

Magnetic Trapping of Ultracold Neutrons Produced Using a Monochromatic Cold Neutron Beam

A thesis presented

by

Carlo Egon Heinrich Mattoni

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

April 2002

©2002 - Carlo Egon Heinrich Mattoni

All rights reserved.

To my parents

Thesis advisor

Author

John Morrissey Doyle

Carlo Egon Heinrich Mattoni

Magnetic Trapping of Ultracold Neutrons Produced Using a Monochromatic Cold Neutron Beam

Abstract

Progress toward a measurement of the neutron lifetime using magnetically trapped ultracold neutrons (UCN) is reported. A monochromatic beam of 8.9 Å neutrons is produced by diffraction of a beam of cold neutrons from a monochromator. The design, fabrication, and characterization of this novel, long-wavelength monochromator, which achieved $> 80\%$ reflectivity for 8.9 Å neutrons, is detailed. Through the superthermal process, a non-equilibrium population of near-rest UCN is produced by scattering the 8.9 Å neutrons from phonons in a superfluid helium bath. The helium bath sits within the bore of an Ioffe-type magnetic trap, which confines UCN with energies less than the trap depth. Decay of the trapped neutrons results in scintillations within the helium which are frequency downconverted using an organic fluor. Visible light is then extracted from the trapping region and detected at room temperature. Magnetic trapping of UCN produced from a monochromatic 8.9 Å neutron beam is established at the 6.1σ level.

Contents

Abstract	iv
Table of Contents	v
List of Figures	vii
List of Tables	xii
Acknowledgements	xiv
1 Introduction	1
1.1 Neutron Beta Decay	2
1.2 CKM Unitarity	6
1.3 Big Bang Nucleosynthesis	11
1.4 Previous Measurements of τ_n	16
1.5 τ_n from Magnetically Trapped UCN	22
2 Materials Selection	39
2.1 Overview	39
2.2 Neutron Activation	42
2.3 Second Order Activation Processes	45
2.4 Neutron Induced Luminescence	46
3 Monochromator	71
3.1 Introduction	71
3.2 Monochromator Design	74
3.3 Long Wavelength Monochromator Materials	77
3.4 Fabrication of IHOPG	84
3.5 Sample Characterization	90
3.6 Assembled Monochromator	97
3.7 Time of Flight Measurements	101
3.8 Neutron Wavelength Filtering	109
3.9 Beam Characterization	112
3.10 Conclusions	118
4 Experimental Apparatus	119
4.1 Neutron Beam	119

4.2	Cryogenic Apparatus	125
4.3	Magnetic Trap	148
4.4	Detection System	174
4.5	Shielding	183
5	Data Acquisition and Analysis	189
5.1	Data Acquisition Hardware	189
5.2	Data Acquisition Software	198
5.3	Data Analysis Software	204
5.4	Data Analysis Procedure	216
6	Results and Discussion	219
6.1	Data Sets	221
6.2	Backgrounds	227
6.3	Evidence of Neutron Trapping	238
6.4	Higher Thresholds	242
6.5	Temperature Variation of the Trap Lifetime	244
6.6	Calibration of the Detection Efficiency	246
6.7	UCN Production Rate	251
6.8	Non-8.9Å production	254
7	Conclusions and Future Directions	255
7.1	Magnetic Trapping of UCN	255
7.2	Time Dependent Background	256
7.3	Separate Production and Detection Regions	258
7.4	Permanent Monochromator Installation	263
7.5	Source Improvements	265
7.6	Increased Detection Efficiency	266
7.7	Larger, Deeper Magnetic Trap	267
7.8	Prospects for a Lifetime Measurement	269
A	Activation	273
B	Dewar Boiloff	281
C	Neutron Beam Simulations	287
C.1	Diffraction Theory	287
C.2	Monte Carlo Simulations	295
D	Monochromator Materials Selection	319
D.1	Diffractometers	319
D.2	Fluorophlogopite Characteristics	321
D.3	Intercalated Graphite Characteristics	324

