1.5+}) \cdot \overline{\mu}$. The oscilloscope recorded data at 250 megasamples per second for 1 μ s. After 100 trigger events were recorded, the scope was read out by the computer, giving a deadtime of nearly a second between sets of 100 triggers.

This deadtime can be broken down more specifically. The computer records a start time and then signals the oscilloscope to begin acquiring data. There is an initial delay before the first trigger, measured to be about 435 ms. The oscilloscope then continues to record data until its buffer is full (at 100 events). The computer then records a stop time and requests the data. After about 200 ms, the scope begins to transfer the data at about 200 megabytes/second. In all, the readout takes about 440 ms, for a total of 875 ms of deadtime between each set of 100 triggers.

3.12 Computer Control

The same computer that reads out the binner also controls, directly or indirectly, several other components of the experiment. As described in section 4.3, half of

the data runs were taken with the magnetic trap off while the beam was on. In these runs, the magnetic trap was turned on when the beam was turned off, the magnetic trap was turned on. The magnet was controlled by the computer, and the current was read from the magnet current supply and recorded each second with the binner totals. Similarly, the beam was controlled by this computer, giving consistent relative timing of all events. The slider was controlled using a signal from the computer. For safety reasons, a human had to manually press the button opening and closing the neutron beam shutter. The computer indirectly controlled the timing, beeping once a second for ten seconds before the beam shutter was to be opened, and again when it was to be closed. Finally, the computer turned off the high voltage to the photomultiplier tubes while the neutron beam was on, because the neutron capture events produced significant amounts of light in the detector.

Chapter 4

Results

4.1 Overview

The signal of a neutron decaying in the trap is 1) the simultaneous registration of a pulse in each photomultiplier tube corresponding to an event which generates at least a two photoelectron signal, when 2) no cosmic ray event has been simultaneously registered in any of the muon paddles. The first cut eliminates single photon backgrounds, primarily from luminescence; the second cut eliminates muon-induced (multi-photon) events. Removing the muon-induced events also eliminates a very small fraction (less than a percent) of the true signal, when a muon happens to be coincident with a neutron decay.* Suppressing the single-photon backgrounds, however, carries a higher cost, requiring coincident detection above a two photoelectron threshold. Only about a third of the neutron decay events lead to the ionization of two photoelectrons in each photomultiplier tube.

*A few thousand “muon” triggers per second, each producing a deadtime of $1.15\mu\text{s}$, give a total deadtime of a few ms each second, or a few tenths of a percent.




Figure 4.1: A plot of the predicted trapped neutron population over the course of a data run. Three times are indicated: the *start* of the run ($t=0$ s), the times at which the neutron beam is turned *on* ($t=100$ s) and *off* ($t=1450$ s). The run ends at $t=5050$ s.

After applying these constraints, the number of background events is still more than ten times the number of detected trapped neutron decays. To extract evidence of neutron trapping, without the ability to distinguish signal from background on an event-by-event basis, requires measuring the background separately. Thus, half of the data was taken with the magnetic trap initially off, thereby not confining the UCN produced by the neutron beam. This non-trapping data was subtracted from the trapping data. The difference should be the component of the data which results from the magnetic trapping of the UCN.

Data was taken in a series of runs, each divided into three phases (figure 4.1). For the first 100 s, the neutron beam is initially turned off, and therefore there are no neutrons in the trap. This data was used for diagnostic purposes, to detect shifts in the background event rate (such as from neutron activation of materials with long lifetimes in the trapping region). In the loading phase, for the next 22.5 minutes (1350 s, or approximately 1.5 neutron lifetimes), the neutron beam is on, and the trapped

UCN population should slowly increase, asymptotically approaching the theoretical maximum (appendix B.3). During this phase, the photomultipliers were turned off, so no events were recorded. In the last, or observation phase, data is recorded as the trapped neutron population decays for one hour (3600 s, or approximately 4 neutron lifetimes) through beta decay. Data is recorded both in the form of binner totals each second (section 3.11.2) and sets of one hundred oscilloscope events (section 3.11.3). In a trap-on run the magnetic trap is energized throughout the entire run. In a trap-off run the magnet is initially off, but is brought up to full current during the last 50 seconds of the loading phase (making conditions during detection identical between trap-on and trap-off runs). This resulted in a small number of trapped UCN in the trap-off runs, estimated to be at most 2% of that in the trap-on runs. With each run lasting approximately an hour and a half, fifteen runs could be taken in a single day.

To verify that the difference in signal observed between trap-on and trap-off runs is in fact due to trapped UCN decay events, and not from backgrounds, we take trap-on and trap-off data with other parameters altered to eliminate any signal from trapped UCN. Specifically, we took data with the temperature of the helium bath at 1.2 K, rather than at 250 mK (“warm runs”) and data with the helium bath doped with a small amount of ^3He , enough to quickly absorb trapped UCN, but not enough to significantly attenuate the neutron beam (“ ^3He runs”) (section 4.5). In the warm runs and the ^3He runs, trap-on and trap-off data were both taken.

The trapping data was acquired over the course of two reactor cycles. During the first reactor cycle (from 3/23/99 to 4/5/99) a series of trapping runs was taken, followed by a set of warm runs, more trapping runs, and finally a set of ^3He runs. Between reactor cycles, the cell was pumped out, and flushed several times with isotopically pure helium in order to cleanse the cell of residual ^3He from the ^3He

runs. During the next reactor cycle (from 4/16/99 to 5/23/99) a longer series (about 18 days) of trapping runs was taken, followed by a series (about 20 days) of ^3He runs. In the first reactor cycle, the magnet operated at a current of 180 A (giving a trap depth of 0.75 mK). During the second cycle, the magnet consistently quenched well below this current, and we were forced to operate the magnet at a current of 120 A (0.5 mK). We expect the number of trapped neutrons to scale with magnetic field as $B^{3/2}$ (appendix D) or a factor of $\left(\frac{120\text{ A}}{180\text{ A}}\right)^{\frac{3}{2}} \approx 0.5$ relative to the first cycle. The background count rate remained the same. This meant that even with four times as many runs, the signal extracted from the second set of data was only as statistically significant as that from the first cycle.

After throwing out runs in which the magnet quenched during the run, line noise was detected from a neighboring welder or the data files were corrupted, 353 pairs (trap-on and trap-off) of data runs remained, 87 with a trap depth of 0.75 mK and 266 with a trap depth of 0.5 mK. In the first data set, there are 32 pairs of trapping runs, 31 pairs of warm runs, and 24 pairs of ^3He runs. In the second set, there are 128 pairs of trapping runs and 136 pairs of ^3He runs.

4.2 Processing of the Data

As described in section 3.11, the raw data was collected using two parallel systems: digitization of the pulses using a digital oscilloscope (the scope) and storage using an array of counters (the binner). The binner data included the rate of the coincidence between the PMTs with a threshold set at 1.5 photoelectrons (1.5+ p.e.) with muons excluded (the “trigger” signal). Between half and a third of these coincidence events will survive the more stringent cut applied to the scope data that each

PMT have a pulse with integrated area corresponding to at least two photoelectrons.[†] The scope data contains the digitized pulses for each PMT to integrate the area and make a sharper cut to eliminate these background events. The timing data on the oscilloscope, however, is less accurate – we only know the starting and ending times of each set of one hundred trigger events and not the times of the individual events themselves. This is equivalent to binning the data in variable width bins. The variable bin sizes alone do not present an obstacle to fitting, but combining the data from multiple runs in which the bins start and stop at different times presents difficulties in the analysis.

In addition to recording the “trigger” signal, the binner recorded the rate at which the scope actually triggered. In principle, we have timing information, within a one second bin, on each event which was recorded with the scope. The 885 ms of scope deadtime between frames of 100 triggers can be seen in the binner data as one or two adjacent bins in which there were more NIM triggers than scope triggers (Table 4.1). When the scope deadtime occurred overlapping two bins, the assignment of counts to scope frames is straightforward. When the scope deadtime was contained completely in one bin, the few remaining counts were equally divided between the two scope frames.

The binner had a deadtime of 20 ms each second, causing it to miss 2% of the triggers. Thus, instead of counting exactly 100 scope triggers between readouts, the binner typically counted about 98. In the sample shown (Table 4.1), one scope bin of 100 triggers corresponds to scope triggers counted from bins 1593 through 1601, a total of 96. To assign times to each of the 100 triggers in the scope frame, the times

[†]Most of the events which are removed by this cut arise from the luminescence signal, which is strongly time-dependent, and differs between “trap-on” and “trap-off” runs (about 10% larger in the “trap-off” runs).

Table 4.1:

Binner Data

bin	NIM	scope
(s)	(s ⁻¹)	(s ⁻¹)
1589	16	16
1590	7	7
1591	8	8
—1592	13	5
—1593	8	6
1594	17	17
1595	11	11
1596	9	9
1597	11	11
1598	12	12
1599	7	7
1600	17	17
—1601	10	6
—1602	7	5
1603	10	10
1604	14	14
1605	12	12

Sample binner data: In the bins marked with “—” the scope missed triggers while being read out.

of the binner counts were assigned to the scope triggers skipping the appropriate number, evenly spaced among the 100 triggers. Thus, in the example the 96 times recorded by the binner were sequentially assigned to scope triggers 1–19, 21–39, 41–59, 61–79 and 81–100, with the remaining four scope triggers (numbers 20, 40, 60 and 80) assigned the same time as the previous trigger. The relative timing between the scope and binner data (i.e. which 100 scope events corresponds to which group in the

binner data) is set by timing of the initial pulses seen just after the photomultipliers were turned on. During the first few seconds, the trigger rate was several hundred per second[‡], filling the initial scope frame within the same second as the first bin which records a non-zero trigger rate.

Once the scope data is binned into 1 s bins, it is corrected for the deadtime during the scope readout. This deadtime is strongly dependent on the trigger rate, which is decreasing in time. Fortunately, we know exactly in which one second bin the scope was being read. In principle we could use the counts observed in these bins to estimate the number of counts which would have been observed. However, we do not know how much of the 885 ms deadtime fell within each bin. Since it is less likely to introduce a bias, we simply threw out these events, and declared the entire one second bin as deadtime. When we histogram the data into bins of fifty or one hundred seconds, the deadtime is small, and can be corrected by scaling the data by one over the live fraction ($1/(1 - \frac{\# \text{ dead bins}}{\text{total } \# \text{ bins}})$).

During the first twenty or thirty seconds after the beam was turned off, the trigger rate is greater than fifty counts per second, with deadtime in nearly every bin, making this correction inaccurate. The first fifty seconds of the observation phase (a small fraction of the expected neutron decays) of the run are thus thrown out.

All of the data of a given type (trapping, warm or ³He; trap-on or trap-off; 0.75 mK or 0.5 mK) is pooled together in fifty second bins, the first bin is thrown out and the rest are corrected for the scope deadtime. These sets of data were then compared, subtracting the pooled trap-off data from the pooled trap-on data, to look

[‡]This high transient count rate was probably due to a combination of a number of effects: radioactive ²⁰F (created by neutron activation in the Teflon windows) with a lifetime of 16 s, triplet state helium dimers, produced in the helium bath, with a lifetime of 13 s [46], and false coincidences from the initial luminescence rate (passing the threshold either by spillover or afterpulsing).

for a signal indicating the decay of trapped neutrons (section 4.4).

4.3 Backgrounds

The pooled data from the first reactor cycle is shown in figure 4.2. The average of the 32 trap-on runs is shown with the “ \odot ” symbol and the average of the 32 trap-off runs is shown with the “—” symbol (both binned in 200 s bins starting 50 s after the beam was turned off). The difference between the two is shown, with error bars (purely statistical uncertainty). There are several significant features to the “raw” data from the trap-on and trap-off runs. After the neutron beam has been turned off, both signals approach a level of about two counts per second. This component of the background is referred to as the “flat” background, since it is relatively constant over the course of a given run. At earlier times in the run, a time-varying background can be seen as well. Each of these two components of the background arises from different sources.

The flat background in the earliest runs of the cycle was 1.4 s^{-1} , but rose quickly to 1.8 s^{-1} by the second day, and averaged 2 s^{-1} over the entire data set. This shift was due to the slow buildup, from neutron activation, of radioactive nuclei with lifetimes on the order of a day (such as ^{64}Cu , with a half-life of 12.7 hours, in the copper endcaps and cupronickel tube, and ^{24}Na with a half-life of 14.95 hours, a major impurity in acrylic). We estimate that 0.6 s^{-1} resulted from such long-lifetime activation. Cosmic ray events should not vary over the course of a 1.5 hour run, and so unvetted muon events would give a constant background signal. This should account for no more than 0.1 s^{-1} of our background. Similarly, the gamma-ray background in the guide hall was to first order constant in time, so gamma rays penetrating




Figure 4.2: A plot of the trapping signal data from the first reactor cycle, with the average of the pooled trap-on runs and the pooled trap-off runs shown separately. The difference is shown below, with error bars.

the lead shielding around the cryostat would also contribute to this component of the background. No more than 1 s^{-1} could have come from this source, just taking the total background rate before the lead shielding ($\approx 10 \text{ s}^{-1}$) and the fractional area of the lead housing which was uncovered ($\approx 5 - 10\%$), however the gamma background was predominantly from the direction of the reactor and the neutron guide, so the actual background from gammas was probably much less than this. Finally, any long-lived radioactive isotopes (like ^{238}U , ^{232}Th , and ^{40}K all with half-lives longer than 10^9 years) occurring in the materials in the cell, or in the cryostat in general, would give rise to an effectively constant background. In estimating our *possible* background rate from natural radioactivity, it is easy to get numbers of this

magnitude. For example, just 200 g of aluminum 6061 gives about 1.4 decays of ^{232}Th each second[65]. Thus, without further measurements, we cannot rule out the possibility that the majority of our present “flat” background comes from naturally occurring radioactive isotopes in our apparatus.

The time-varying component of the background could include the activation of shorter-lived isotopes, either main constituents of the apparatus such as aluminum ($t_{1/2} = 2.25$ min) (most of the dewar), copper ($t_{1/2} = 5.10$ min) (the cupro-nickel tube and cell endcaps), vanadium ($t_{1/2} = 3.76$ min) (in the magnet form) or trace impurities in or near the detector (appendix E). No one isotope appeared to dominate this component of the background, as it did not fit well to a single exponential. Any remnant of the luminescence signal would also decrease with time, though we believe that this signal is eliminated by our threshold.

4.3.1 Luminescence

Neutron absorption in the BN (and perhaps other materials) creates color centers which relax by emitting light. This results in an initial count rate in excess of 10^5 s^{-1} in each photomultiplier tube ($5 \times 10^4 \text{ s}^{-1}$ at the start of the used data, 50 s into the observation phase) (figure 4.3). This signal varies with time, temperature and magnetic field, and is therefore the most problematic of our backgrounds. Cuts which were required to reduce the luminescence background also eliminated about two-thirds of the events arising from trapped neutrons. In addition, the temperature dependence of the luminescence made the “warm” data, in which the helium was warmed to 1.2 K, potentially suspect. To a lesser extent, the dependence of the luminescence on magnetic field raised possible doubts about the trap-off runs as well.




Figure 4.3: The counting rate with a one-photoelectron threshold. Also shown is the actual coincidence rate of this signal, and the expected rate from the accidental overlap of uncorrelated events. (from a trapping run in the first data set)

Thus, because of the luminescence, data in which the helium was doped with ^3He was required to confirm the presence of trapped neutrons.

Since the luminescence most likely results from the recombination of electrons and holes in the boron nitride, it should emit single photons uncorrelated in time. In principle, neutron decays should result in several photons being detected in coincidence for each decay, providing an easy way to distinguish the neutron decay signal from the luminescence background. In practice, it was more difficult. The rate of the luminescence signal is more than five orders of magnitude greater than the trapping signal. With a single-photon counting rate of around $50,000 \text{ s}^{-1}$ in each PMT, the accidental coincidence rate (with a coincidence window of 43 ns) was still 100 s^{-1} , three orders of magnitude larger than the expected neutron decay signal. The probability

of three single photons arriving independently within the coincidence window is still larger than the expected trapped neutron signal. The luminescence background can be reduced to a tolerable level (assuming that the luminescence photons are always uncorrelated) when coincidence is required between four photons (less than two false coincidences in each 1 h run). Recalling that a neutron decay in the center of the trapping region emitting a 330-keV electron would produce on average a signal of five photoelectrons, a 4 p.e. threshold is a fairly strenuous constraint.

In practice, such a reduction in the luminescence background cannot be obtained by setting a four-photoelectron threshold on a single photomultiplier tube. Even with the high gain dispersion of the Burle 8850 (section 3.9 and figure 3.33), there is not a sharp distinction between a three- and a four-photoelectron event; that is three-photoelectron events overlap the four-photoelectron range of pulse heights (i.e. three-photoelectron events “spill over” into the four-photoelectron region). Even with a higher threshold, a single photoelectron event can result in a much larger “after-pulse” in the PMT[66],[67]. This afterpulse results from contamination of the PMT vacuum over time (especially by helium, which diffuses through glass) and occurs 400-600 ns after the initial pulse. As the photoelectron is accelerated from the photocathode towards the first dynode, it can ionize the helium. The positive helium ion is accelerated towards the photocathode, where it knocks off several additional “photoelectrons”. This creates a secondary pulse that can be larger than the initial pulse. Single photoelectron events can therefore create multiple photoelectron events in a single PMT by means of afterpulsing.

The total effect of afterpulsing and spillover were studied by illuminating a PMT with light from an LED (light-emitting diode) at low current (and thus, producing predominantly uncorrelated single photons), and looking at the pulse-height

distribution. It was found that from two to four percent of the pulses were of at least a size corresponding to two photoelectrons, depending on the level of helium contamination in the PMT vacuum.

To eliminate single-photon backgrounds, two separate photomultiplier tubes in coincidence were used, with a threshold of two photoelectrons on each. This ensures that at least two actual photons arrived in coincidence to cause the signal. Although the rate at which two luminescence photons are detected in coincidence is quite high ($\approx 100 \text{ s}^{-1}$, falling to $\lesssim 20 \text{ s}^{-1}$ by 200 s), the probability of each producing a signal corresponding to two or more photoelectrons, either by afterpulsing or the finite gain dispersion of the PMT, is reasonably low ($\lesssim 0.03 \text{ s}^{-1}$ by 200 s). In the end, to verify that such backgrounds were not producing a false signal, data was taken with ^3He in the trapping region (section 4.5). This requirement also reduces the efficiency for detection of neutron decays as well, giving a final detection efficiency of 31% (section 3.9). In addition, requiring coincidence between two PMTs eliminated the potential background of dark counts in the PMTs.

4.4 Data Analysis

Figure 4.4 shows the pooled background subtracted data, the difference between the pooled trap-on data and the pooled trap-off data, divided by the number of runs which were pooled together. In these plots the data are binned in one hundred second

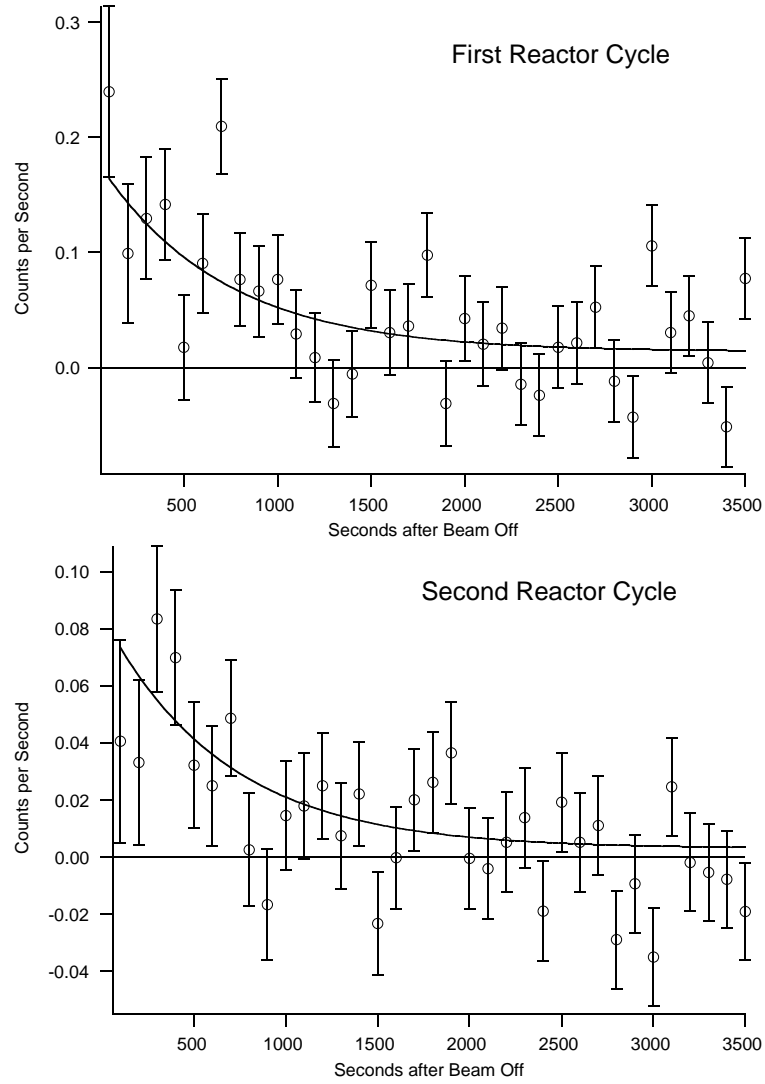


Figure 4.4: Plot of the difference between the averaged “trap-on” and the averaged “trap-off” data for the two reactor cycles. The solid lines indicate the result of the fits described in the text.

bins, beginning fifty seconds after the beam was turned off. The 32 pairs of runs in the first cycle and the 128 pairs of runs in the second cycle were each combined to produce these two data sets. Combining these two sets into a single data set would be less straightforward. One could scale the second cycle’s data to match that of the first cycle using the assumption that the signal scales with magnetic field as $B^{\frac{3}{2}}$, however this assumption is true only if the signal comes from trapped UCN — which is exactly what we intend to show. Instead, the data from the two cycles was simultaneously fit to the equations:

$$\begin{aligned} W_1 &= A_1 e^{-\frac{t}{\tau}} + C_1 \\ W_2 &= A_2 e^{-\frac{t}{\tau}} + C_2 \end{aligned} \tag{4.1}$$

where W is the experimental counting rate and the subscripts correspond to the first and second reactor cycle. The A ’s correspond to the amplitude of the neutron decay signal, τ corresponds to the lifetime of the neutrons in the trap (assumed to be the same in each reactor cycle) and the C ’s correspond to any residual offset between trap-on and trap-off runs, such as might arise from small changes in the “flat” background rate over the course of many runs.

To fit the data from the two cycles to equation 4.1, the data from each cycle was binned into 70 fifty-second bins, with the first starting fifty seconds after the beam was turned off. The five parameters, A_1, A_2, C_1, C_2 , and τ were varied[§] to minimize the χ^2 [¶] of the fit. The most likely values of the fit parameters are those which yield to the minimum χ^2 , and the “1 σ ” (68% probability) range of each parameter is given by the range over which the χ^2 (varying the other four parameters) increases at most

[§]using a routine adapted from the grid fit algorithm described in reference [68]

[¶]Where $\chi^2 = \sum (\frac{x_i^{fit} - x_i^{data}}{\sigma_i})^2$ is a dimensionless measure of the difference between the data and the value predicted by the fit function.




Figure 4.5: Plot of the minimum χ^2 as a function of τ from the fit of the trapping data to equation 4.1.

by one relative to the minimum (varying all five parameters).

The fit gives a minimum χ^2 of 150.67 with 135 degrees of freedom (2 runs \times 70 bins minus 5 fit parameters) at: ^{||}

$$\begin{aligned}
 A_1 &= 0.175 \pm 0.045 \text{ s}^{-1} \\
 A_2 &= 0.082 \pm 0.021 \text{ s}^{-1} \\
 C_1 &= 0.014 \pm 0.011 \text{ s}^{-1} \\
 C_2 &= -0.003 \pm 0.005 \text{ s}^{-1} \\
 \tau &= 660^{+290}_{-170} \text{ s}
 \end{aligned}
 \tag{4.2}$$

^{||}N.B.: The fit values here, and for the ^3He runs differ from those cited in reference [69], because several sets of data, excluded in reference [69] are included in the fits here. The excluded data was from the first six hours of data (four runs) after any long period in which the apparatus was not irradiated. The flat background was changing rapidly between these runs, giving a non-zero C_1 and C_2 in the fits shown here.




Figure 4.6: Plot of the minimum χ^2 of the fit of the data to equation 4.1 as a function of A_1 and A_2 .

The 1σ error bars for each parameter are the values at which the minimum χ^2 accessible by variation of the other four parameters is just one greater than the minimum (i.e. specifically, < 151.67). The χ^2 of this four-parameter fit (with τ held constant), relative to the global minimum, is plotted as a function of τ in figure 4.5. The $\Delta\chi^2$ is less than one for a value of τ between 490 seconds and 950 seconds. Similar graphs for A_1 and A_2 are shown in figure 4.6. C_1 and C_2 (not shown) had roughly parabolic $\Delta\chi^2$'s, like A_1 and A_2 . In equation 4.2, uncertainties are quoted as symmetric if the asymmetry is less than 10%.

4.5 Non-trapping Data

The data shown in figure 4.4 are evidence of an increase in the counting rate in runs in which the magnetic trap is turned on. However, from this evidence alone we cannot conclude that the increase represents an additional source of counts from trapped UCN and not from small, magnet-dependent changes in the rate of neutron-

induced backgrounds. Our assumption has been that the trap-on run data contains both the neutron decay signal and background events, while the trap-off data contains no decay events, but contains the same backgrounds. If this assumption is true, then trap-on data taken under similar conditions, but with a different parameter altered to eliminate the trapped UCN, will result in a data set consistent with the trap off data taken under the same conditions.

To keep the background rate the same in the new configuration, the detector, including the helium bath (as far as its scintillation properties), must remain unchanged. Similarly, the neutron beam, which causes all the time-varying backgrounds, must remain unchanged as well. Two changes which meet these conditions and remove the tarpped neutrons within one second are warming the helium target to a temperature of 1.2 K and doping a small amount of ^3He ($x_3 \approx 10^{-7}$) into the isotopically purified ^4He . At a temperature of 1.2 K, the UCN upscattering rate is $> 0.02 \text{ s}^{-1}$ [41]. Similarly, a small concentration of ^3He will absorb the UCN in less than one second, absorbing less than 1% of the neutron beam passing through the helium (appendix G.3). Since we observed that the luminescence background varied with temperature, as well as with magnetic field, the warm runs could produce a false negative. The ^3He technique did not have this problem, making it a more conclusive negative than the “warm” data.

In reference [41], the upscattering rate for UCN stored in superfluid helium was measured to be about $2.5 \times 10^{-2} \text{ s}^{-1}$ at 1.16 K, giving a storage lifetime of about 40 s. During the loading phase, the continuous upscattering of UCN reduces the expected peak trapped UCN from $P\tau_\beta$ to $P\tau_{up}$, a factor of $\tau_{up}/\tau_\beta \approx 0.05$. With $P = 0.73$ (for the $E_T = 0.8 \text{ mK}$ runs), this gives a maximum of 29 trapped UCN when the beam is turned off, or a beta-decay rate of 0.03 s^{-1} initially, decreasing with a lifetime of

about 40 s, to under 0.01 s^{-1} before the first analyzed data, 50 s after the beam is turned off, compared to an expected rate of (0.16 s^{-1}) at that time in the trapping data.

For the ^3He runs we put 10^{-4} atm-liters of ^3He into the 800 atm-liters of ^4He filling the experimental cell, yielding a concentration $x_3 \equiv n(^4\text{He})/n(^3\text{He}) \approx 10^{-7}$. This concentration gives a UCN lifetime of 0.4 s and an absorption length of nearly 400 m for 4 Å neutrons (appendix G.3), the lowest wavelength passing the beryllium filter (section 3.2), and a factor of two shorter at 8 Å on the tail of the spectrum. Over the 1 m length of helium before the beam dump, the ^3He should absorb less than 1% of the beam. Thus any background arising from the cold neutrons in the beam, rather than from UCN remaining after the beam is turned off, should be affected by less than a percent.

The data from the warm runs is shown in figure 4.7. There was no warm data from the second reactor cycle, so this data was fit only to A_1 and C_1 , with $\tau = 660 \text{ s}$ held fixed. This two parameter fit gives a result of:

$$A_1 = 0.042 \pm 0.043 \text{ s}^{-1}$$

$$C_1 = 0.001 \pm 0.009 \text{ s}^{-1},$$

with a χ^2 of 76.1 for 68 degrees of freedom. This is consistent with the prediction of $A_1 \lesssim 0.01 \text{ s}^{-1}$.




Figure 4.7: Plots of the difference between the average pooled magnetic “trap-on” and “trap-off” in the 1.2 K “warm” data.

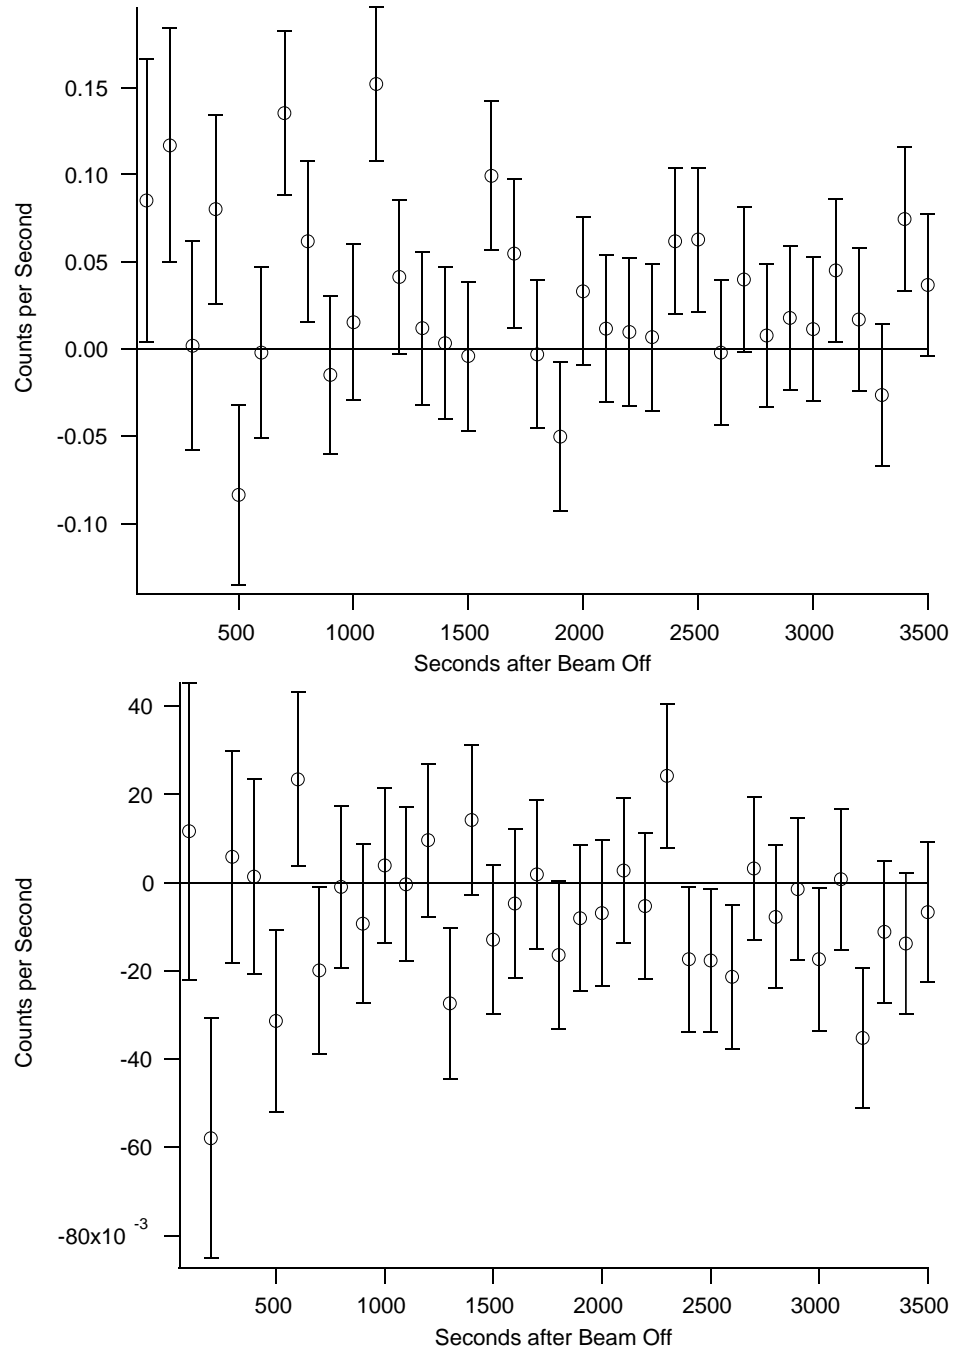


Figure 4.8: Plots of the difference between the average pooled magnetic “trap-on” and “trap-off” in the ^3He data for each reactor cycle.

The data from the helium runs is shown in figure 4.8. This data was fit to equation 4.1 with τ held fixed at 660 seconds (the fit value from the trapping data). This four parameter fit gives the following results, with a minimum χ^2 of 119.3 for 136 degrees of freedom:

$$\begin{aligned} A_1 &= 0.045 \pm 0.041 \text{ s}^{-1} \\ A_2 &= -0.001 \pm 0.017 \text{ s}^{-1} \\ C_1 &= 0.023 \pm 0.009 \text{ s}^{-1} \\ C_2 &= -0.007 \pm 0.0035 \text{ s}^{-1}. \end{aligned}$$

The amplitudes from both cycles are roughly consistent with no trapped UCN in the trapping region, and definitely inconsistent with the signal observed in the trapping data. The remaining constant offsets arise from the slow change in the “flat” background rate from run to run.

4.6 Discussion

For an infinite loading time, the initial number of trapped neutrons is given by $P\tau_n$, where P is the production rate of trapped UCN and τ_n is the storage time. For a finite loading time, the initial number of trapped UCN, N , is given by (appendix B.3):

$$N = P\tau_n(1 - \exp(-t_{\text{load}}/\tau_n)). \quad (4.3)$$

The production rate of trapped UCN, P , for the 0.75 mK runs, was calculated to be $0.73 \pm 0.24 \text{ s}^{-1}$ (appendix D). If we assume that beta decay is the only significant loss, then $\tau_n = \tau_\beta = 885.8 \text{ s}$ [21]. A loading time of $t_{\text{load}} = 22.5 \text{ min}$ then gives

$N_1 = 500 \pm 170$. In the 0.5 mK runs the production rate, P , and hence N_0 , is scaled by a factor of $\left(\frac{120}{180}\right)^{3/2} = 0.544$, giving $N_2 = 270 \pm 90$.

The initial count rate in our detector is given by $W = \frac{\epsilon N}{\tau_\beta}$, where $\epsilon = 31\%$ is the detector efficiency (section 3.9). We can thus calculate the number of trapped neutrons observed in each data set by taking $\hat{N}_i = \tau_\beta A_i / \epsilon$. For the first data set, $A_1 = 0.175 \pm 0.045$ gives $\hat{N}_1 = 500 \pm 155$ (predicted $N_1 = 500 \pm 170$). For the second data set, $A_2 = 0.082 \pm 0.021$ gives $\hat{N}_2 = 230 \pm 70$ (predicted $N_2 = 270 \pm 90$). The fit to each data set is thus consistent with the expected number of trapped UCN. Combined with the vanishing of the signal under conditions inimical to UCN, this constitutes strong evidence that we have magnetically trapped neutrons.