E The Parser	329
F Detection Efficiency Calculation	341
Bibliography	347

List of Figures

1.1	Feynman diagrams of neutron beta-decay	3
1.2	Experimental constraints on λ and $ V_{ud} $	7
1.3	Beam Lifetime Experiment	17
1.4	Ioffe configuration concept	25
1.5	Dispersion curves of the free neutron and of superfluid helium	27
1.6	Overview of helium scintillation	31
2.1	The materials interacting with the neutron beam	42
2.2	Luminescence test apparatus	50
2.3	Signal observed with the aluminum can up and down and no sample .	58
2.4	BN and LiF luminescence signals	59
2.5	BN luminescence signal in coincidence mode	60
2.6	BN and BN encapsulated in graphite luminescence signals	61
2.7	boron carbide and UVT acrylic luminescence signals	62
2.8	LiF luminescence in two experimental setups	63
2.9	Luminescence of TPB evaporated on GoreTex	64
2.10	Luminescence of the boron/lithium glass	65
2.11	Luminescence of boron oxide	66
3.1	Structure of Graphite	80
3.2	Staging of graphite intercalation compounds	81
3.3	Simulated output spectra of fluorophlogopite, stage 1, and stage 2 IHOPG	83
3.4	Nominal temperature dependence of stage formation in potassium in- tercalated graphite	85
3.5	An intercalation ampule	85
3.6	Vacuum system used for loading potassium into intercalation cells . .	86
3.7	Intercalation setup	88
3.8	Position dependence of furnace temperature	89
3.9	Stage purity of IHOPG samples	91
3.10	Pinwheel sample holder	93
3.11	Vertical mosaic before compression	94
3.12	Photograph of IHOPG showing bend	94
3.13	Compression rig	95

3.14	Vertical mosaic after compression	96
3.15	Sample tiling of the monochromator	99
3.16	Monochromator Holder	100
3.17	Monochromator assembly setup	100
3.18	Time of flight measurement setup	103
3.19	TOF spectra with varying gate and delay output width	105
3.20	TOF spectra at the center of the monochromator	107
3.21	TOF spectra in transmission and reflection	109
3.22	Transmission spectra of graphite and beryllium filters	111
3.23	TOF spectra with the $\lambda/2$ filter	113
3.24	TOF spectra of the assembled beamline	114
3.25	TOF spectrum seven months after monochromator assembly	115
3.26	Monochromatic beam flux vs. position	117
4.1	The reactor and coldsource	120
4.2	Neutron beam layout	121
4.3	Image of the monochromatic neutron beam	124
4.4	The old apparatus	126
4.5	The old IVC	127
4.6	The new apparatus	128
4.7	Old and new bypass lines for ultrapure helium	131
4.8	Superfluid helium-filled heat link	133
4.9	The new buffer cell	135
4.10	A teflon window seals the beam entrance side of the cell	137
4.11	The snout	138
4.12	Kevlar strings for positioning the cell	139
4.13	Thermal anchoring of the detection system	141
4.14	Direct cooling of the 77 K shield	143
4.15	Wires embedded in the 77 K lightguide	146
4.16	The 4K window	147
4.17	Ioffe trap realization	149
4.18	Magnetic field profile of our trap	150
4.19	Prestressing scheme for quadrupole assembly	153
4.20	Quench detection circuit	154
4.21	Trigger circuit for the quench protection system	155
4.22	Quench protection circuit	156
4.23	Quadrupole and octupole geometries	159
4.24	Octupole training curve	160
4.25	Coil voltages during a quench	162
4.26	Octupole trap first assembly training curve	163
4.27	Solenoid assemblies training curve	165
4.28	Octupole trap second assembly training curves	166