Chapter 5

Conclusions and Future Directions

5.1 Lifetime Measurement

5.1.1 Prospects for Improvement

Having demonstrated the magnetic trapping of UCN in a bath of liquid helium, we can now consider the prospects for making a measurement of the neutron lifetime in such an apparatus. At first, the size of the error bars on the measurement quoted in equation 4.2 ($\tau_n = 660^{+290}_{-170}$ s) may seem discouraging, compared to a present uncertainty of ± 0.9 s. However, the uncertainty in the data (i.e. the error bars in figure 4.4) is entirely due to statistical fluctuations in the large backgrounds being subtracted, not in the data itself. With no background to contend with, four weeks of data like that in the first data set would give error bars of ± 20 s. In addition to this direct improvement, a substantial reduction in the single-photon backgrounds could easily increase the detection efficiency (presently 31%, section 3.9.1) by a factor of 2 or more, by allowing the thresholds to be relaxed.

There is every reason to believe that such a drastic reduction in at least the time-varying background is possible, since these backgrounds are created by neutrons of all wavelengths, while the trapped UCN are produced only by neutrons in a narrow band around 8.9 Å. Thus, a monochromatic beam around 8.9 Å could reduce these backgrounds by as much as two orders of magnitude relative to the neutron signal.

Another possible area of improvement is the flux of the cold neutron beam used to produce the UCN. A higher 8.9 Å flux clearly translates directly to a higher UCN production rate. The NIST hydrogen cold source will be upgraded in late 2000/early 2001, giving a factor of 1.8–2 increase in 8.9 Å flux[1]. As well, the PF1 cold neutron beam at the Institut Laue Langevin quotes a capture flux of 1×10^{10} n/cm²/s at the end of a ⁵⁸Ni-coated guide (to be compared to roughly $2 - 3 \times 10^9$ n/cm²/s at NG-6, with an empty filter cryostat)[70], possibly giving a factor of 3–5 higher 8.9 Å flux*.

Finally, the number of trapped UCN could be increased by building a magnetic trap either physically larger than that described here, “deeper” (in terms of magnetic trap depth), or both. From the calculation of the production rate (equation D.8), the number of trapped UCN produced is approximately proportional to the volume of the trap (lr_T^2) and to $k_TE_T \propto B_T^{\frac{3}{2}}$. In fact, collimation constraints (the dependence of r_c on l and r_T , section 3.2.3) make the UCN population depend slightly more strongly on r_T and less so on l . An increased trap radius is also likely to increase the trap depth. The present magnet has a bore of 5.08 cm and the diameter of the trap is just over 3 cm. In a larger-bore magnet, this ratio should improve, giving a trap depth that is a larger fraction of the field at the bore of the magnet. An improvement of more than two orders of magnitude in the number of trapped UCN, giving a potential

*The actual improvement possible at the ILL could be as much as a factor of two or three different from this, depending on the actual beam spectrum and divergence.

measurement at the 1–2 s level, seems likely in the next several years[71].

5.1.2 Limits to the Method

Ultimately, this method of measuring the neutron lifetime is limited by losses of neutrons from the trap, which would give the appearance of a lifetime shorter than the actual beta-decay lifetime. It has also been suggested that the beta-decay lifetime itself might be altered by the presence of the helium bath. Several possible mechanisms of such a shift are discussed and estimated in appendix H. Ways in which neutrons could be lost from the trap, other than beta-decay, include: scattering by excitations in the liquid helium, absorption by ^3He , Majorana spin-flip transitions and the marginal trapping of neutrons. Scattering of trapped neutrons by excitations in the liquid helium can be reduced to below 10^{-5} of the beta-decay rate by cooling the helium to about 150 mK (section 2.2, [41, 44]). Absorption by ^3He can be reduced by using isotopically purified ^4He (section 2.5.1). At the level to which our supply of ^4He has been verified, we can only set a limit of 10^{-2} of the beta-decay rate, however there is good reason to believe that the purity is far better than this[51]. This purity will need to be measured independently to verify a less than 10^{-2} effect on the lifetime. Several methods of isotopic analysis, including accelerator mass spectrometry[72] and isotopic detection in magneto-optical traps[73], seem likely to be able to measure the isotopic purity of our helium orders of magnitude better than has been done in the past. Majorana spin-flip transitions occur near zero-field regions, and can be suppressed in a Ioffe trap by the application of a bias field (section 2.1, appendix G.2). Marginally trapped neutrons can be eliminated by ramping the magnetic field down to 30% of its maximum value and then back up, resulting in a population of

UCN all with energies less than the trap depth and thus all truly trapped (section 2.5.2, appendix G.1). There is no reason to expect any of these effects to prevent a measurement of the neutron lifetime at the 10^{-5} level, given sufficient statistics.

In order to extract $|V_{ud}|^2$ from a measurement of the neutron lifetime, several corrections need to be applied (appendix A). While $f(1 - \delta^R)$ has been calculated at the 10^{-5} level, $(1 + \Delta_R^V)$ is known only at the 8×10^{-4} level[11, 74], limiting the precision with which $|V_{ud}|^2$ can be determined from g_v^2 . However, this $(1 + \Delta_R^V)$ term affects *all* attempts to determine $|V_{ud}|^2$ by measuring g_v^2 , including $0^+ \rightarrow 0^+$ decays and the beta decay of the pion. In addition, even at a level of 8×10^{-4} , a measurement of g_v using just the neutron would solidly answer the question of CKM non-unitarity raised by measurements of the $0^+ \rightarrow 0^+$ decay rates.

5.2 Other Applications

In addition to their immediate use in a measurement of the neutron lifetime, the techniques developed for this experiment are applicable to a wide range of future experiments. Most generally, the experiment described here, even without an upgrade, demonstrated the production of polarized UCN at densities comparable to those available for experiments anywhere, with the advantage of operating from a standard cold neutron beam, rather than requiring a dedicated UCN source.

In addition to its use in measurements with neutrons, the detection of EUV scintillations in liquid helium (or possibly neon) has been proposed for use in a solar neutrino experiment[31]. In such an experiment, the downconverted scintillation light from neutrino scattering events in a bath of helium could be detected from many sides, rather than at the end of a long, narrow magnet bore, resulting in a much

higher signal than in our experiment. Since the number of ionized helium atoms, and thus the EUV signal, would be proportional to the kinetic energy of the scattered electron, the energy spectrum of the detected neutrinos could potentially be extracted at energies much lower than those detectable using Čerenkov light.

Both the storage of UCN in a helium bath and the detection techniques described herein are relevant to a proposed experiment to measure the electric dipole moment (EDM) of the neutron using polarized UCN in a bath of liquid helium[30]. As described in section 2.1, the magnetic moment of the neutron causes it to precess around a magnetic field. Similarly, an electric dipole moment would cause the neutron to precess around an electric field as well, with a total frequency of $\omega = \frac{2\mu B + 2d_e E}{\hbar}$, where d_e is the EDM. In this experiment, the neutron spin precesses around a magnetic field, with an electric field oriented alternately parallel and anti-parallel to the magnetic field, inducing a shift in the precession frequency proportional to the electric dipole moment. The precession rate would be measured using nuclear-spin-polarized ^3He , also precessing around the magnetic field and absorbing neutrons at a rate strongly dependent on the relative spin of the neutron and ^3He nucleus. The absorption rate, oscillating as the spins of the neutrons and ^3He precess, would be measured by detecting the scintillations in the helium bath of the emitted protons in the capture reaction $^3\text{He} + n \rightarrow ^3\text{H} + p$, just as the neutron decays in our experiment are observed by detecting the decay electrons. The detection of neutron capture by ^3He using this technique was explicitly demonstrated in reference [47].

Appendix A

Calculation of Neutron Beta Decay Rate

(N.B. Throughout Appendix A, unless otherwise noted, units are chosen such that $\hbar = c = 1$.)

Fermi's Golden Rule tells us that any transition rate is given by:

$$\Gamma_{i \rightarrow f} = k \int |V_{if}|^2 d\Phi_n. \quad (\text{A.1})$$

For the decay of the neutron in the neutron's center-of-mass frame, the constant, k , is given by $(2\pi)^4/(2m_n)$, the phase space factor, $d\Phi_n$ is given by:

$$\delta(m_n - E_p - E_e - E_\nu) \delta^3(\vec{p}_p + \vec{e}_p + \vec{p}_\nu) \frac{d^3 p_p}{(2\pi)^3 2E_p} \frac{d^3 p_e}{(2\pi)^3 2E_e} \frac{d^3 p_\nu}{(2\pi)^3 2E_\nu}, \quad (\text{A.2})$$

and the matrix element V_{if} can be separated into terms for each of the two vertices in figure 1.1 and the “propagator” term for the exchanged W-boson, using the standard Feynman rules.

The full propagator term is given by $-i \frac{g_{\mu\nu} - q_\mu q_\nu / M_W^2}{q^2 - M_W^2}$, where M_W and q are the mass and the four-momentum of the W-boson respectively, $g_{\mu\nu} = \frac{1}{2} \gamma_\mu \gamma_\nu$ with

γ_μ ($\mu = 0..3$) the gamma matrices, the spinor equivalents of the Pauli spin matrices. Since the energy scale of the decay (≈ 1 MeV) is so much less than $M_W = 80.4$ GeV, the propagator term can be reduced to $ig_{\mu\nu}/M_W^2$, giving an effective four-fermion interaction with a strength $\frac{G_F}{\sqrt{2}} = \frac{g^2}{8M_W^2}$, where the two factors of g come from the two vertex factors.

The standard 2-fermion weak charged current is given by $\frac{-ig}{2\sqrt{2}}\bar{\psi}_1\gamma^\mu(c_v + c_a\gamma_5)\psi_2$, where g is the weak force coupling constant, ψ_1 and ψ_2 are the spinor wavefunctions of the two fermions, and $\gamma_5 = i\gamma^0\gamma^1\gamma^2\gamma^3$. The constants c_v and c_a correspond to the relative weight of the vector and axial-vector currents respectively. For the leptonic current (the electron and the anti-neutrino), the weak interaction is pure “V-A”, that is $c_v = +1$ and $c_a = -1$. For the hadronic current, the strong interaction between the decaying d-quark and the rest of the neutron can perturb these factors. By the conserved vector current hypothesis, $c_v = +1$ despite this interaction (at least in the limit $q \ll M_W$), but the axial current is only partially conserved and c_a is not exactly -1. In addition, the weak interaction couples the vector of $q=+2/3$ quarks (u c t) to a slightly rotated vector of the $q=-1/3$ quarks (d' s' b') given by a 3x3 matrix (the Cabibbo-Kobayashi-Maskawa or CKM mixing matrix) acting on the vector (d s b). This has the effect of weakening the coupling between the up and down quarks by a factor of “ V_{ud} ” (the first element of the CKM matrix), which has a value between 0.9745–0.9760 (90% confidence limits). This gives a hadronic current of $\frac{igV_{ud}}{2\sqrt{2}}\bar{\psi}_p\gamma^\mu(c_v + c_a\gamma_5)\psi_n$. Traditionally, c_v and c_a are subsumed into constants g_v and g_a , where $g_v = G_F V_{ud} c_v$, and similarly for g_a .

Defining p , n , u , and v as the spinor wavefunctions for the proton, neutron, electron and anti-neutrino, respectively (i.e. dropping the ψ 's), the matrix element is

given by:

$$V_{n \rightarrow pe\nu} = \frac{-igV_{ud}}{2\sqrt{2}}(\bar{p}\gamma^\mu(c_v + c_a\gamma_5)n)i\frac{g_{\mu\nu}}{M_W^2}\frac{-ig}{2\sqrt{2}}(\bar{u}\gamma^\nu(1 - \gamma_5)v) \quad (\text{A.3})$$

where a sum over the indices μ and ν is implied, $g_{\mu\nu} = \frac{1}{2}\gamma_\mu\gamma_\nu$, and $q = p_i - p_f$. Pulling all the constants out front, we get:

$$V_{n \rightarrow pe\nu} = -i\frac{g^2}{M_W^2}\frac{V_{ud}}{8}(\bar{p}\gamma^\mu(c_v + c_a\gamma_5)n)g_{\mu\nu}(\bar{u}\gamma^\nu(1 - \gamma_5)v) \quad (\text{A.4})$$

or

$$V_{n \rightarrow pe\nu} = -i\frac{g_v}{\sqrt{2}}(\bar{p}\gamma^\mu(1 + \lambda\gamma_5)n)(\bar{u}\gamma_\mu(1 - \gamma_5)v), \quad (\text{A.5})$$

where $\lambda \equiv \frac{c_a}{c_v} = \frac{g_a}{g_v}$, and $g_v = c_v G_F V_{ud} = 8\sqrt{2}c_v \frac{g^2}{M_W^2} V_{ud}$.

Multiplying by the complex conjugate, summing over final spin states and averaging over initial spin states, this becomes:

$$|V_{n \rightarrow pe\nu}|^2 = \frac{1}{2} \sum_{s_n, s_p, s_e, s_\nu} \frac{g_v^2}{2} (\bar{p}\gamma^\mu(1 + \lambda\gamma_5)n)(\bar{u}\gamma_\mu(1 - \gamma_5)v)(\bar{v}\gamma_\mu(1 + \gamma_5)u)(\bar{n}\gamma^\mu(1 - \lambda\gamma_5)p). \quad (\text{A.6})$$

Performing the sums and using ‘‘Casimir’s Trick’’ to turn the spin sums into traces, this becomes:

$$|V_{n \rightarrow pe\nu}|^2 = \frac{g_v^2}{4} \text{Tr}((e_a\gamma^a + m_e)\gamma_\mu(1 - \gamma_5)k_b\gamma^b\gamma_\nu(1 - \gamma_5)) \quad (\text{A.7})$$

$$\begin{aligned} & \text{Tr}((p_c\gamma^c + m_p)\gamma^\mu(1 + \lambda\gamma_5)(n_d\gamma^d + m_n)\gamma^\nu(1 + \lambda\gamma_5)) \\ &= \frac{g_v^2}{2} \text{Tr}(e_a\gamma^a\gamma_\mu k_b\gamma^b\gamma_\nu(1 - \gamma_5)) \\ & \quad (\text{Tr}(p_c\gamma^c\gamma^\mu n_d\gamma^d\gamma^\nu(1 + \lambda^2 + 2\lambda\gamma_5)) + m_n m_p (1 - \lambda^2 \text{Tr}(\gamma^{mu}\gamma^\nu))) \end{aligned} \quad (\text{A.8})$$

where p , n , e and k are the four-momenta of the proton, neutron, electron and neutrino, respectively.

Taking the traces gives:

$$|V_{n \rightarrow pe\nu}|^2 = 8g_v^2 (e_\mu k_\nu + e_\nu k_\mu - (e \cdot k)g_{\mu\nu} - ie^a k^b \epsilon_{a\mu k\nu}) \quad (\text{A.9})$$

$$((1 + \lambda^2)(p^{mu} n^\nu + p^\nu n^\mu - (p \cdot n)g^{\mu\nu}) + 2i\lambda p_c n_d \epsilon^{c\mu d\nu} + (1 - \lambda^2)m_n m_p g^{\mu\nu})$$

Finally, multiplying out and grouping terms gives a final result for the square of the matrix element:

$$|V_{n \rightarrow pe\nu}|^2 = 16g_v^2 m_n ((E_e(k \cdot p)(1 + \lambda)^2 + E_\nu(e \cdot p)(1 - \lambda)^2 + m_p(e \cdot k)(\lambda^2 - 1)) \quad (\text{A.10})$$

Performing the integration over the phase space factor (equation A.2), we can use the delta-function representing momentum conservation to completely integrate out the proton momentum and convert $d^3e = p_e^2 dp_e d\Omega_e$, and $p_e^2 dp_e / E_e = p_e dE_e$, giving:

$$d\Gamma_{n \rightarrow pe\bar{\nu}_e} = \frac{(2\pi)^4}{2m_n} |V_{n \rightarrow pe\nu}|^2 \frac{\delta(m_n - E_p - E_e - E_\nu)}{8(2\pi)^9 E_p} p_e dE_e d\Omega_e p_\nu dE_\nu d\Omega_\nu, \quad (\text{A.11})$$

and then use the delta-function representing energy conservation to integrate over E_ν giving:

$$d\Gamma_{n \rightarrow pe\bar{\nu}_e} = \frac{|V_{n \rightarrow pe\nu}|^2}{16(2\pi)^5 m_n E_p} p_e p_\nu dE_e d\Omega_e d\Omega_\nu. \quad (\text{A.12})$$

Finally, we can plug in the squared matrix element from equation A.10 giving:

$$d\Gamma_{n \rightarrow pe\bar{\nu}_e} = \frac{g_v^2}{(2\pi)^5 E_p} \quad (\text{A.13})$$

$$((E_e(k \cdot p)(1 + \lambda)^2 + E_\nu(e \cdot p)(1 - \lambda)^2 + m_p(e \cdot k)(\lambda^2 - 1))$$

$$p_e p_\nu dE_e d\Omega_e d\Omega_\nu.$$

Now, we can plug in $-\vec{p}_p = \vec{p}_e + \vec{p}_\nu$, and expand out each dot product giving:

$$\begin{aligned} d\Gamma_{n \rightarrow pe\bar{\nu}_e} &= \frac{g_v^2 p_e p_\nu}{(2\pi)^5 E_p} \left(E_e(E_\nu E_p + p_\nu^2 + \vec{p}_e \cdot \vec{p}_\nu)(1 + \lambda)^2 + \right. \\ &\quad E_\nu(E_e E_p + p_e^2 + \vec{p}_e \cdot \vec{p}_\nu)(1 - \lambda)^2 + \\ &\quad \left. m_p(E_e E_\nu - \vec{p}_e \cdot \vec{p}_\nu)(\lambda^2 - 1) \right) dE_e d\Omega_e d\Omega_\nu. \end{aligned} \quad (\text{A.14})$$

Regrouping, and recalling that $E_\nu = p_\nu$, this gives:

$$\begin{aligned} d\Gamma_{n \rightarrow pe\bar{\nu}_e} &= \frac{g_v^2 p_e E_e E_\nu^2}{(2\pi)^5} \left((1 + 3\lambda^2) + \frac{E_\nu}{E_p}(1 + \lambda)^2 + \frac{E_e}{E_p}(1 - \lambda)^2 \right. \\ &\quad \left. - \frac{E_p - m_p}{E_p}(\lambda^2 - 1) + \frac{(\vec{p}_e \cdot \vec{p}_\nu)}{E_e E_\nu} \left(\frac{m_p}{E_p}(1 - \lambda^2) + \frac{E_e}{E_p}(1 + \lambda)^2 + \frac{E_\nu}{E_p}(1 - \lambda)^2 \right) \right) dE_e d\Omega_e d\Omega_\nu. \end{aligned} \quad (\text{A.15})$$

Finally, recalling that $m_n - m_p = E_e + E_\nu + (E_p - m_p) = 1.2936 \text{ MeV}/c^2$ and $E_p^{min} = m_p = 938.3 \text{ MeV}/c^2$, we know that:

$$\frac{E_e}{E_p} \lesssim 1.4 \times 10^{-3} \quad (\text{A.16})$$

$$\frac{E_\nu}{E_p} \lesssim 1 \times 10^{-3} \quad (\text{A.17})$$

$$\frac{E_p - m_p}{E_p} \lesssim 1 \times 10^{-3}. \quad (\text{A.18})$$

(And actually, $E_p - m_p$ has a maximum of less than 1 keV/ c^2 , making this term $\lesssim 1 \times 10^{-6}$.) The remaining energy dependence is traditionally absorbed into a function:

$$\begin{aligned} F(E_e) &\equiv p_e E_e E_\nu^2 \\ &\approx E_e \sqrt{E_e^2 - m_e^2} (E_{max} - E_e)^2, \end{aligned} \quad (\text{A.19})$$

where $E_{max} \approx (m_n^2 - m_p^2 + m_e^2)/2m_n = 1.293 \text{ MeV}/c^2$, or the more familiar kinetic energy $E_{max} - m_e = 782 \text{ keV}/c^2$. (figure 2.5) Thus, we can write:

$$\frac{d\Gamma_{n \rightarrow pe\bar{\nu}_e}}{dE_e d\Omega_e d\Omega_\nu} \approx \frac{g_v^2 F(E_e)}{(2\pi)^5} (1 + 3\lambda^2) \left(1 + a \frac{(\vec{p}_e \cdot \vec{p}_\nu)}{E_e E_\nu} \right), \quad (\text{A.20})$$

where $a \equiv \frac{1-\lambda^2}{1+3\lambda^2}$.

If one carries through σ_n , rather than averaging over the neutron spin, two additional significant terms arise, giving the standard:

$$\frac{d\Gamma_{n \rightarrow pe\bar{\nu}_e}}{dE_e d\Omega_e d\Omega_\nu} \approx \frac{g_v^2 F(E_e)}{(2\pi)^5} (1 + 3\lambda^2) \left(1 + a \frac{(\vec{p}_e \cdot \vec{p}_\nu)}{E_e E_\nu} + A \frac{(\vec{\sigma}_n \cdot \vec{p}_e)}{E_e} + B \frac{(\vec{\sigma}_n \cdot \vec{p}_\nu)}{E_\nu} \right), \quad (\text{A.21})$$

where $A \equiv -2\frac{\lambda^2+\lambda}{1+3\lambda^2}$ and $B \equiv 2\frac{\lambda^2-\lambda}{1+3\lambda^2}$.

Finally, one can integrate equation A.20 to give the total decay rate. Since we averaged over the neutron spin, we still haven't selected an orientation, and the $\int d\Omega_\nu$ just gives a trivial factor of 4π . We can then convert $d\Omega_e = \sin\theta_{e\nu} d\theta_{e\nu} d\phi_e$, and the $\int d\phi_e$ gives a factor of 2π , leaving:

$$\frac{d\Gamma_{n \rightarrow pe\bar{\nu}_e}}{dE_e d\theta_{e\nu}} \approx \frac{2(g_v^2 + 3g_a^2)F(E_e)}{(2\pi)^3} (\sin\theta_{e\nu} + a \frac{|p_e|}{E_e} \cos\theta_{e\nu} \sin\theta_{e\nu}) \quad (\text{A.22})$$

Taking the integral over $d\theta_{e\nu}$ of the a term gives $\int_0^\pi d\theta_{e\nu} \cos\theta_{e\nu} \sin\theta_{e\nu} = 0$. For the other term, $\int_0^\pi d\theta_{e\nu} \sin\theta_{e\nu} = 2$. This leaves the integral over dE_e :

$$\begin{aligned} \frac{d\Gamma_{n \rightarrow pe\bar{\nu}_e}}{dE_e} &= \frac{4(g_v^2 + 3g_a^2)F(E_e)}{(2\pi)^3} \\ &= \frac{4(g_v^2 + 3g_a^2)}{(2\pi)^3} p_e E_e E_\nu^2 \end{aligned} \quad (\text{A.23})$$

$$\begin{aligned} \Gamma_{n \rightarrow pe\bar{\nu}_e} &= \frac{4(g_v^2 + 3g_a^2)}{(2\pi)^3} \int_{m_e}^{E_{max}} dE_e p_e E_e (E_{max} - E_e)^2 \\ &= \frac{4(g_v^2 + 3g_a^2)}{(2\pi)^3} \int_0^{p_{max}} dp_e p_e^2 (E_{max}^2 - 2E_{max}E_e + m_e^2 + p_e^2) \\ &= \frac{4(g_v^2 + 3g_a^2)}{(2\pi)^3} (0.0569 \text{MeV}^5) \\ &= \frac{4|V_{ud}|^2 G_F^2 (1 + 3\lambda^2)}{(2\pi)^3} (0.0569 \text{MeV}^5) \end{aligned} \quad (\text{A.24})$$

Finally, plugging in $G_F = 1.1664 \times 10^{-11} \text{MeV}^{-2}$ gives:

$$\Gamma_{n \rightarrow pe\bar{\nu}_e} = |V_{ud}|^2 (1 + 3\lambda^2) 1.25 \times 10^{-25} \text{MeV}, \quad (\text{A.25})$$

and dividing this width by \hbar to get a rate gives:

$$\Gamma_{n \rightarrow pe\bar{\nu}_e} = |V_{ud}|^2(1 + 3\lambda^2)1.89 \times 10^{-4} s^{-1}. \quad (\text{A.26})$$

Plugging in $|V_{ud}| = 0.975$ and $\lambda = -1.26$ gives a final rate of $\Gamma_{n \rightarrow pe\bar{\nu}_e} = 1.04 \times 10^{-5} s^{-1}$ or $\tau_n = 960$ seconds.

The various corrections to this are usually subsumed into a corrected value for $\int F(E_e)$, given as $f^R(1 + \delta_R)(1 + \Delta_R^V)$, where δ_R is the nucleus-dependent radiative correction for the neutron and Δ_R^V is the nucleus-independent radiative correction. f is usually given in terms of $(m_e/c^2)^5$, in order to render it unitless. For comparison, $\int F(E_e) = 0.0569 \text{ MeV}^5$ from equation A.24 gives $f = 1.63$. $f^R(1 + \delta_R)$ has been evaluated more rigorously to be 1.71489 ± 0.00002 , and Δ_R^V to be $2.40 \pm 0.08 \times 10^{-2}$ [11, 74]. Making these corrections, we get a neutron lifetime of $\tau_n = 893$ s, where the leading uncertainties now arise from $|V_{ud}|$ and λ . Taking $|V_{ud}| = 0.973$ (the 1σ lower limit from $0^+ \rightarrow 0^+$ decays) gives $\tau_n = 889.1$ s, whereas $|V_{ud}| = 0.976$ (the 1σ upper limit from CKM unitarity) gives $\tau_n = 894.6$ s. Similarly, using $\lambda = -1.274$ [13] gives $\tau_n = 909$ s, while $\lambda = -1.2594$ [12] gives $\tau_n = 892$ s. Constraints on $|V_{ud}|$ and λ from various measurements are shown in figure 1.3 with reference to the neutron lifetime.

Appendix B

Miscellaneous Math

B.1 Magnetic Field Maxima

It is here demonstrated that Maxwell's Equations preclude the creation of a local maximum in the magnitude of a magnetic field in free space (i.e. $\vec{J} = 0$) and in a steady-state (i.e. $d\vec{E}/dt = 0$). Under such conditions, Maxwell's Equations for the magnetic field simplify to:

$$\begin{aligned}\vec{\nabla} \cdot \vec{B} &= 0 \\ \text{and } \vec{\nabla} \times \vec{B} &= \vec{0}.\end{aligned}\tag{B.1}$$

We can write these in summation notation, where ∂_i represents $\frac{d}{dx_i}$, B_i represents the i th component of \vec{B} (i.e. $B_i = \hat{x}_i \cdot \vec{B}$), ϵ_{ijk} is the anti-symmetric tensor (defined by $\epsilon_{123} = 1$ and $\epsilon_{ijk} = -\epsilon_{jik} = -\epsilon_{kji}$) and a sum over all repeated indices is implied as:

$$\begin{aligned}\partial_i B_i &= 0 \\ \partial_i B_j \epsilon_{ijk} \hat{x}_k &= 0.\end{aligned}\tag{B.2}$$

Note that the second of these can be expanded to three equations of the form $\partial_i B_j - \partial_j B_i = 0$, or $\partial_i B_j = \partial_j B_i$, which will be used below.

A local extremum or saddle point of a function (Φ) exists wherever $\frac{d}{dx_i} = 0$ for all three axes. A local maximum exists if in addition, $\frac{d^2}{dx_i^2} \Phi \leq 0$ for all three axes. A local maximum in the magnitude of the magnetic field, $|\vec{B}| = \sqrt{B_x^2 + B_y^2 + B_z^2}$ implies a local maximum in the square of the magnitude of the magnetic field $B_x^2 + B_y^2 + B_z^2$. Thus, at a local maximum of the magnitude of the magnetic field, $\nabla^2(B^2) \leq 0$.

However, expanding $\nabla^2(B^2)$ gives:

$$\begin{aligned}
 \nabla^2(B^2) &= \partial_i(\partial_i(B_j B_j)) \\
 &= \partial_i(2B_j \partial_i B_j) \\
 &= 2(\partial_i B_j)(\partial_i B_j) + 2B_j(\partial_i \partial_i B_j) \\
 &= 2(\partial_i B_j)^2 + 2B_j(\partial_i \partial_j B_i) \\
 &= 2(\partial_i B_j)^2 + 2B_j(\partial_j \partial_i B_i) \\
 &= 2(\partial_i B_j)^2 + 2B_j \partial_j(\partial_i B_i)
 \end{aligned} \tag{B.3}$$

The first term in equation B.3 is explicitly non-negative and the second $(2\vec{B} \cdot \vec{\nabla}(\vec{\nabla} \cdot \vec{B}))$ is zero, since $\vec{\nabla} \cdot \vec{B} = 0$. Thus, $\nabla^2(B^2)$ is non-negative everywhere for a static magnetic field in free space. Further, the trivial case $\nabla^2(B^2) = 0$ everywhere in a region can only corresponds to a constant \vec{B} , which obviously lacks a local maximum. Hence, no static magnetic field maximum can be created in free space.

B.2 Kinematics of 1-phonon Scattering

If the incident neutron has an initial momentum \vec{p}' and a final momentum of \vec{p} , and the phonon has a momentum \vec{q} , then momentum conservation requires:

$$\vec{q} + \vec{p} = \vec{p}' \longrightarrow (p')^2 = p^2 + q^2 + 2pq\cos\theta \quad (\text{B.4})$$

and energy conservation requires:

$$cq + \frac{p^2}{2m} = \frac{(p')^2}{2m}. \quad (\text{B.5})$$

Combining the two, we get:

$$p^2 + q^2 + 2pq\cos\theta = p^2 + 2mcq \quad (\text{B.6})$$

$$q(q + 2p\cos\theta - 2mc) = 0 \quad (\text{B.7})$$

Thus, either $q = 0$, in which case no momentum is transferred, or else:

$$\cos\theta = \frac{2mc - q}{2p} \quad (\text{B.8})$$

which must, of course, fall between -1 and 1. Now, $2mc$ has units of momentum, and corresponds to the point at which the two curves in figure 2.4 cross. That is, if $p = q$ and $\frac{p^2}{2m} = cq$, then $q^2 = 2mcq$ and either $q = 0$ (at which point the curves also cross) or else $q = q^* \equiv 2mc$.

So, from equation B.8, we can conclude that:

$$-2p \leq (q^* - q) \leq 2p \quad (\text{B.9})$$

$$(q^* - 2p) \leq q \leq (q^* + 2p) \quad (\text{B.10})$$




Figure B.1: 1-phonon scattering

Thus, for UCN, where $p \ll q^*$, only phonons with $q \approx q^*$ can scatter with UCN by “1-phonon” processes (1st order diagrams). This result is used in section 2.2.

Quantitatively, the crossing point occurs at $E \approx 11$ K or a neutron wavelength $\lambda \approx 8.9$ Å.

B.3 Buildup of Neutrons in the Cell

If trapped UCN are produced at a constant rate P , and decay with a lifetime of τ_n , then the change in the number of trapped neutrons is given by:

$$\frac{dN}{dt} = P - \frac{N}{\tau_n} \quad (\text{B.11})$$

or

$$\tau_n \frac{dN}{dt} = P\tau_n - N \quad (\text{B.12})$$

Since $P\tau_n$ is a constant, we can make the substitution

$$N' = N - P\tau_n \quad (\text{B.13})$$

$$\frac{dN'}{dt} = \frac{dN}{dt} \quad (\text{B.14})$$

This gives us:

$$\frac{dN'}{dt} = -\frac{N'}{\tau_n} \quad (\text{B.15})$$

which can be integrated to give:

$$N' = Ae^{-t/\tau_n} \quad (\text{B.16})$$

Plugging equation B.13 into equation B.16 gives:

$$N - P\tau_n = Ae^{-t/\tau_n} \quad (\text{B.17})$$

$$N = P\tau_n + Ae^{-t/\tau_n} \quad (\text{B.18})$$

Finally, requiring that $N(t = 0) = 0$ gives $A = -P\tau_n$ or:

$$N = P\tau_n(1 - e^{-t/\tau_n}) \quad (\text{B.19})$$

Thus, the production of UCN begins nearly linear in time, but saturates with a time constant of τ_n at a peak trapped population approaching $N_{max} = P\tau_n$, where P is the production rate and τ_n is the lifetime of the trapped neutrons. For a plot of the buildup, see figure 4.1.

B.4 Attenuation of Neutron Absorbers

The attenuation length for thermal neutrons in a neutron absorber is given by:

$$\lambda_{\text{abs}} = \frac{1}{n\sigma} \quad (\text{B.20})$$

where σ is the absorption cross section and n is the number density of absorbing nuclei. The number density n is just given by:

$$n = \rho \frac{N_A}{A} x_{\text{abs}}, \quad (\text{B.21})$$

where ρ is the mass density of the material, N_A is Avogadro's number, A is the atomic weight and x_{abs} is the weight fraction of the absorbing isotope. Similarly, the absorption cross section is given by:

$$\sigma = \frac{v_{300 \text{ K}}}{v} \sigma_{th} = \frac{\lambda}{\lambda_{300 \text{ K}}} \sigma_{th} \quad (\text{B.22})$$

Combining, the final absorption length is given by:

$$\begin{aligned}\lambda_{\text{abs}} &= \frac{\lambda_{300 \text{ K}}}{\lambda \sigma_{th}} \frac{A}{\rho N_A x_{\text{abs}}} \\ &= \frac{2.83A}{\lambda \sigma_{th} \rho x_{\text{abs}}},\end{aligned}\tag{B.23}$$

where λ_{abs} is the absorption length in cm, λ is the neutron wavelength in Å, σ_{th} is the thermal neutron absorption cross section in barns, ρ is the density in g cm⁻³ and x_{abs} is the weight fraction of the absorbing isotope.

As an example, for hexagonal boron nitride (solid, rather than aerosol), with the natural isotopic ratio, A is 10 for the absorbing ¹⁰B, σ_{th} is 3840 barns, ρ is 2.2 g cm⁻³, and the weight fraction of the absorbing isotope is $x_{\text{abs}} = \frac{0.199 \times 10}{0.199 \times 10 + 0.801 \times 11 + 1.0 \times 14} = 0.08$, giving:

$$\lambda_{\text{BN}} = \frac{0.042 \text{ cm}}{\lambda},\tag{B.24}$$

or, about 100 μm for 4 Å neutrons, and about 40 μm for 10 Å neutrons. ⁶Li has a thermal neutron absorption cross section of 941 barns, and an isotopic ratio of 7.5% (with the remainder ⁷Li).

Appendix C

Beam Calculations

In this appendix, $\frac{d\Phi}{d\lambda}$ at the critical wavelength, $\lambda \approx 8.9$ Å, (the 8.9 Å flux, needed in section D to calculate the trapped UCN production rate) is derived, from a measurement of the capture flux at the end of the filter cryostat. At the end of the filter cryostat (less than 1 m from the end of NG-6), the capture flux was measured to be $9.0 \pm 0.45 \times 10^8$ n cm⁻²s⁻¹[57].