4.29	Efficiency of the quench protection system	169
4.30	$R(t)$ after quenches at several currents	170
4.31	Quadrupole trap training curve	172
4.32	Quadrupole trap second training curve	175
4.33	The detection cell	176
4.34	The detection system optics	179
4.35	Spillover in a Burle 8850	180
4.36	Magnetic shielding for PMTs	182
4.37	Neutron shielding	184
5.1	Overview of the data acquisition system	190
5.2	Block diagram of NIM trigger logic	193
5.3	A typical digitized PMT pulse	195
5.4	Time sequence of a data run	199
5.5	Count rate during the first 100 s of a run	201
5.6	Typical pulse height spectrum of a Burle 8854	205
5.7	Pulse height spectrum from integrator	207
5.8	Operation of the parser on a digitized pulse	208
5.9	A typical “average event”	209
5.10	Distorted and corrected average event from the slave card	210
5.11	Operations for correcting misordered slave card bytes	211
5.12	Time dependence of backgroundave	213
5.13	Pulse height spectrum of parsed events	216
6.1	Pulse height spectra with cell empty and full	229
6.2	Negative Cold1 runs cut at 1+ photoelectron	232
6.3	Singles rate for PMT β	233
6.4	Negative Cold1 runs cut at 2+ photoelectrons	235
6.5	Cold1 and ^3He data cut at 1+ photoelectron	239
6.6	Cold - ^3He data cut at 1+ photoelectron	240
6.7	^3He data cut at 2+ photoelectrons	241
6.8	Cold1 data cut at 2+ photoelectrons	242
6.9	Decay Rate as a function of temperature	245
6.10	Cold2 data cut at 2+ photoelectrons	246
6.11	Pulse height spectrum from ^{113}Sn source	248
6.12	Beta decay spectra of the neutron and ^{113}Sn	249
6.13	Weighted decay amplitude	251
7.1	Apparatus with separate production and detection regions	259
7.2	NG-6 layout with permanent monochromator installation	264
B.1	Boiloff rate above and below transfer tube	283

C.1	$ \vec{G} /2$ sphere	289
C.2	Crystal orientations \hat{n}_M in a mosaic crystal	290
C.3	Comparison of measured and simulated neutron spectra for NG-6	297
C.4	Momentum of a randomly chosen neutron in the (x, y, z) coordinate system.	300
C.5	Position and orientation of the monochromator with respect to the end of NG-6.	301
C.6	Phase space diagram of reflection in the (x'', y'', z'') coordinate system	302
C.7	Diffraction of a neutron from a crystal mosaic	305
C.8	The primed and unprimed coordinate systems	306
C.9	Geometry of the trapping apparatus	307
C.10	Monte carlo simulated reflectivity compared to implied reflectivity from Z1	308
C.11	Simulated rocking curve of stage 1 IHOPG	310
C.12	Simulated vs. Measured transmission spectra of sample Z3	311
D.1	Rocking curve for “flat” fluorophlogopite sample	322
D.2	θ - ω contour plot for 1.3 cm thick fluorophlogopite	324
D.3	Rocking curve for stage 2 GIC sample Z1	325
D.4	8.9 Å transmission through stage 2 GIC sample Z1	326
D.5	Rocking curve for stage 1 GIC sample Z3	327
D.6	TOF Spectrum of transmission through sample Z3	328

List of Tables

2.1	Stable isotopes with large neutron capture cross sections	41
2.2	Manufacturer's elemental analysis for Beryllium windows	44
3.1	Measured characteristics of 15 IHOPG samples	98
4.1	Quench protection efficiency	168
4.2	Quench protection efficiency for quadrupole	173
6.1	Summary of the data sets	222
6.2	Fraction of constant background due to helium	230
6.3	Activation fit parameters vs. temperature	236
6.4	Activation fit parameters vs. threshold	237
6.5	Extracted fit parameters at various thresholds	243
6.6	Extracted fit parameters vs. temperature	245
6.7	Detection efficiency as a function of threshold	250
7.1	Lifetime measurement accuracy for possible future experimental setups	270
A.1	Neutron activation properties of selected stable isotopes	275
B.1	Helium boiloff due to the transfer tube	282
C.1	Crystallographic parameters of 8.9 Åmonochromator materials	307

Acknowledgements

The completion of any large, long-term project, such as a doctoral thesis, cannot be accomplished without the assistance of a great many people. I gratefully acknowledge the help I have received from all of the following, as well as those who I am sure to have mistakenly overlooked. Their contribution has been not only to the technical and scientific aspects of my graduate career, but also to the other elements of life, which have made these years ones I will look back on fondly.

My thanks, first, must be to those who I have worked alongside on the neutron trapping experiment: to Dan McKinsey, my colleague and “twin” from the first to the last, without whose creativity and dedication this work would not have been a success; to Sergei Dzhosyuk, who has been a constant example not only of competence and hard work, but also of how to have fun regardless of the circumstances; to Liang Yang, whose contributions are only increased by the good nature with which they are accomplished; to Bob Golub, who has taught me more than I ever wanted to know about neutrons, while exemplifying the merit of broad interests and critical thinking; to Paul Huffman, my mentor and friend, who has taught me everything I know about how to make things work; and to my advisor, John Doyle, from whom I have learned so much more than just physics — when I walked through the door of a brand new junior faculty member, I had no idea of the journey I was embarking upon, but it has been one of growth from youth to adulthood as much as from an undergraduate to a PhD, and I appreciate your guidance throughout. I count as the highest point of my graduate school years the opportunity to work with such fine people.