C.1 Beam Spectrum at end of NG-6

The velocity spectrum of the neutrons in the cold source is given, in the limit that the neutrons are perfectly thermalized, by:

$$\begin{aligned} d\rho(v) &\propto \exp(-\frac{1}{2}mv^2/k_BT)v^2dv d\Omega \\ &\propto \exp(-v^2/v_T^2)v^2dv d\Omega, \end{aligned} \tag{C.1}$$

where $v_T \equiv \sqrt{k_BT/2m}$. The NIST cold source (section 3.2.1) is at 20 K, but the neutrons are not fully thermalized, having a spectrum closer to 30 K[18]. The neutron

flux leaving the cold source with a velocity v is given by:

$$\begin{aligned}
 d\Phi &\propto v\rho(v) \\
 &\propto v^3 \exp(-v^2/v_T^2) dv d\Omega \\
 &\propto \frac{1}{\lambda^5} \exp(-\lambda_T^2/\lambda^2) d\lambda d\Omega,
 \end{aligned} \tag{C.2}$$

substituting $v = h/m\lambda$.

These neutrons are coupled to the experiment by a neutron guide whose inner surfaces are coated with ^{58}Ni . The walls of the guide reflect neutrons with a glancing incidence up to an angle of $\theta_c = \sin^{-1}(\lambda/\lambda_c)$, where λ is the neutron wavelength and $\lambda_c \approx 500 \text{ \AA}$. Thus, for small $\theta_c \approx \lambda/\lambda_c$ we can write the total flux coupled into the guide as:

$$d\Phi \propto \int_{-\theta_c}^{\theta_c} d\theta_x \int_{-\theta_c}^{\theta_c} d\theta_y \frac{1}{\lambda^5} \exp(-\lambda_T^2/\lambda^2) d\lambda. \tag{C.3}$$

The integrals $d\theta_x d\theta_y$ would just give a factor proportional to λ^2 , except for the finite size of the cold source and the angular dependence of the transmission of the guide. The cold source fills approximately 17 mrad (corresponding to 8.5 \AA) as seen from the end of the guide, so that no neutrons enter the guide at higher angles. Higher angle neutrons are also subject to a greater attenuation in the guide, which can be modelled as a 2% loss per bounce. The physical dimensions of the guide (60 m long, 15 cm tall by 6 cm wide) give a total number of bounces $N = 400\theta_y + 1000\theta_x$, giving an additional factor of 0.98^N in the transmission. This turns out to be a small correction (see section C.3), so for the moment we ignore the transmission of the guide and can

write the total flux reaching the end of NG-6 as:

$$\begin{aligned}\Phi &\propto \int_0^\infty d\lambda \int_{-\theta_c}^{\theta_c} d\theta_x \int_{-\theta_c}^{\theta_c} d\theta_y \frac{1}{\lambda^5} \exp(-\lambda_T^2/\lambda^2) \\ &\propto \int_0^{8.5 \text{ \AA}} d\lambda \frac{1}{\lambda^3} \exp(-\lambda_T^2/\lambda^2) + \int_{8.5 \text{ \AA}}^\infty d\lambda \frac{(8.5 \text{ \AA})^2}{\lambda^5} \exp(-\lambda_T^2/\lambda^2)\end{aligned}\quad (\text{C.4})$$

Finally, at the end of NG-6, the beam passes through a filter cryostat, containing 10 cm each of beryllium and bismuth. To first order, the effect of these on the neutron spectrum can be approximated by eliminating all neutrons with $\lambda < 4.0 \text{ \AA}$, by Bragg scattering in the beryllium. Thus, the first integral is taken from a lower limit of 4 \AA instead of to zero.

When neutrons pass through a “thin” material (i.e. one that attenuates a small fraction of the neutrons), the probability of absorption is proportional to $1/v \propto \lambda$ (where v is the velocity of the neutron). Cross-sections are usually expressed as σ_{th} for neutrons with an energy corresponding to 300 K. Thus, to find the total absorption in a thin material, the number of interest is the total “capture flux”, $\Phi_c \equiv \int_0^\infty \lambda \frac{d\Phi}{d\lambda} d\lambda$. We can write the flux at the end of NG-6, after accounting for the filtering of the beam, as:

$$\frac{d\Phi}{d\lambda} = \begin{cases} 0 & \lambda < 4 \text{ \AA} \\ C \frac{1}{\lambda^3} \exp(-\lambda_T^2/\lambda^2) & 4 \text{ \AA} < \lambda < 8.5 \text{ \AA} \\ C \frac{(8.5 \text{ \AA})^2}{\lambda^5} \exp(-\lambda_T^2/\lambda^2) & \lambda \geq 8.5 \text{ \AA} \end{cases} \quad (\text{C.5})$$

Or a total capture flux of:

$$\begin{aligned}\Phi_C &= C \int_{4 \text{ \AA}}^{8.5 \text{ \AA}} \frac{1}{\lambda^2} \exp(-\lambda_T^2/\lambda^2) \\ &\quad + \int_{8.5 \text{ \AA}}^\infty \frac{(8.5 \text{ \AA})^2}{\lambda^4} \exp(-\lambda_T^2/\lambda^2) d\lambda.\end{aligned}\quad (\text{C.6})$$

Taking $\lambda_T = \lambda_{300\text{ K}} = \frac{h}{\sqrt{2mk_BT}} = 1.78\text{ \AA}$, we find:

$$\frac{d\Phi}{d\lambda}|_{\lambda=8.9} = C \times \frac{8.5^2}{8.9^5} \exp(-1.78^2/8.9^2) = C \times 8.67 \times 10^{-4} \text{ n cm}^{-2} \text{ s}^{-1} \text{ \AA}^{-1} \quad (\text{C.7})$$

$$\Phi_C = C \int_{4\text{ \AA}}^{8.5\text{ \AA}} \frac{1}{\lambda^2} \exp(-1.78^2/\lambda^2) \quad (\text{C.8})$$

$$+ \int_{8.5\text{ \AA}}^{\infty} \frac{(8.5\text{ \AA})^2}{\lambda^4} \exp(-1.78^2/\lambda^2). \\ = C \times 7.79 \times 10^{-2} \text{ n cm}^{-2} \text{ s}^{-1} \quad (\text{C.9})$$

Considering the beam divergence (section C.2), the capture flux measured in the center of the beamline at the end of the filter cryostat is the same as the capture flux in the guide, but the 8.9 Å flux in the trapping region is reduced to about 46% of the 8.9 Å flux in the guide. For the calculation in appendix D, we require units of $\text{n cm}^{-2} \text{ s}^{-1} \text{ K}^{-1}$, so we need to convert equation C.7 from Å to K. To do this, we note that $T = \frac{h^2}{2m\lambda^2 k_B}$ which gives $dT = \frac{-h^2}{m\lambda^3 k_B} d\lambda$. Multiplying these two factors by equation C.7, plugging in $\lambda = 8.9\text{ \AA}$, $\frac{d\lambda}{dT} = 0.39\text{ \AA/K}$, and taking C from C.9 and the measured value of Φ_C , we get:

$$\begin{aligned} \frac{d\Phi}{d\lambda} &= 8.67 \times 10^{-4} \text{ n cm}^{-2} \text{ s}^{-1} \text{ \AA}^{-1} \times 0.46 \times 0.39\text{ \AA/K} \times \frac{9 \times 10^8}{7.79 \times 10^{-2}} \quad (\text{C.10}) \\ &= 1.8 \times 10^6 \text{ n cm}^{-2} \text{ s}^{-1} \text{ K}^{-1} \end{aligned}$$

Comparing equation C.5 to figure C.3 makes clear the limitations of this simple model. Though it agrees quite well at longer wavelengths, outside the range of scattering resonances in most materials, at the lower wavelengths the agreement is quite poor. Though the transmission of the beryllium filter rises rapidly near 4 Å, the wavelength spectrum falls very rapidly as well, making the beryllium “edge” quite soft. In addition, the 4.8 Å dip from the upstream monochromator and several other absorption dips are clearly visible.




Figure C.1: Sketch of diverging beam, for one wavelength, with $\theta_{\max} = \sin^{-1}(\lambda/494 \text{ \AA})$.

A fit to the spectrum measured at a different guide at the NCNR was extrapolated to the trapping region to give a value about 5% higher than equation C.10, $2.0 \pm 0.4 \times 10^6 \text{ n cm}^{-2} \text{ K}^{-1}$ [18]. A measurement of our beam spectrum combined with the measurement of the total capture flux gives a result of $1.6 \times 10^6 \text{ n cm}^{-2} \text{ s}^{-1} \text{ K}^{-1}$ (section C.4).

C.2 Beam Divergence

The divergence of the beam in one dimension (and at one wavelength) is shown in figure C.1. The solid lines indicate the range of angles from each end of the guide. The transverse profile is shown (under the assumption that the flux is constant in θ up to θ_{\max}) at three points. Right at the guide exit, the flux is constant across the width of the guide. Until the point at $L' = 6 \text{ cm} / (2 \times \tan(\theta_{\max}))$, (where the solid

lines cross), all the neutrons diverging from the very center of the distribution are replaced by neutrons converging in from the sides. Beyond this point, the central flux falls to $\Phi(L) = \frac{6\Phi_0}{2 \tan(\theta_{max})L}$.

For $\lambda = 4 \text{ \AA}$ (the lowest wavelength passing the beryllium filter), $L' = 6 \text{ cm} / (2 \times \tan(\sin^{-1}(\lambda/494 \text{ \AA}))) \approx 6 \text{ cm} / 2 \times .008 = 3.7 \text{ m}$, and for the maximum divergence (set by the angular size of the source), $\theta_{max} = 17 \text{ mrad}$, $L' = 1.75 \text{ m}$.

Thus, the measurement of the capture flux was well inside the first region, in which the flux (at the center of the beam) is essentially constant, and this measurement can be taken to be the flux directly at the end of NG-6, corrected for the transmission of the filter cryostat.

In the trapping region of the experiment, 2.6 m from the end of NG-6, the flux of 8.9 \AA neutrons is decreased by a factor of $\Phi(L)/\Phi_0 = 6 \text{ cm} / (2 \times 0.017 \times 260 \text{ cm}) = 0.679$ along each axis, or $\frac{d\Phi}{d\lambda}(L = 260 \text{ cm}) = 0.46 \frac{d\Phi_0}{d\lambda}$.

C.3 Transmission Losses in the Guide

Most of the losses of neutrons along the length of the guide occur as the neutrons reflect off the walls. Thus, the loss rate is strongly angular dependent. Given the aspect ratios of the guide $z:x = 1000$ and $z:y = 400$ (where z is along the length of the guide, x is the width and y is the height), in the small angle approximation, the number of bounces is given by $N = 400\theta_y + 1000\theta_x$. Assuming a 2% loss per bounce, this gives a transmission of:

$$(0.98)^{400\theta_y + 1000\theta_x} \tag{C.11}$$

The total transmission as a function of θ_{\max} can thus be evaluated:

$$\mathcal{T} = \frac{1}{\theta_{\max}^2} \int_0^{\theta_{\max}} d\theta_x \int_0^{\theta_{\max}} d\theta_y (0.98)^{400\theta_y + 1000\theta_x}. \quad (\text{C.12})$$

Evaluating this numerically for the largest acceptance of the guide at 8.5 Å, $\mathcal{T} = 0.79$. For the shortest wavelength passing the beryllium filter (4 Å, or 0.008 mrad), $\mathcal{T} = 0.89$. The effect at any wavelength will fall somewhere between these two, giving a net reduction of about 15% in the total capture flux calculated above, and of about 20% in the 8.9 Å flux. Thus, the ratio of the two (8.9/capture) will decrease by about 6%.

C.4 Measurement of the Beam Spectrum

Ideally, rather than approximating the beam spectrum, one would measure the flux at 8.9 Å directly. In practice, we measured the relative flux as a function of wavelength, and combined this with an absolute measurement of the total capture flux.

The beam spectrum was measured by the time-of-flight (TOF) method, using a rapidly rotating chopper disc to turn the continuous neutron beam into a pulsed beam and detecting the neutrons with a ^3He detector (figure C.2). Just in front of the chopper was a small collimator to define the beam size at a diameter 400 μm . The collimator itself was made from a sheet of cadmium, the hole drilled by hand using a drill which measured 360 μm in diameter, giving a hole with a probable diameter around 410 μm . In front of the cadmium sheet was a piece of lithium plastic with a larger hole, to decrease the gamma radiation produced by neutron absorption in the cadmium. Just behind the collimator, a rotating cadmium disc blocked the beam




Figure C.2: Schematic of the time-of-flight measurement of the neutron spectrum.

except when either of a pair of slits passed in front of the collimator. The slits were each made using a $200\text{ }\mu\text{m}$ thick saw, giving a probable width of the slit of $250\text{ }\mu\text{m}$. The disc was rotating at 67 rotations per second, passing a slot in front of the collimator 134 times each second for about $23\text{ }\mu\text{s}$ (depending on the actual width of the slit). All of the neutrons passing through the slit were detected in a ^3He detector placed 69 cm beyond the chopper. The output of the ^3He detector was recorded with a mutli-channel scaler card* running on a PC, with a $5\text{ }\mu\text{s}$ bin width. The scaler was triggered using an LED and a photodiode detecting the passage of the slit approximately 180° from the neutron beam.

Since the velocity of the neutron is inversely proportional to its wavelength, the arrival time of the neutrons at the detector is directly proportional to the wavelength.

*A Canberra accuspec Fast Multi-channel Scaler.




Figure C.3: Data from the time-of-flight measurement.

A velocity of $\frac{4000}{\lambda}$ m s⁻¹ Å and a flight length of 69 cm gives an arrival time:

$$t_{\text{arr}} = t_0 + 171 \mu\text{s}/\text{\AA} \times \lambda. \quad (\text{C.13})$$

Two points in the spectrum are available to set t_0 , the beryllium edge, at 3.95 Å, and the narrow slice of the beam removed by the monochromator in the neutron guide, which was set at 4.8 Å. In fact, this slice was narrow enough that it could be used to confirm the approximately 25 μs neutron pulse-width. Both points gave roughly the same value for t_0 . The data is shown in figure C.3.

The collimator and chopper were 1.3 m from the end of NG-6. At this distance, at the center of the beam, the flux is still the same as in NG-6, at all wavelengths passing the beryllium (section C.2). However, for the longer wavelengths, this constant-flux region is only 1.5 cm across at 1.3 m. At the time of the measurement,

there was no opportunity to align the beam using a theodolite, and the aperture was instead aligned using a piece of string held taut between two apertures aligned on the beam. If the aperture were misaligned, the 8.9 Å flux would be reduced by a factor of $\frac{6.75-r}{6}$, with r the distance in cm from the beam axis.

If we assume that the spectrum shown in figure C.3 is correct, we can extract both a capture flux and a flux at 8.9 Å, just as we did from the predictions of a thermal distribution, and use the ratio of the two to scale the measured capture flux. Thus, we scale the data by λ and integrate it from $\lambda = 0$ until the noise begins to dominate the signal (above 20 Å, not shown, where the flux is expected to be extremely low in any case). We count a total of 4.83×10^7 counts for the total capture flux, with a value of 4.86×10^5 counts/Å for the flux at 8.9 Å. Following the method of equation C.10, we can use these to calculate an 8.9 Å flux in the trapping region of:

$$\begin{aligned} \frac{d\Phi}{d\lambda} &= 4.86 \times 10^5 \text{ Å}^{-1} \times 0.46 \times 0.39 \text{ Å/K} \times \frac{9 \times 10^8 \text{ n cm}^{-2} \text{ s}^{-1}}{4.83 \times 10^7} \quad (\text{C.14}) \\ &= 1.62 \pm 0.08 \times 10^6 \text{ n cm}^{-2} \text{ s}^{-1} \text{ K}^{-1}, \end{aligned}$$

with the error bars corresponding to uncertainties in the wavelength calibration.

In principle, we can use the time-of-flight measurement as an independent measurement of the flux at 8.9 Å, however the accuracy of such a measurement is limited by the uncertainty in the sizes of the collimation aperture and the chopper slit. Thus, we calculate,

$$\frac{d\Phi}{d\lambda} = \frac{4.86 \times 10^5 \text{ counts/Å} \times 0.39 \text{ Å/K}}{36000 \text{ s} \times \pi(0.0205 \text{ cm})^2 \times (134 \times 23 \text{ μs/s})} = 1.3 \times 10^6 \text{ n cm}^{-2} \text{ s}^{-1} \text{ K}^{-1}, \quad (\text{C.15})$$

with significant uncertainty both from the actual size of the slit and collimator aperture and from the alignment. Any error in the alignment greater than 0.75 cm would

cause this number to be less than the actual flux, by a factor of $\frac{6.75 \text{ cm}-r}{6 \text{ cm}}$. (The difference between the results in equations C.14 and C.15 could be completely explained if the chopper aperture was 1.9 cm off axis.)

Appendix D

Calculated UCN Production Rate

The cross section for neutron scattering from the critical momentum, $\hbar k_0^*$, to a much lower momentum, $\hbar k_{ucn}$, in a bath of superfluid helium is given by[43]:

$$\sigma(k_0^* \rightarrow k_{ucn}) = \sigma_{coh} \alpha S_{coh}(k_0^*) \frac{k_{ucn}}{k_0^*}, \quad (\text{D.1})$$

where σ_{coh} is the neutron scattering cross section for a bound helium nucleus, and $\alpha S_{coh}(k_0^*)$ is the structure factor for superfluid helium, which is just $S(\vec{k}, \omega)$ integrated over ω . The total rate of downscattering to an energy below the trap energy at some point in the trap is thus given by:

$$\Phi^* n \int_0^{E_{\max}(\vec{r})} dE_{ucn} \sigma(k_0 \rightarrow k_{ucn}), \quad (\text{D.2})$$

where Φ^* is the spectral flux density of the neutron beam at the critical momentum, n is the atomic density of liquid helium at 250 mK and E_{\max} is the maximum kinetic energy to which a neutron can scatter at that point and still be trapped. Thus $E_{\max}(\vec{r}) = E_T - U(\vec{r})$ is the difference between the maximum total energy of a trapped UCN, E_T and the magnetic potential at that point $U(\vec{r})$.

Since the incident beam is unpolarized, half of the neutrons scattering to UCN energies will have the correct spin state to be trapped, giving a total production of trapped neutrons of:

$$\begin{aligned}
P &= \frac{1}{2} \int dV \Phi^* n \int_0^{E_{\max}(\vec{r})} \sigma(k_0^* \rightarrow k_{ucn}) dE_{ucn} \\
&= \frac{1}{2} \int dV \Phi^* n \frac{\sigma_{coh} \alpha S_{coh}(k_o^*)}{k_0^*} \int_0^{E_{\max}(\vec{r})} k_{ucn} dE_{ucn} \\
&= \frac{1}{2} n \frac{\sigma_{coh} \alpha S_{coh}(k_o^*)}{k_0^*} \int \Phi^* dV \int_0^{E_{\max}(\vec{r})} k_{ucn} dE_{ucn}. \tag{D.3}
\end{aligned}$$

The trapping field (section 3.3) can be approximated by a cylindrically symmetric potential, with infinite potential for $z < 0$, $z > l$ and rising linearly with radius to a maximum of E_T at r_T , where r_T is the radius of the innermost material in the trapping region:

$$U(\vec{r}) = \begin{cases} E_T \frac{r}{r_T} & \text{for } 0 < z < l \\ \infty & \text{for } z < 0, z > l \end{cases} \tag{D.4}$$

Since neutrons able to reach r_T will be absorbed by materials found there, only neutrons with a total kinetic energy, E , such that $E + U(\vec{r}) < E_T$ will remain trapped indefinitely. Thus $E_{\max}(\vec{r})$, the maximum kinetic energy of a neutron which is trapped, can be expressed as $E_{\max}(\vec{r}) = E_T(1 - \frac{r}{r_T})$ for $0 < z < l$.

Plugging $E_{\max}(\vec{r})$ into equation D.3 and integrating over dz gives:

$$P = \frac{1}{2} l n \frac{\sigma_{coh} \alpha S_{coh}(k_o^*)}{k_0^*} \int_0^{r_T} \int_0^{2\pi} \Phi^* r dr d\theta \int_0^{E_T(1-\frac{r}{r_T})} k_{ucn} dE_{ucn} \tag{D.5}$$

where $k_{ucn} = \sqrt{\frac{2m}{\hbar^2}} E_{ucn}^{\frac{1}{2}}$.

$$\begin{aligned}
P &= \frac{1}{2}l \, n \frac{\sigma_{coh} \alpha S_{coh}(k_o^*)}{k_0^*} \int_0^{r_T} \int_0^{2\pi} \Phi^* \, r \, dr \, d\theta \int_0^{E_T(r)} \sqrt{\frac{2m}{\hbar^2}} E_{ucn}^{\frac{1}{2}} dE_{ucn} \\
&= \frac{1}{2}l \, n \frac{\sigma_{coh} \alpha S_{coh}(k_o^*)}{k_0^*} \int_0^{r_T} \int_0^{2\pi} \Phi^* \, r \, dr \, d\theta \frac{2}{3} \sqrt{\frac{2m}{\hbar^2}} E_T^{\frac{3}{2}}(r) \\
&= \frac{1}{3}l \, n \sigma_{coh} \alpha S_{coh}(k_o^*) \frac{k_T}{k_0^*} E_T \int_0^{r_T} \int_0^{2\pi} \Phi^* \left(1 - \frac{r}{r_T}\right)^{\frac{3}{2}} r \, dr \, d\theta
\end{aligned} \tag{D.6}$$

Φ^* is approximately constant across the transverse area of the trap, until it is cut off sharply by the collimator, at $r = r_c$:

$$\Phi^*(r) = \begin{cases} \Phi_0^* & \text{for } r < r_c \\ 0 & \text{for } r \geq r_c \end{cases} \tag{D.7}$$

Thus, we can integrate over the transverse area of the trap to get a production rate of:

$$\begin{aligned}
&= \frac{1}{3}l \, n \sigma_{coh} \alpha S_{coh}(k_o^*) \frac{k_T}{k_0^*} E_T \int_0^{r_T} \int_0^{2\pi} \Phi^* \left(1 - \frac{r}{r_T}\right)^{\frac{3}{2}} r \, dr \, d\theta \\
&= \frac{1}{3}l \, n \sigma_{coh} \alpha S_{coh}(k_o^*) \frac{k_T}{k_0^*} E_T \Phi_0^* \int_0^{r_c} \left(1 - \frac{r}{r_T}\right)^{\frac{3}{2}} r \, dr \int_0^{2\pi} d\theta \\
&= \frac{2\pi}{3}l \, n \sigma_{coh} \alpha S_{coh}(k_o^*) \frac{k_T}{k_0^*} E_T \Phi_0^* \int_0^{r_c} \left(1 - \frac{r}{r_T}\right)^{\frac{3}{2}} r \, dr \\
&= \frac{2\pi}{3}l \, n \sigma_{coh} \alpha S_{coh}(k_o^*) \frac{k_T}{k_0^*} E_T \Phi_0^* \left[-\frac{2}{5} r_c r_T \left(1 - \frac{r_c}{r_T}\right)^{\frac{5}{2}} + 0 + \frac{2}{5} r_T \int_0^{r_c} \left(1 - \frac{r}{r_T}\right)^{\frac{5}{2}} dr \right] \\
&= \frac{2\pi}{3}l \, n \sigma_{coh} \alpha S_{coh}(k_o^*) \frac{k_T}{k_0^*} E_T \Phi_0^* \left[-\frac{2}{5} r_c r_T \left(1 - \frac{r_c}{r_T}\right)^{\frac{5}{2}} - \frac{4}{35} r_T^2 \left(1 - \frac{r_c}{r_T}\right)^{\frac{7}{2}} + \frac{4}{35} r_T^2 \right] \tag{D.8}
\end{aligned}$$

Plugging the values of several measured parameters:

$$\begin{aligned}
n &= 2.17 \times 10^{22} \text{ atoms/cm}^3 \\
\sigma_{coh} &= 1.3 \times 10^{-24} \text{ cm}^2 \\
\alpha &= 1.45 \\
S_{coh}(k_0^*) &= 0.12
\end{aligned}$$

and of several parameters of the experimental trap:

$$\begin{aligned}
 l &= 30 \pm 1 \text{ cm} \\
 r_c &= 0.889 \pm 0.1 \text{ cm} \\
 r_T &= 1.588 \pm 0.1 \text{ cm} \\
 B_T &= B_{max} - B_{min} = 1.15 \text{ T} - 0.11 \text{ T} = 1.04 \pm 0.1 \text{ T} \\
 E_T &= 0.7 \times B_T = 0.728 \pm 0.07 \text{ mK}
 \end{aligned}$$

into equation D.8 yields:

$$\begin{aligned}
 P &= \frac{2\pi}{3} 30 \cdot 2.17 \times 10^{22} \cdot 1.3 \times 10^{-24} \cdot 1.45 \cdot 0.12 \sqrt{\frac{0.728 \times 10^{-3} \text{ K}}{11 \text{ K}}} \cdot 0.728 \times 10^{-3} \text{ (D.9)} \\
 &\quad \times \Phi_0^* \left[-\frac{2}{5} \cdot 0.889 \cdot 1.588 \left(1 - \frac{0.889}{1.588}\right)^{\frac{5}{2}} - \frac{4}{35} \cdot 1.588^2 \left(1 - \frac{0.889}{1.588}\right)^{\frac{7}{2}} + \frac{4}{35} \cdot 1.588^2 \right] \\
 &= 3.63 \times 10^{-7} \times \Phi_0^*,
 \end{aligned}$$

where P has units of s^{-1} and Φ_0^* has units of $\text{n cm}^{-2} \text{ K}^{-1}$. Finally plugging in $\Phi_0^* = 2.0 \pm 0.4 \times 10^6 \text{ n cm}^{-2} \text{ K}^{-1}$, $P = 0.73 \pm 0.24 \text{ s}^{-1}$.

Appendix E

Activation Testing Results

The table below gives the results of several “rabbit tube” activation tests of materials used in the experiment. In such a measurement, the material moved in a pneumatic “rabbit” tube into the reactor core (with a flux of 4×10^{14} n/cm²/s) for a few minutes, and then quickly placed in front of a germanium counter, which detects gamma rays with sufficient energy resolution to distinguish the isotopic decay which produced the gamma. The spectrum of gamma radiation coming off the sample thus indicates what radioactive nuclei have been produced in the sample. The isotopes detected are listed at the top of the chart, with their half-life (or half-lives) listed below.* Each material is listed along with the proportion in parts per million detected of each impurity.

UVT PMMA refers to the ultra-violet transmitting acrylic, purchased from Boedeker Plastics. The TPB is the tetraphenyl butadiene scintillator, purchased from Aldrich. The PFA is the DuPont fluoropolymer film. Aerodag G is the aerosol graphite from Acheson Colloids, while “graphite” refers to a piece of pyrolytic graphite

*Units: m=min, h=hour.

from Poco Graphite. The UVT PMMA was used to form the detector and lightpipes; the TPB formed the inner layer of the detector; the PFA formed the neutron vacuum windows; the Aerodag was sprayed on all of the boron nitride shielding; and the solid graphite formed an additional light-blocking shield inside the boron nitride tube around the detector.

	Al	V	Cu	Ti	Br	Mg	I	Cl	Mn	Na
$t_{1/2}$	2.25 m	3.76 m	$\frac{5.10 \text{ m}}{12.7 \text{ h}}$	5.76 m	$\frac{6.1 \text{ m}}{17.66 \text{ m}}$	9.45 m	25 m	37.2 m	2.578 h	14.95 h
UVT PMMA	3.5	0.182	3.4	6.4	< 3	< 10	< 0.2	1.4	0.028	0.36
TPB	160		0.18					8	0.09	9.4
PFA								41.4		
Aerodag	3.7	0.75	8.9	3.4		13.8			1.6	
graphite	0.22	0.28	0.20	0.80	< 0.03	< 2.5			0.008	

Appendix F

Some Computer Routines

F.1 Detection Efficiency Calculation

```
#include <stdio.h>
#include <stdlib.h>
#include <math.h>

/* Cell Parameters */
#define len (34.0) /* length of detection region in cm */
#define att (0.0164) /* Attenuation per cm in tube */
#define bd (0.0) /* Beam Dump transmission */
#define eff (0.8) /* overall efficiency scaling factor */
/* 0.8 transmission to pmts through Y-piece */
#define pewalls (14.5) /* p.e. per MeV through walls (i.e. bd=0.0) */
/* Beta Spectrum */
#define E_max (1.2926) /* Maximum electron energy in MeV */
#define M_e (0.511) /* electron mass in MeV */
/* Program Parameters */
#define N 1001
/* Constants for Ran1, from Numerical Recipes */
```

```

#define IA 16807
#define IM 2147483647
#define AM (1.0/IM)
#define IQ 127773
#define IR 2836
#define NTAB 32
#define NDIV (1+(IM-1)/NTAB)
#define EPS (1.2e-7)
#define RNMX (1-EPS)
/* Functions */
void Open_File();
double iBeta(double E); /* Integrated Beta-decay spectrum of neutron */
float ran1(long *idum); /*From Numerical Recipes */
/* File Global Variable – output file */
FILE *fout;

main (int argc, char **argv)
{
    int i,j;
    double z,p,q,E,step,tot,fac;
    double ex;
    double pe[40],lnfac[40],coinc[40];
    double cdet,ctot;
    long idum;

    /* Open Output File */
    if (argc > 2) {
        fprintf(stderr, "Error in command line args – monte <filename>\n");
        exit (1);
    }
    else if((argc == 2) && ((fout = fopen(argv[1], "w")) == NULL)) {
        fprintf(stderr, "cannot open %s\n", argv [1]);
        exit(-6);
    }
}

```

```

}
if (argc == 1) {
    fout=stdout;
}

/* coinc[j] is the probability of getting 2+ pe per pmt
given that there will be j total pe in the two pmts */
coinc[0]=0.0;
coinc[1]=0.0;
coinc[2]=0.0;
coinc[3]=0.0;
coinc[4]=0.3750;
coinc[5]=0.6253;
coinc[6]=0.7793;
coinc[7]=0.8754;
coinc[8]=0.9292;
coinc[9]=0.9604;
coinc[10]=0.9789;
coinc[11]=0.9885;
coinc[12]=0.9939;
coinc[13]=0.9967;
for (j=14;j<40;j++) coinc[j]=1.0;

/* Set the seed for the random number generator */
idum=-3241498;
tot=0.0;
for (j=0; j<40; j++) {
    pe[j]=0;
    if (j==0) {lnfac[0]=0.0;}
    else {lnfac[j]=lnfac[j-1]+log(j);}
}
for (i=1; i<N; i++) {
    p=ran1(&idum);

```



```

    step=(E_max-M_e)/2;
    E=M_e;
    q=-1000.0;
    for (j=0; fabs(q-p)>0.00001; j++) {
        if (q<p) {E+=step;}
        else {E-=step;}
        step=step/2;
        q=iBeta(E);
    }
    E=E-M_e;
    z=len*ran1(&idum);
    /* E=0.370;
    z=len/2.0;*/
    fac=1+(z-len/2)*att;
    ex=(fac*pewalls*E)*eff;
    for (j=0; j<40; j++) {
        pe[j]+=exp(j*log(ex)-ex-lnfac[j]);
    }
}
for(j=0; j<40; j++) fprintf(stdout,"%e\n",pe[j]);
cdet=0.0;
ctot=0.0;
for(j=0; j<40; j++) {
    ctot+=pe[j];
    cdet+=pe[j]*coinc[j];
}
fprintf(stdout,"total: %e %e %e\n",cdet,ctot,cdet/ctot);

} /* Main */

double iBeta(double E) /* Integrated Beta spectrum */
{
    double d;

```

```

if (E>E_max) E=E_max;
if (E<M_e) E=M_e;
d=E*E-M_e*M_e;
return ((
    sqrt(d)*( -8*M_e*M_e*M_e*M_e
        +M_e*M_e*(-4*E*E+15*E*E_max-20*E_max*E_max)
        +2*E*E*(6*E*E-15*E*E_max+10*E_max*E_max) )/60
        +M_e*M_e*M_e*M_e*E_max*(log(E+sqrt(E*E-M_e*M_e))-log(M_e))/4
    )*17.6076);
}

```

```

float ran1(long *idum) /* From Numerical Recipes */
{
    int j;
    long k;
    static long iy=0;
    static long iv[NTAB];
    float temp;

    if (*idum <= 0 —— !iy) {
        if (-(*idum) < 1) *idum=1;
        else *idum = - (*idum);
        for (j=NTAB+7; j>=0; j-) {
            k=(*idum)/IQ;
            *idum=IA*(*idum-k*IQ)-IR*k;
            if (*idum < 0) *idum += IM;
            if (j<NTAB) iv[j]=*idum;
        }
        iy=iv[0];
    }
    k=(*idum)/IQ;
    *idum=IA*(*idum-k*IQ)-IR*k;
    if (*idum < 0) *idum += IM;
}

```

```
j=iy/NDIV;  
iy=iv[j];  
iv[j]=*idum;  
if ((temp=AM*iy)>RNMX) return RNMX;  
else return temp;  
} /* ran1 */
```

Appendix G

Trap Losses

G.1 Marginal Trapping

UCN inside the trapping region with energies below the trap depth are trapped. Neutrons with energies greater than the trap depth, though not trapped, can sometimes be found in trajectories such that they remain in the trapping region for many seconds. The decay of such neutrons will give a signal indistinguishable from that of the decay of trapped neutrons, and this component of our signal will fall off more rapidly than through beta decay alone. Thus, these “marginally trapped” neutrons are a potential source of a systematic error in a precise measurement of the neutron lifetime.

The magnetic potential in our trap is shown in figures 3.7 and 3.8. In considering marginal trapping, it is useful to consider a slightly simpler trap, identical in the x-y projection (figure 3.7) with perfectly flat walls at the ends (a pair of square waves in place of the rounded potentials from the solenoids in the top of figure 3.8, or a rectangular well in the bottom figure). In practice, the actual rounded ends of the trap make marginal trapping harder in the real trap by directly coupling momentum along the axis of the trap with momentum in the radial direction[75]. Thus, a method which eliminates all the neutrons which would be marginally trapped in this simplified version will certainly eliminate all the neutrons which would be marginally trapped in our real trap.

In this simplified version of the trap, we can separate the total Hamiltonian

into functions of the longitudinal and transverse coordinates:

$$H = H_L + H_{xy} = \left(\frac{p_z^2}{2m} + V_{sol}(z) \right) + \left(\frac{mv_x^2 + mv_y^2}{2} + \mu\beta\sqrt{x^2 + y^2} \right), \quad (\text{G.1})$$

where μ is the neutron's magnetic moment and β is the magnetic potential gradient. Further, if we write the transverse energy as a function of $r = \sqrt{x^2 + y^2}$ and θ , the Hamiltonian is independent of θ . Thus $\ell \equiv \frac{d}{d\theta}H = mr^2\dot{\theta}$ is a constant of the motion and we can write the total energy as:

$$H = \left(\frac{p_z^2}{2m} + V_{sol}(z) \right) + \left(\frac{p_r^2}{2m} + \frac{\ell^2}{2mr^2} + \mu\beta r \right). \quad (\text{G.2})$$

Now the condition for a neutron to be confined is no longer $E_{tot} \leq E_T$, but rather ($E_L \leq V_{sol}$) while ($E_{xy} \leq E_{xy}^{max}$). In the transverse plane, we can consider the possible trajectories to fall into three categories: for $\ell = 0$ neutrons follow one-dimensional trajectories straight through the center of the trap; for $\ell^2 = \mu\beta mr^3$ the trajectories are circular ($\frac{\ell^2}{mr^3} = \frac{mv^2}{r} = \frac{d}{dr}V = \mu\beta$); for all intermediate ℓ s the trajectories can be thought of as a combination of these two. For the $\ell = 0$ trajectories, the trapping condition becomes $E_{xy} \leq \mu\beta r_T$. Since ℓ is conserved, trajectories with $\ell > 0$ have can have up to $\frac{\ell^2}{2mr_T^2}$ additional energy, without being able to escape the trap (up to $\frac{\ell^2}{2mr_T^2} = \frac{\mu\beta mr_T^3}{2mr_T^2} = \frac{\mu\beta r_T}{2}$ for the circular orbit with radius r_T). Thus, the maximum energy which can be confined in this idealized version of our trap is $V_{sol} + \frac{3}{2}\mu\beta r_T$. For $V_{sol} = \mu\beta r_T = E_T$ this corresponds to a maximum confined energy of $2.5E_T$, considerably higher than the maximum trapped energy.