The group at NIST has also been a pleasure to work with. Working as a graduate student at a research lab rather than in a university environment has its advantages

and disadvantages. As with my Harvard colleagues, interacting both socially and professionally with good people has made the overall experience an extremely positive one. In particular, Alan Thompson, in addition to building the final DAQ, has been a never-failing source of hardware and software advice, and will always be my idol as the ur-computer-nerd. While not working directly on this project, Jeff Nico, whose contributions too often go unpraised, has been an example of the value of dedication and competence, not to mention a little good-hearted sarcasm and cynicism.

Many others have contributed in ways too numerous to list, including, but by no means limited to, Clint Brome, Kyle Alvine, James Butterworth, Bob Michniak, Fred Weitfeldt, Klaus Habicht, Katerina Korobkina, and Kevin Coakley. On the development of the 8.9 Å monochromator, I must thank both Hartmut Zabel and Dan Neumann for their invaluable advice and fruitful suggestions.

I have had the great fortune of having colleagues who have also been good friends. Nevertheless, I would be remiss not to mention the support of some of my other friends, who have helped make life outside the lab fun: Roxanne, Venk, Joe, Dean, Koji, Eric, and Matthew and Sarah and the whole Furner clan.

I don't know how to offer enough thanks to my parents. Their example and support span my entire life and cannot be simply summarized. Without a doubt their encouragement sustained me through the discouraging times of graduate school.

Saving the best for last, I offer my undying gratitude to my wife, Becca. From the first weeks of graduate school, when we began dating, you have been by my side the whole way. I couldn't have made it through the tough times nor would I have enjoyed the good times half as much without you. Your love and kindness have been my constant strength and inspiration.

Chapter 1

Introduction

Neutron decay is an essential feature of the universe. At the known most fundamental level, the neutron decays because the down quark is more massive than the up quark. The neutron, which consists of two down and one up quarks, decays into a proton, containing two up and one down quarks. Were the up quark to be more massive than the down quark, then the proton would decay into the neutron and the universe around us would be radically different, without stable hydrogen. Without hydrogen, stars, planets, and life would be significantly different, if they existed at all.

The rate at which the neutron decays was a crucial determinant in the early minutes of the universe, as the elements began to form. If Nature had bestowed on us a larger rate, less helium and more hydrogen would have resulted. Conversely, if the neutron decayed slower, the early universe could have been dominated by helium rather than hydrogen. Stellar formation and burning would have been decidedly different under either circumstance. While it is perhaps a bit of a stretch to claim that the sun shines because the neutron lifetime is what it is, it is not hard to imagine that it might not shine at all, and that, moreover, we would not exist to bask in its

light, were the neutron lifetime significantly different.

The sun has shined, and hydrogen has been stable for billions of years, despite the fact that the neutron was unknown to humanity prior to 1932 [1]. The business of physics, however, is to understand the nature and laws of the universe. This understanding is the result of a sometimes clunky dance between theory and experiment. Experiments demonstrate a phenomenon, one or more theoretical models tries to explain that phenomenon, and further experiments probe the implications of those models to determine if they are consistent with reality. In the shuffle between theory and experiment, the neutron lifetime still has much to tell us. As the simplest system which decays by the weak interaction, the neutron is an important probe of the Standard Model of electroweak physics. Together with other measurements, the neutron lifetime can test the unitarity of the Cabibbo-Kobayashi-Maskawa (CKM) quark-mixing matrix. Measurement of the neutron lifetime is also an important input to the theory of Big Bang Nucleosynthesis (BBN), whereby we may understand the formation of elements in the primordial universe.

1.1 Neutron Beta Decay

The neutron decays into a proton, an electron, and an electron anti-neutrino with the release of approximately 782 keV of energy through the reaction:

$$n \rightarrow p^+ + e^- + \bar{\nu}_e + (\sim 782 \text{ keV}) \quad (1.1)$$

This decay process is now understood to result from the operation of the weak interaction upon the constituent quarks of the neutron. Through the exchange of a virtual W boson with an electron and anti-neutrino, one of the neutron's down quarks becomes an up quark (see figure 1.1A). At low energies, beta decay may be adequately




Figure 1.1: Feynman diagrams of two models of neutron beta decay. A) At the quark level, the electroweak interaction is mediated by the exchange of a virtual W boson. B) At low neutron energies, a four particle point interaction adequately describes the decay process.

described using a phenomenological four particle point interaction (see figure 1.1B).

In this four fermion point interaction model, due originally to Fermi at a time when the quark structure was unknown [2], the weak interaction hamiltonian can be written as

$$H = \frac{G_F}{\sqrt{2}} J_\mu^W J_\mu^{W\dagger} \quad (1.2)$$

where J_μ^W is the sum of leptonic and hadronic currents, j_μ^l and j_μ^h , respectively [3].

The leptonic current

$$j_\mu^l = \bar{\Psi}_e \gamma_\mu (1 - \gamma_5) \Psi_{\nu_e} \quad (1.3)$$

couples leptonic isospin doublets such as $\begin{pmatrix} \nu_e \\ e^- \end{pmatrix}$ directly.

In the hadronic sector, the interaction current becomes much more complicated. The simplest hadronic interaction current, that for quark interactions, can be written

as

$$j_\mu^h = \bar{\Psi}_{d'} \gamma_\mu (1 - \gamma_5) \Psi_u \quad (1.4)$$

where we notice that the interaction couples quark doublets of the form $\begin{pmatrix} u \\ d' \end{pmatrix}$. The primed quark states, which are the eigenstates of the weak interaction, are not equivalent to the unprimed “mass” eigenstates (the eigenstates of the strong interaction). The weak force and mass eigenstates are related by the Cabibbo-Kobayashi-Maskawa (CKM) quark mixing matrix:

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

where the letters used represent the various quarks (d for down, s for strange, and so on). Since free quarks have extremely short lifetimes, this model is not particularly relevant for real calculations. All hadronic interactions of interest occur for composite hadrons, which consist of multiple quarks bound together by the strong force. The interaction of the strong force between quarks modifies the weak interaction vertices so that the weak hadronic current takes on the more general form

$$j_\mu^h = \bar{\Psi}_1 \gamma_\mu (c_v + c_a \gamma_5) \Psi_2 \quad (1.6)$$

where c_v and c_a are the relative strengths of the new vector and axial-vector components of the current. In general, scalar and tensor currents may also be present, however symmetry arguments require them to be zero and they are therefore excluded from the Standard Model (SM), although certain extensions to the SM include them. The conserved vector current (CVC) hypothesis states that the vector current is the same for all interactions, and hence $c_v = 1$. No such constraint limits the value of c_a , so it must be determined experimentally.

It is often convenient to discuss the vector and axial vector form factors, g_v and g_a rather than c_v and c_a . The form factors are defined so that

$$g_v = G_F |V_{ud}| c_v \quad (1.7)$$

and likewise for g_a . The dimensionless parameter λ is defined as g_a/g_v ($= c_a$ if CVC holds). The overall weak interaction coupling parameter $G_F = (1.16637 \pm 0.00001) \times 10^{-5} \text{GeV}^{-2}$ can be accurately calculated with the measured muon lifetime as the limiting experimental input, since muon decay is a purely leptonic weak decay and hence is not subject to uncertainties of the hadronic sector.

With the complete weak interaction Hamiltonian, one can compute the neutron lifetime. From Fermi's Golden Rule, the rate of an interaction is found by integrating the matrix element of the hamiltonian over the phase space of the final states. The differential decay rate in a V-A theory, for a fixed initial neutron spin, averaging over the spins of the products can be written in the well-known form given by Jackson, Treiman, and Wyld [4] as:

$$\frac{d\Gamma}{dE_e d\Omega_e d\Omega_\nu} \propto (g_v^2 + 3g_a^2) F(E_e) \left(1 + a \frac{\vec{p}_e \cdot \vec{p}_\nu}{E_e E_\nu} + A \frac{\vec{\sigma}_n \cdot \vec{p}_e}{E_e} + B \frac{\vec{\sigma}_n \cdot \vec{p}_\nu}{E_\nu} \right) \quad (1.8)$$