A reasonable way to eliminate any systematic effects arising from neutrons in such trajectories would be to eliminate all such neutrons— by ramping the magnetic field. After loading the trap, one can lower the magnetic field part of the way, wait a short time for neutrons with energies greater than $2.5E_T$ at the decreased trap depth to all leave the trap, and then raise the field back up. As the field increases, the potential energy of the neutrons rises as well. If the field were somehow raised instantaneously, then a neutron with $r = r_T$ at that instant would gain just as much energy as the increase in the trap depth. Thus, some marginally trapped neutrons would remain. However, for a ramping rate that is slow compared to the frequency of the orbits, the energy of the neutrons will increase more slowly than does the trap depth. We can consider the effect of the field ramping on the neutrons in each of the three categories.

In the $\ell = 0$ trajectories, the neutron initially (just before ramping back up) sees magnetic potentials ranging from $V = 0$ to $V = \beta\mu r_0 = E$. As the neutron travels up the field gradient, it loses kinetic energy and thus velocity, spending more time, $t \propto (E - \mu\beta r)^{-1/2}$, at the higher magnetic field. If the ramping rate is slow compared to the orbit frequency, we can simply consider $\langle V \rangle$, the time-averaged magnetic potential, to be the field seen by that trajectory. Integrating the potential $(\mu\beta r)$ with the weighting $(E - \mu\beta r)^{-1/2}$, we get:

$$\left(\int_0^{E/\mu\beta} \frac{1}{\sqrt{E - \mu\beta r}} dr \right) \langle V \rangle = \int_0^{E/\mu\beta} \frac{\mu\beta r}{\sqrt{E - \mu\beta r}} dr. \quad (\text{G.3})$$

Substituting $x = 1 - \frac{\mu\beta r}{E}$, $dx = -\frac{\mu\beta}{E} dr$ gives:

$$\begin{aligned} \left(\frac{E}{\mu\beta\sqrt{E}} \int_0^1 x^{-1/2} dx \right) \langle V \rangle &= \frac{E^2}{\mu\beta\sqrt{E}} \int_0^1 x^{-1/2} dx - \frac{E^2}{\mu\beta\sqrt{E}} \int_0^1 x^{1/2} dx \\ \langle V \rangle &= -\frac{2}{3}E. \end{aligned} \quad (\text{G.4})$$

As the field is ramped up ($\beta \rightarrow \beta + d\beta$), the potential increases as well, as $\langle V \rangle \rightarrow \langle V \rangle (1 + \frac{d\beta}{\beta})$, but the kinetic energy remains constant. Thus $E \rightarrow E(1 + \frac{2}{3}\frac{d\beta}{\beta})$, or (with r_0 representing the maximum radius):

$$\begin{aligned} dE &= \frac{2}{3}\mu(d\beta)r_0 \\ \mu(d\beta)r_0 + \mu\beta dr_0 &= \frac{2}{3}\mu(d\beta)r_0 \\ \frac{dr_0}{r_0} &= -\frac{1}{3}\frac{d\beta}{\beta} \rightarrow \ln r_0 = -\frac{1}{3}\ln \beta \rightarrow r_0 \propto \beta^{-1/3} \end{aligned} \quad (\text{G.5})$$

Thus, the maximum radius of the trajectory decreases as $\beta^{-1/3}$ as the magnetic field is ramped up.

For neutrons in circular orbits, by definition $\dot{r} = 0$. Thus, $\frac{dH}{dr} = 0$ as well, giving $0 = -\frac{\ell^2}{mr^3} + \mu\beta$ or $r^3\beta = \frac{\ell^2}{\mu m}$. If the field is ramped slowly compared to the orbit frequency, the neutron will spiral inward slowly, gaining an insignificant amount of radial velocity. The final trajectory will thus still be circular. Since ℓ is a constant, as β increases, r will again decrease as $\beta^{-1/3}$.

Neutrons in the third category can be thought of as a combination of the two bounding cases. Like the two pure trajectories, provided that the ramping is sufficiently slow, $\beta^{1/3}r_{max}$ will remain a constant as the field is ramped.

We determined above that the maximum transverse energy of a marginally trapped neutron is $1.5E_T$. That is to say, a neutron with an energy greater than this will reach r_T during each orbit. If we assume that any neutron reaching r_T is lost (or more realistically, that after one hundred bounces, $\approx 1\text{--}2$ s, nearly every neutron will be either absorbed or scattered) then we need to ramp down enough that neutrons with $E = 1.5E_T$ at the minimum of the ramping, that is the maximum energy that isn't lost, will have an energy less than the (increased) trap depth after ramping back up, that is r_{max} after ramping should be no more than $\frac{2}{3}r_T$. Although for most trajectories this is a more stringent condition than is required, it is always sufficient to show that no neutrons will be marginally trapped by the quadrupole. Since $\beta_i^{1/3}r_i = \beta_f^{1/3}r_f$, this condition becomes $\frac{\beta_f}{\beta_i} = \left(\frac{2}{3}\right)^3 \approx 0.295$. Thus, if we ramp the magnet to less than 0.295 of its normal maximum, wait long enough for all the neutrons with energy greater than $1.5E_T$ to be lost, and then ramp back up, we can be certain that there are no longer any marginally trapped neutrons.

In addition to removing all marginally trapped neutrons this procedure will also remove many neutrons which would otherwise have been trapped. We can estimate the fraction lost by considering the effect on the trapped neutrons as we ramp down. We know that r_f , the maximum radius of the trajectory after ramping, is given by $r_i \left(\frac{\beta_i}{\beta_f}\right)^{1/3} = 1.5$. Thus, any neutron with a maximum radius initially greater than $\frac{2}{3}r_T$ will be able to reach r_T during the minimum of the ramping, and could potentially be lost. We calculated in appendix D that the number of trapped UCN produced varies as $E_T^{1.5}$. Throwing out all neutrons with a maximum radius of $\frac{2}{3}r_T$ is equivalent to having r_T (and hence E_T) two-thirds as large during the loading (since those which don't fulfill this condition will be lost during ramping). This leaves a trapped population $\left(\frac{2}{3}\right)^{1.5} \approx 0.54\times$ as large as without ramping. So throwing out all of the marginally trapped neutrons will cost roughly a factor of two in trapped UCN.

G.2 Majorana Transitions

The interaction between a neutron and a magnetic field can be written as:

$$H = \mu_n \vec{\sigma}_n \cdot \vec{B}. \quad (\text{G.6})$$

In a field with a spatially varying magnitude, this leads to a force on the neutron $F = -\mu_n \vec{\nabla}(\vec{\sigma}_n \cdot \vec{B})$, which can be used to form a trapping potential, as described in section 2.1. The same Hamiltonian yields an oscillation in $\vec{\sigma}_n$ with frequency $\omega = \frac{2\mu_n B}{\hbar}$. Where this precession is much faster than the change in the direction of \vec{B} (i.e. where $\frac{2\mu_n B}{\hbar} \gg \frac{d\vec{B}/dt}{|\vec{B}|}$), equation G.6 reduces to the much simpler:

$$H = \mu_n B, \quad (\text{G.7})$$

giving the magnetic potential which forms any static magnetic trap.

When $\frac{2\mu_n B}{\hbar} \gg \frac{d\vec{B}/dt}{|\vec{B}|}$ is not the case, the spin no longer perfectly follows the field. In a magnetic trap, when the spin “flips” to the opposite orientation relative to the field, it is no longer confined by the trapping field, and is quickly lost from the trap. This can typically occur where the magnitude of the field is low, or especially near a zero-field point. Such points exist in many traps, such as the anti-Helmholtz and spherical hexapole traps described in section 2.1. The Ioffe trap used in this experiment has a non-zero field everywhere in order to suppress the loss of trapped neutrons through these “Majorana” spin-flip losses.

In this section, we discuss the level at which a transverse field suppresses these losses. In order to make a worst case estimate, we take the very worst possible trajectory. Considering the neutron trajectories described in section G.1, clearly the $\ell = 0$ trajectories will most closely approach the field-minimum at the central axis of the quadrupole. $d\vec{B}/dt$ in these trajectories will be given roughly by $\beta\dot{r}$. Thus, the trajectory most likely to undergo a Majorana spin-flip is the $\ell = 0$ state with the highest velocity (v_0) as it approaches the center, that is the one with $E = E_T$. The force on the neutron is given by $-\mu\beta$, so we can write down the time dependence of the neutron’s position as $r(t) = v_0 t - \frac{\mu\beta}{m} t^2$. The magnetic field is given by the quadrupole field, $\beta r \hat{r}$, and a constant bias field, B_0 , perpendicular to the quadrupole field. Choosing axes where $\hat{r} = \hat{z}$ and the bias field is given by $B_0 \hat{x}$, we can write down the magnetic field as a function of position as $\vec{B}(z) = B_0 \hat{x} + \beta z \hat{z}$. Finally, we can write the Hamiltonian (equation G.6) using the standard Pauli spin matrices as:

$$H = \mu_n \vec{\sigma} \cdot \vec{B} = \mu_n \begin{pmatrix} \beta z & B_0 \\ B_0 & -\beta z \end{pmatrix} = \begin{pmatrix} \beta v_0 t - \frac{\beta^2 \mu t^2}{m} & B_0 \\ B_0 & \frac{\beta^2 \mu t^2}{m} - \beta v_0 t \end{pmatrix}. \quad (\text{G.8})$$

Making the simplifying substitutions:

$$a \equiv \frac{\mu^2 \beta^2}{\hbar m} \quad (\text{G.9})$$

$$b \equiv \frac{\mu \beta v_0}{\hbar} \quad (\text{G.10})$$

$$c \equiv \frac{\mu B_0}{\hbar}, \quad (\text{G.11})$$

we can then write the equations of motion for the neutron's spin as:

$$\frac{d}{dt} \begin{pmatrix} x \\ y \end{pmatrix} = -i \begin{pmatrix} bt - at^2 & c \\ c & at^2 - bt \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}, \quad (\text{G.12})$$

where x and y are each complex, giving a set of four coupled first-order differential equations.

Plugging in the actual parameters of our trap (and the specified trajectory): $\beta = 60$ T/m, $v_0 = 3.3$ m/s, $B_0 = 0.11$ T, give:

$$\begin{aligned} a &= 1.898 \times 10^{12} \text{s}^{-3} \\ b &= 1.812 \times 10^{10} \text{s}^{-2} \\ c &= 1.007 \times 10^7 \text{s}^{-1}. \end{aligned} \quad (\text{G.13})$$

These equations can be integrated numerically, starting with a spin in the \hat{x} direction at the center of the axis, integrating these equations out to r_T and then taking $\langle \vec{\sigma} \cdot \vec{B} \rangle$. This gives a $\langle \vec{\sigma} \cdot \hat{B} \rangle = 1.0 - 4 \times 10^{-9}$, or a transition into the anti-parallel spin state of 2×10^{-9} . A passage from one side to the other ($z = -r_T$ to $z = +r_T$) should give four times this transition rate or 8×10^{-9} per crossing. This trajectory crosses through the center of the trap once every 100 ms, giving a transition rate of $8 \times 10^{-8} \text{s}^{-1}$ *, or about 10^{-4} of the neutron beta-decay rate.

Clearly this rate can be further suppressed (or unsuppressed) by the addition of a constant field parallel to the present bias field. The addition of a coil to add to the bias field in this manner could also be used to decrease the bias field, intentionally raising the Majorana spin-flip rate to a more easily detectable range. Thus, the

*Assuming that the spin-flipped component, which is pushed out to the walls by the magnetic field, rapidly loses coherence with the trapped component.

predicted transition rate could be tested in the real apparatus. Calculations using the numerical method described above give a transition rate of $7.3 \times 10^{-9} \text{ s}^{-1}$ ($6 \times 10^{-6} \times \frac{1}{\tau_\beta}$) for a bias field of 0.2 T.

G.3 Absorption of Trapped UCN by ^3He

The thermal cross section of ^3He for neutron capture is 5330 barns ($5.33 \times 10^{-21} \text{ cm}^2$). This is the effective average cross section of one unpolarized ^3He nucleus seen by a thermal neutron (one with an energy corresponding to 300 K). The cross section scales with the neutron velocity as $1/v$, increasing with decreasing velocity.

The attenuation length, λ_{abs} , of a neutron beam of velocity v_n passing through a density nx_3 of ^3He nuclei (where $n = 2.17 \times 10^{22} \text{ cm}^{-3}$ is the density of He nuclei and x_3 is the ratio of ^3He to ^4He), is given by:

$$\frac{1}{\lambda_{abs}} = \frac{v_{300K}}{v_n} \sigma_{th} n x_3 \quad (\text{G.14})$$

$$\lambda_{abs} = \frac{8.6 \times 10^{-3} \text{ cm}}{x_3} \frac{\lambda_{300K}}{\lambda n} = 1.73 \text{ \AA}$$

For trapped UCN it makes sense to multiply the capture per unit length, $1/\lambda_{abs}$ by the velocity of the UCN to get capture per unit time, and invert to get a lifetime:

$$\frac{1}{\tau} = v_{UCN} \cdot \frac{1}{\lambda} = v_{300K} \sigma_{th} n x_3$$

$$\tau = \frac{1}{v_{300K} \sigma_{th} n x_3} \quad (\text{G.15})$$

Unlike the attenuation length, which is proportional to the neutron's velocity, the capture lifetime is the same for all UCN energies. Plugging in some numbers:

$$\frac{1}{\tau} = (2.25 \times 10^5 \text{ cm/s})(5.33 \times 10^{-21} \text{ cm}^2) \quad (\text{G.16})$$

$$\times (2.17 \times 10^{22} \text{ atoms/cm}^3) x_3$$

$$= \frac{1}{x_3} \cdot 3.9 \times 10^{-8} \text{ s}$$

Appendix H

Superfluid Helium and the Neutron Lifetime

It has been suggested, in connection with this experiment, that the beta-decay lifetime of a neutron in a bath of liquid helium might differ from that of a neutron in vacuum. In this appendix, several possible mechanisms leading to such a shift are considered, and the magnitudes of the corrections due to each are estimated relative to the uncorrected beta-decay lifetime.

One set of mechanisms which will affect the beta-decay lifetime of the neutron relate to small changes in the phase space factor, due to an increased or decreased amount of energy available to the decay products. The leading order term in the phase space factor is given by $f = \int_{m_e}^{E_{max}} p_e E_e (E_{max} - E_e)^2 dE_e$, where E_{max} is the maximum electron energy, with a magnitude of 1.2926 MeV (appendix A). Simple integration of this term with different values of E_{max} shows that $\frac{df}{f} \approx 5 \times 10^{-6} \text{ eV}^{-1} dE_{max}$, giving a corresponding fractional change in the neutron lifetime. (Reference [11] confirms this dependence in the full f^R for dE_{max} up to 250 eV.)

The energy of a neutron with a sufficiently low energy that its wavelength is longer than the interatomic spacing experiences an effective potential arising from the averaged interaction with the nuclei in the material. In liquid helium, this potential is of order $2 \times 10^{-8} \text{ eV}$, giving a relative decrease in the neutron lifetime of 10^{-13} .

The energy of the electron is also slightly higher in the helium. This energy has been measured to be 1.3 eV, by measuring the minimum energy required for electrons to penetrate a liquid helium surface[76]. The electron can be thought of

as occupying a self-formed void in the helium, and the 1.3 eV associated with the zero-point energy of the electron confined to this void. The net effect of this shift is to increase the neutron lifetime by about 6.5×10^{-6} , by reducing the maximum kinetic energy available to the electron.

The energy of the proton is actually slightly lower in the helium than in the vacuum. This can be thought of as a screening of the proton's charge by the partial polarization of helium atoms clustered near the proton. The energy shift is of order 2 eV*, which could decrease the neutron lifetime by as much as 1×10^{-5} . However, it seems quite plausible that the clustering and polarization of the helium atoms will happen subsequent to the decay, and have no effect. In any case the effect will be no more than about 1×10^{-5} .

In addition to changes in the phase space, there is the possibility of the matrix element for the decay being influenced by the presence of the helium nuclei. We can approach this effect by considering the decay of neutrons in even closer proximity to other nucleons: those bound in nuclei. In comparing beta decay in a radioactive nucleus to that of the free neutron, the change the lifetime can be separated into a change in the phase space (both in the allowed states and the different interaction energy) and the change in the coupling constants. The vector coupling turns out to be the same for all nuclei (experimentally verified at the level of 4×10^{-4} [6]). The axial-vector coupling does vary, by as much as 20-30% in different nuclei (specifically, it is 26% greater in neutron decay than in muon decay). Obviously the effect of other nucleons in the general vicinity of the decaying neutron will be less than the effect of nucleons bound into a nucleus with the nucleon. We can estimate the dilution of the effect by considering the effective range of the strong force and the density of nuclei in the neutron's environment. By definition, a neutron bound into a nucleus is entirely within strong-interaction range of other nucleons. For trapped neutrons in a bath of liquid helium, the wavelength of the neutron ($\gtrsim 100$ nm) is much larger than the interatomic distance in liquid helium (≈ 0.35 nm). Thus the neutron will overlap many helium nuclei, with a magnitude of roughly $1/\lambda^{1.5}$, or as $V^{-.5}$, where V is the volume containing the neutron. For the length scale of each helium nucleus we might use the neutron scattering length ($\approx 3.2 \times 10^{-15}$ m), or the range of a pion ($cdt \gtrsim c\hbar/2M_\pi = 7 \times 10^{-16}$ m). Taking the larger of these two, the effect of the helium

*Calculated for a series of positive ions in reference [77].

nuclei on the neutron will be roughly $2.2 \times 10^{22} \text{ cm}^{-3} \times (3.2 \times 10^{-13} \text{ cm})^3 \approx 10^{-15}$. Thus, given a shift of order 10–20% for a neutron actually bound in a nucleus, the effect of the helium nuclei in the environment of the neutron should have an effect of roughly $1\text{--}2 \times 10^{-16}$.