where $F(E_e)$ is the phase space factor due to the electron energy. The correlation coefficients, which form experimental observables, have the following dependence on λ :

$$a = \frac{1 - \lambda^2}{1 + 3\lambda^2} \quad (1.9)$$

$$A = -2 \frac{\lambda^2 + \lambda}{1 + 3\lambda^2} \quad (1.10)$$

$$B = 2 \frac{\lambda^2 - \lambda}{1 + 3\lambda^2} \quad (1.11)$$

One can also see that the neutron lifetime is a measurement of the quantity $g_v^2(1+3\lambda^2)$. Thus to determine g_v and g_a one needs to measure both the neutron lifetime and any

one of the other decay parameters. The parameter B is relatively insensitive to λ and the parameter a is extremely difficult to measure due to its dependence on the neutrino momentum and energy, which cannot be observed directly. Hence, the most sensitive limits on λ come from measurements of the neutron asymmetry coefficient, A . The current accepted value of $\lambda = -1.2670 \pm 0.003$ is determined from four measurements which disagree by an amount inconsistent with their individual quoted errors [5]. The two outlying measurements are $\lambda = -1.2594 \pm 0.0038$ [6] and $\lambda = -1.274 \pm 0.003$ [7], which differ from each other by almost 10 standard deviations¹.

1.2 CKM Unitarity

An important purpose for measuring the neutron lifetime with improved accuracy is to test the theory of CKM unitarity. The CKM Matrix (see equation 1.5) is simply a rotation of the quark eigenstates. As a rotation it must be unitary. Non-unitarity of the CKM matrix would be strong evidence for physics beyond the Standard Model, such as a fourth generation of quarks.

If the CKM matrix is unitary, the sum of the squares of the first row matrix elements must be one:

$$|V_{ud}|^2 + |V_{us}|^2 + |V_{ub}|^2 = 1 \quad . \quad (1.12)$$

The off-diagonal matrix elements are measured in high energy physics experiments and are much smaller than V_{ud} . The matrix element $|V_{us}| = 0.2196 \pm 0.0023$ is determined from studies of K decays [5]. From analysis of rare B meson decays, the matrix element $V_{ub} = 0.0036 \pm 0.0010$ can be determined [5]. Plugging these values

¹A second run of the apparatus in [7], with lower backgrounds, has recently been published [8]. The new result of $\lambda = -1.274 \pm 0.0021$ is not considered in the analysis herein.




Figure 1.2: Current experimental constraints on the values of λ and $|V_{ud}|$ are shown. Experimental limits are derived from the following sources: neutron lifetime [5], Nuclear $0^+ \rightarrow 0^+$ beta decays [9], CKM Unitarity (assumes CVC, CKM unitarity, and values of $|V_{us}|$ and $|V_{ub}|$ from [5]), A_1 and A_2 are measurements of the neutron beta-decay correlation A from [6] and [7], respectively.

into (1.12) we can rewrite the unitarity condition in terms of a constraint on V_{ud} as

$$|V_{ud}| = 0.9756 \pm 0.0005 \quad . \quad (1.13)$$

Experimental constraints on the value of $|V_{ud}|$ can also be made by studying neutron beta decay, superallowed $0^+ \rightarrow 0^+$ beta decays, or rare pion decays [9]. The current experimental constraints on V_{ud} and λ are shown in figure 1.2.

It is possible to determine $|V_{ud}|$ by studying the rare pion decay $\pi^+ \rightarrow \pi^0 e^+ \nu_e$. This $0^- \rightarrow 0^-$ decay has no axial-vector component and requires no nuclear-structure dependent theoretical corrections. Unfortunately the extremely small branching ratio for the decay, $\sim 10^{-8}$, makes measurements in this system extremely difficult

experimentally. The current limit on $|V_{ud}|$ using pions is

$$|V_{ud}| = 0.9670 \pm 0.0161 \quad . \quad (1.14)$$

A measurement is underway with the goal of improving the sensitivity of these pion measurements by about a factor of ten [10].