Appendix I

Useful Numbers

The Neutron at 1 mK

K.E.	86 neV = 1.4×10^{-26} J	$E = k_B T$
v	4 m/s	$v = \sqrt{\frac{2k_B T}{m_n}}$
λ	1000 Å	$\lambda = \frac{h}{\sqrt{2m_n k_B T}}$
B	1.4 T	$B = \frac{k_B T}{\mu_n}$
gravity	85 cm	$z = \frac{k_B T}{m_n g}$

$$\begin{aligned} m_n &= 939.56 \text{ MeV}/c^2 \\ &= 1.675 \times 10^{-27} \text{ kg} \\ \mu_n &= -1.91 \mu_N \\ &= -6.03 \times 10^{-8} \text{ eV/T} \\ &= -0.700 \text{ mK/T} \end{aligned}$$

$$\begin{aligned} k_B &= 1.38 \times 10^{-23} \text{ J/K} \\ &= 8.617 \times 10^{-5} \text{ eV/K} \\ \hbar &= 6.58 \times 10^{-22} \text{ MeV s} \\ &= 1.054 \times 10^{-34} \text{ J s} \\ &= 7.638 \times 10^{-12} \text{ K s} \end{aligned}$$

Bibliography

- [1] <http://www.ncnr.nist.gov>.
- [2] R.E. Lopez and M.S. Turner. Precision prediction for the big-bang abundance of primordial ^4He . *Physical Review D*, 59:103502, 1999.
- [3] C. Caso, et. al. (Particle Data Group). Review of particle physics. *European Physical Journal C*, 3:1–794, 1998.
- [4] J. P. Alexander, et. al. (the CLEO collaboration). First measurement of the $B \rightarrow \pi \ell \nu$ and $B \rightarrow \rho(\omega) \ell \nu$ branching fractions. *Physical Review Letters*, 77:5000–5004, 1996.
- [5] H. Leutwyler and M. Roos. *Zeitschrift für Physik C, Particles and Fields*, 25:91, 1984.
- [6] J.C. Hardy, I.S. Towner, V.T. Koslowsky, E. Hagberg, and H. Schmeing. Super-allowed 0^+ to 0^+ nuclear beta-decays: a critical survey with tests of CVC and the standard model. *Nuclear Physics A*, 509:429–460, 1990.
- [7] G. Savard, A. Galindo-Uribarri, E. Hagberg, J. C. Hardy, V. T. Koslowskya, D. C. Radford, and I. S. Towner. ^{10}C superallowed branching ratio and the cabibbo-kobayashi-maskawa matrix unitarity. *Physical Review Letters*, 74:1521–1524, 1995.
- [8] K. P. Saito and A. W. Thomas. Effect of nucleon structure variation in super-allowed fermi beta-decay. *Physics Letters B*, 363:157, 1995.
- [9] <http://pibeta.psi.ch>.

- [10] W.C. McFarlane, L.B. Auerbach, F.C. Gaille, V.L. Highland, E. Jastrzembski, R.J. Macek, F.H. Cverna, C.M. Hoffman, G.E. Hogan, R.E. Morgado, J.C. Pratt, and R.D. Werbeck. Measurements of the pion beta decay. *Physical Review D*, 32:547–565, 1985.
- [11] D.H. Wilkinson. Analysis of neutron β -decay. *Nuclear Physics*, A377:474–504, 1982.
- [12] B. Yerozolimsky, I. Kuznetsov, Yu. Mostovoy, and I. Stepanenko. Corrigendum: Corrected value of the beta-emission asymmetry in the decay of polarized neutrons measured in 1990. *Physics Letters B*, 412:240, 1997.
- [13] H. Abele, S. Baessler, D. Dubbers, J. Last, U. Mayerhofer, C. Metz, T.M. Mueller, V. Nesvizhevsky, C. Raven, O. Schaerpf, and O. Zimmer. A measurement of the beta asymmetry A in the decay of free neutrons. *Physics Letters B*, 407:212, 1997.
- [14] P. Liaud, K. Schreckenbach, R. Kossakowski, H. Nastoll, A. Bussiere, J.P. Guillaud, and L. Beck. The measurement of the beta asymmetry in the decay of polarized neutrons. *Nuclear Physics A*, A612:53, 1997.
- [15] P. Bopp, D. Dubbers, L. Hornig, E. Klemt, J. Last, H. Schutze, S.J. Freedman, and O. Scharpf. Beta-decay asymmetry of the neutron and g_A/g_V . *Physical Review Letters*, 56:919, 1986.
- [16] J. Byrne, P. G. Dawber, J. A. Spain, A. P. Williams, M. S. Dewey, D. M. Gilliam, G. L. Greene, G. P. Lamaze, R. D. Scott, J. Pauwels, R. Eykens, and A. Lamberty. Measurement of the neutron lifetime by counting trapped protons. *Physical Review Letters*, 65:289–292, 1990.
- [17] J. Byrne, P. G. Dawber, C. G. Habeck, S. J. Smidt, J. A. Spain, and A. P. Williams. A revised value for the neutron lifetime measured using a penning trap. *Europhysics Letters*, 33:187–192, 1996.
- [18] Dr. Scott Dewey, personal communication.

- [19] W. Mampe, L. N. Bondarenko, V. I. Morozov, Yu. N. Panin, and A. I. Fomin. Measuring neutron lifetime by storing ultracold neutrons and detecting inelastically scattered neutrons. *JETP Letters*, 57:82–87, 1993.
- [20] V. V. Nesvizhevskii, A. P. Serebrov, R. R. Tal'daev, A. G. Kharitonov, V. P. Alfimenkov, A. V. Streikov, and V. N. Shvetsov. Measurement of the neutron lifetime in a gravitational trap and analysis of experimental errors. *Soviet Physics JETP*, 75:405–412, 1992.
- [21] S. Arzumanov, L. Bondarenko, S. Chernyavsky, W. Drexel, A. Fomin, P. Geltenbort, V. Morozov, Yu. Panin, J. Pendlebury, and K. Schreckenbach. Neutron lifetime measured by monitored storing of ultra-cold neutrons. *Nuclear Instruments and Methods in Physics Research A*, 440:511–516, 2000.
- [22] Yu. B. Zeldovich. *Soviet Physics JETP*, 9:1389–90, 1959.
- [23] V. V. Vladimirkii. Magnetic mirrors, channels and bottles for cold neutrons. *Soviet Physics JETP*, 12:740–746, 1961.
- [24] Yu. G. Abov, V. V. Vasil'ev, V. V. Vladimirkii, and I. B. Rozhnin. Long-term storage of neutrons in a nonuniform magnetic field. *JETP Letters*, 44:472–475, 1986.
- [25] W. Paul, F. Anton, L. Paul, S. Paul, and W. Mampe. Measurement of the neutron lifetime in a magnetic storage ring. *Zeitschrift für Physik C, Particles and Fields*, 45:25–30, 1989.
- [26] N. Niehues. *Untersuchungen an einer magnetischen flasche zur speicherung von neutronen*. PhD thesis, Friedrich Wilhelm University of Bonn, 1983.
- [27] K. J. Kügler, K. Moritz, W. Paul, and U. Trinks. Nestor – a magnetic storage ring for slow neutrons. *Nuclear Instruments and Methods in Physics Research*, 228:240–258, 1985.
- [28] Yu. G. Abov, S. P. Borolëv, V. V. Vasil'ev, V. V. Vladimirkii, and E. N. Mospan. Measurement of the time of storage of ultracold neutrons in a magnetic trap. *Soviet Journal of Nuclear Physics*, 38:70–73, 1983.

- [29] J. M. Doyle and S. K. Lamoreaux. On measuring the neutron beta-decay lifetime using ultracold neutrons produced and stored in a superfluid- ^4He -filled magnetic trap. *Europhysics Letters*, 26:253–258, 1994.
- [30] R. Golub and S.K. Lamoreaux. Neutron electric-dipole moment, ultracold neutrons and polarized ^3He . *Physics Reports*, 237:1, 1994.
- [31] D.N. McKinsey and J.M. Doyle. Liquid helium and liquid neon – sensitive, low background scintillation media for the detection of low energy neutrinos. *Journal of Low Temperature Physics*, 118(3/4):153–165, 2000.
- [32] R. Golub and J.M. Pendlebury. The interaction of ultra-cold neutrons (ucn) with liquid helium and a superthermal ucn source. *Physics Letters A*, 62:337–339, 1977.
- [33] Y.V. Gott, M.S. Ioffe, and V.G. Tel’kovskii. *Nuclear Fusion*, Suppl., Pt. 3:1045, 1962.
- [34] W.D. Phillips. Laser cooling and trapping of neutral atoms. In *Laser Manipulation of Atoms and Ions*. North Holland, 1992.
- [35] J.M. Doyle, Jinha. Kim, B. Friedrich, and D. Patterson. Buffer-gas loading of atoms and molecules into a magnetic trap. *Physical Review A*, 52:R2515, 1995.
- [36] J.M. Doyle. *Energy Distribution Measurements of Magnetically Trapped Spin Polarized Atomic Hydrogen: Evaporative Cooling and Surface Sticking*. PhD thesis, Massachusetts Institute of Technology, 1991.
- [37] O. Zimmer. A method of magnetic storage of ultra-cold neutrons for a precise measurement of the neutron lifetime. *Journal of Physics G*, 26:67, 2000.
- [38] A. Steyerl, H. Nagel, F.X. Schreiber, K.A. Steinhauser, R. Gahler, W. Glaser, P. Ageronand J.M. Astruc, W. Drexel, R. Gervais, and W. Mampe. A new source of cold and ultracold neutrons. *Physics Letters A*, 116:347–352, 1986.
- [39] T.O. Brun, J.M. Carpenter, V.E. Krohn, G.R. Ringo, J.W. Cronin, T.W. Dombeck, J.W. Lynn, and S.A. Werner. Measurement of ultracold neutrons

- produced by using doppler-shifted bragg reflection at a pulsed-neutron source. *Physics Letters A*, 75(3):223–224, 1980.
- [40] R. Golub and J.M. Pendlebury. Super-thermal sources of ultra-cold neutrons. *Physics Letters A*, 53:133–135, 1975.
- [41] R. Golub, C. Jewell, P. Ageron, W. Mampe, B. Heckel, and I. Kilvington. Operation of a superthermal ultra-cold neutron source and the storage of ultra-cold neutrons in superfluid helium⁴. *Zeitschrift für Physik B*, 51:187–193, 1983.
- [42] H. Yoshiki, K. Sakai, M. Ogura, T. Kawai, Y. Masuda, T. Nakajima, T. Takayama, S. Tanaka, and A. Yamaguchi. Observation of ultracold-neutron production by 9-Å cold neutrons in superfluid helium. *Physical Review Letters*, 68:1323–1326, 1992.
- [43] R. Golub, D. Richardson, and S.K. Lamoreaux. *Ultra-Cold Neutrons*. Adam Hilger, 1991.
- [44] R. Golub. On the storage of neutrons in superfluid ⁴He. *Physics Letters A*, 72(4/5):387–390, 1979.
- [45] J.S. Adams, Y.H. Kim, R.E. Lanou, H.J. Maris, and G.M. Seidel. Scintillation and quantum evaporation generated by single monoenergetic electrons stopped in superfluid helium. *Journal of Low Temperature Physics*, 13:1121, 1998.
- [46] D.N. McKinsey, C.R. Brome, J.S. Butterworth, S.N. Dzhosyuk, P.R. Huffman, C.E.H. Mattoni, J.M. Doyle, R. Golub, and K. Habicht. Radiative decay of the metastable He₂(a³Σ_u⁺) molecule in liquid helium. *Physical Review A*, 59:200–204, 1999.
- [47] K. Habicht. *Szintillationen in flüssigem Helium – ein Detektor Für ultrakalte Neutronen*. PhD thesis, Technischen Universität Berlin, 1998.
- [48] Catalog of Luxél sub-micron foils, Friday Harbor, WA 98250.
- [49] W. M. Burton and B. A. Powell. Fluorescence of tetraphenyl-butadiene in the vacuum ultraviolet. *Applied Optics*, 12:87–89, 1973.

- [50] Frank Pobell. *Matter and Methods at Low Temperatures*. Springer-Verlag, second edition, 1965.
- [51] P.C. Hendry and P.V.E. McClintock. Continuous flow apparatus for preparing isotopically pure ^4He . *Cryogenics*, 27:131–138, 1987.
- [52] C. Mattoni. The precision measurement of the neutron lifetime using magnetically trapped neutrons: Marginally trapped neutrons and fluorescent time constants. Undergraduate thesis, Harvard University, 1995.
- [53] Dr. Kevin Coakley, personal communication.
- [54] P.R. Huffman, C.R. Brome, J.S. Butterworth, S.N. Dzhosyuk, R. Golub, S.K. Lamoreaux, C.E.H. Mattoni, D.N. McKinsey, and J.M. Doyle. Magnetically stabilized luminescent excitations in hexagonal boron nitride. *submitted to Physical Review B*, 2000.
- [55] C.E.H. Mattoni, C.R. Brome, S.N. Dzhosyuk, R. Golub, K. Habicht, P.R. Huffman, D.N. McKinsey, and J.M. Doyle. Neutron induced luminescence of selected materials at low temperature. *to be submitted to Nuclear Instrumentation and Methods A*, 2000.
- [56] J.S. Butterworth, C.R. Brome, P.R. Huffman, C.E.H. Mattoni, D.N. McKinsey, and J.M. Doyle. A removeable cryogenic window for transmission of light and neutrons. *Review of Scientific Instruments*, 69:3998, 1998.
- [57] Dr. Jeff Nico, personal communication.
- [58] W.G. Driscoll and W. Vaughan, editors. *Handbook of Optics*, pages 7–62–7–63. McGraw Hill, first edition, 1978.
- [59] L. Duband, L. Hui, and A. Lange. Thermal isolation of large loads at low temperature using kevlar rope. *Cryogenics*, 33(6):643–647, 1993.
- [60] H.S. Sommers, J.G. Dash, and L. Goldstein. Transmission of slow neutrons by liquid helium. *Physical Review*, 97:855, 1955.
- [61] M. Locatelli, D. Arnand, and M. Routin. Thermal conductivity of some insulating materials below 1K. *Cryogenics*, 16:374, 1976.

- [62] R. Winston. Light collection within the framework of geometrical optics. *Journal of the Optical Society of America*, 60:245, 1970.
- [63] D.N. McKinsey, C.R. Brome, J.S. Butterworth, R. Golub, K. Habicht, P.R. Huffman, S.K. Lamoreaux, C.E.H. Mattoni, and J.M. Doyle and. Fluorescence efficiencies of thin scintillating films in the extreme ultraviolet spectral region. *Nuclear Instruments and Methods B*, 132:351, 1997.
- [64] D.N. McKinsey, C.R. Brome, J.S. Butterworth, S.N. Dzhosyuk, P.R. Huffman, C.E.H. Mattoni, J.M. Doyle, R. Golub, and K. Habicht. Detector optimization for an experiment to measure the neutron lifetime using magnetically trapped neutrons. *to be submitted to Nuclear Instrumentation and Methods A*, 2000.
- [65] D.C.Camp, C. Gatrousis, and L.A. Maynard. *Nuclear Instrumentation and Methods*, 117:189, 1974.
- [66] D.W. Müller, G. Best, J. Jackson, and J. Singletary. Afterpulsing in photomultipliers. *Nucleonics*, 10, 6:53, 1952.
- [67] G.A. Morton, H.M. Smith, and R. Wasserman. Afterpulses in photomultipliers. *IEEE Transactions on Nuclear Science*, page 443, 1967.
- [68] Philip R. Bevington and D. Keith Robinson. *Data Reduction and Error Analysis for the Physical Sciences*. McGraw Hill, Inc., 1992.
- [69] P.R. Huffman, C.R. Brome, J.S. Butterworth, K.J. Coakley, M.S. Dewey, S.N. Dzhosyuk, R. Golub, G.L. Greene, K. Habicht, S.K. Lamoreaux, C.E.H. Mattoni, D.N. McKinsey, F.E. Wietfeldt, and J.M. Doyle. Magnetic trapping of neutrons. *Nature*, 403:62–64, 2000.
- [70] <http://www.ill.fr>.
- [71] <http://www.doylegroup.harvard.edu/~neutron>.
- [72] W. Henning, W. Kutschera, M. Paul, R.K. Smither, E.J. Stephenson, and J.L. Yntema. Accelerator mass spectrometry and radioisotope detection at the Argonne FN tandem facility. *Nuclear Instruments and Methods*, 184(1):1247–1268, 1981.

- [73] C.Y. Chen, Y.M. Li, K. Bailey, T.P. O'Connor, L. Young, and Z.T. Lu. Ultrasensitive isotope trace analyses with a magneto-optical trap. *Science*, 286(5442):1139–1141, 1999.
- [74] I.S. Towner and J.C. Hardy. The current status of V_{ud} . *nucl-th/9809087* on xxx.lanl.gov preprint archive, 1998.
- [75] E.L. Surkov, J.T.M. Walraven, and G.V. Shlyapnikov. Collisionless motion of neutron particles in magnetostatic traps. *Physical Review A*, 49(6):4778–4786, 1994.
- [76] W.T. Sommer. Liquid helium as a barrier to electrons. *Physical Review Letters*, 12:271, 1964.
- [77] M. Beau, H. Günther, G. zu Putlitz, and B. Tabbert. Atoms and ions in superfluid helium. *Zeitschrift für Physik B*, 101:253–262, 1996.