The study of superallowed $0^+ \rightarrow 0^+$ Fermi transitions (SFT) currently gives the tightest experimental constraints on $|V_{ud}|$. For these types of decays, $c_a = 0$ and $c_v = 1$ (by CVC). It is thus possible to determine $|V_{ud}|$ from any SFT by measuring the half-life of the state, $t_{1/2}$, the branching ratio of the SFT, and f , the phase space factor integral. The ft value, or partial half-life of the SFT, for different isotopes can be compared if nucleus-dependent theoretical corrections are applied. The relationship between ft for any nucleus and the nucleus-independent term $\mathcal{F}t$ is written as [9]:

$$\mathcal{F}t \equiv ft(1 + \delta_R)(1 - \delta_C) \quad (1.15)$$

where δ_R is the nucleus-dependent, or “outer,” radiative correction, and δ_C is the isospin symmetry-breaking (Coulomb) correction. The outer radiative corrections are $\delta_R \sim (1.4 \pm 0.05)\%$, while the Coulomb corrections are $\delta_C \sim (0.5 \pm 0.03)\%$ [9].

The $\mathcal{F}t$ value has been determined at the 10^{-3} level for 9 SFTs. The 9 $\mathcal{F}t$ values are statistically consistent, with $\chi^2/\nu = 1.1$ when the distribution is fit to a single mean value [9]. From the mean, $\bar{\mathcal{F}}t$, of the measured $\mathcal{F}t$ values, g_v , and hence $|V_{ud}|$ can be determined from:

$$\bar{\mathcal{F}}t = \frac{K}{2g_v^2(1 + \Delta_V^R)} \quad (1.16)$$

where $K = 2\pi^3 \ln 2 \hbar (\hbar c)^6 / (m_e c^2)^5$ is a constant and $\Delta_V^R = (2.40 \pm 0.08)\%$ is the nucleus-independent, or inner, radiative correction [9].

An analysis using the entire data set of 9 SFTs yields [9]

$$|V_{ud}| = 0.9740 \pm 0.0005 \quad . \quad (1.17)$$

This value of $|V_{ud}|$ is inconsistent with the requirement of CKM unitarity at the 2.8σ level. The largest contributors to the error bars on $|V_{ud}|$ determined from SFTs are from Δ_V^R (± 0.0004) and δ_C (± 0.0003) [9]. In order to shift $|V_{ud}|$ so that it would be consistent with unitarity, significant shifts in the radiative corrections would be necessary: Δ_V^R would have to decrease by 0.3%, or all values of δ_R would have to decrease by 0.3%, or all values of δ_C would have to increase by 0.3% [9].

Since the leading order terms in the radiative corrections are well understood, most attention has focussed on uncertainties in the calculation of the Coulomb correction [9]. In particular, a recent proposal by Saito and Thomas [11] may point the way towards an increase of δ_C by 0.15 to 0.2%, hence making up most of the amount needed to restore consistency with unitarity. It has even been argued [5] that the quoted error on $|V_{ud}|$ determined from SFTs, which does not include the correction proposed in [11], is thus an underestimate, and should be doubled due to the uncertainties in the Coulomb correction. Due to the difficulty in improving confidence in the calculated value of δ_C , it will likely prevent the uncertainty in the determination of $|V_{ud}|$ from SFTs from improving beyond the present limit.

In order to improve the determination of $|V_{ud}|$ beyond the precision obtained in SFT systems, one can study neutron beta decay. An attraction of the neutron system is the relatively more robust understanding of the necessary corrections. The nucleus dependent radiative correction δ_R has been calculated to the 10^{-5} level for the neutron [12], and there is no coulomb correction [9]. Hence the limiting theoretical uncertainty on $|V_{ud}|$ determined from neutron decay is simply that due to the inner radiative correction, Δ_V^R (± 0.0004). The inner radiative correction, Δ_V^R is calculated as [9]:

$$\Delta_V^R = \frac{\alpha}{2\pi} [4 \ln(m_Z/m_p) + \ln(m_p/m_A) + 2C_{Born}] + O(\alpha^2), \quad (1.18)$$

where m_Z is the mass of the Z -boson, m_p is the mass of the proton, m_A is the mass of the a_1 -meson, and C_{Born} is the $O(\alpha)$ axial-vector photonic contribution for a single nucleon. Currently the uncertainty in the calculation of Δ_V^R is dominated by the uncertainty in the cutoff mass m_A , which is determined by the details of the strong interaction at low energy, and is hence restricted only weakly, to the range 400–1600 MeV [13]. Improvements in the calculation of Δ_V^R will rely upon restricting this mass range further.

The overall radiative corrections to the neutron decay rate have been calculated at the 10^{-4} level [12, 14]. At this level, the neutron beta decay rate can be written [14] as

$$\Gamma = 0.1897|V_{ud}|^2(1 + 3\lambda^2)(1 + 0.0739 \pm 0.0008) \times 10^{-3} \text{ s}^{-1} \quad (1.19)$$

The CKM matrix element $|V_{ud}|$ can thus be determined from the neutron system by measuring the neutron lifetime ($1/\Gamma$) and λ , the ratio of g_a/g_v . The ratio λ is most precisely determined from measurements of the electron-neutron spin-asymmetry coefficient, A . Currently the limiting error on $|V_{ud}|$ measured with neutrons comes from A . As is discussed above, the various measurements of A are not consistent with each other to within the stated errors. Using the accepted value for τ_n and a weighted average of the 4 most precise measurements of A yields [5]

$$|V_{ud}| = 0.9728 \pm 0.0012 \quad . \quad (1.20)$$

This result is 2.3σ discrepant from unitarity. A new measurement of A using UCN is currently underway which aims to reduce the error on A by a factor of 4 initially, with further improvements possible[15].

The combination of an improved measurement of A with improved measurements of the neutron lifetime could allow determination of $|V_{ud}|$ with errors as small as the ± 0.0004 set by current theoretical calculations of the inner radiative correction,

Δ_V^R [9, 14]. Determination of $|V_{ud}|$ at that level in the neutron system would be of comparable or greater precision than the determination of $|V_{ud}|$ from SFTs, and would be based upon a more robust theoretical foundation. Therefore, improved understanding of CKM unitarity is possible by studying neutron beta-decay, including measuring τ_n with increased precision.

1.3 Big Bang Nucleosynthesis

The theory of Big Bang Nucleosynthesis (BBN) describes the formation of the light nuclei in the first few minutes following the Big Bang [16]. The story of the early universe, according to the standard theory of Big Bang cosmology, follows a monotonic path of expansion and decreasing energy at increasing times. Until the the universe cools to ~ 0.1 MeV ($t \sim 100$ s), significant nucleosynthesis does not occur. The relative abundance of the isotopes produced in the early universe depends on physics at even earlier times, back to around one second after the Big Bang. The BBN theory is thus a window into the earliest moments of time. The neutron lifetime is an important input to the theory of BBN and dominates the theoretical uncertainty in the abundance of ${}^4\text{He}$ produced through primordial nucleosynthesis [17]. For reasons described below, the ${}^4\text{He}$ mass fraction depends linearly on the neutron lifetime.

At energies more than about 1 MeV, or times earlier than about 1 s, neutrons and protons interchange freely. Interchange occurs through weak interaction processes such as

$$p + \bar{\nu}_e \leftrightarrow n + e^+ \quad (1.21)$$

and

$$n + \nu_e \leftrightarrow p + e^- \quad . \quad (1.22)$$

The rates for all such interchange reactions can be calculated, using Fermi's Golden Rule, in much the same manner as the neutron decay rate is calculated, by integrating the square of the matrix element over the available phase space. Since these phase space integrals can be calculated to very high precision, and the matrix element

$$|M|^2 \propto G_F^2(1 + 3g_a^2) \quad (1.23)$$

is the same for all neutron-proton interchange reactions (including neutron beta decay), then the knowledge of all of these interchange rates is limited by the knowledge of the matrix element. The matrix element can be determined experimentally by measuring the neutron lifetime. So long as the total interchange rate ($\propto T^5/\tau_n$) remains large compared with the expansion rate ($\propto T^2$), the neutrons and protons remain in thermal equilibrium and their relative abundance, x_n is given by the Boltzmann factor:

$$x_n = \frac{n}{p} = e^{-Q/T} \quad (1.24)$$

where $Q \equiv m_n - m_p = 1.3$ MeV is the mass difference between the neutron and proton. At high temperatures, the numbers of neutrons and protons are almost equal, with the relative number of neutrons falling as the temperature decreases.

At a temperature of about 0.8 MeV, occurring at about $t = 1$ s, the universe has expanded sufficiently that interchange reactions do not happen often enough to maintain thermal equilibrium, and the neutron to proton ratio “freezes out” at a value of about 1/6. The exact temperature at which freezeout occurs, and hence the precise value of x_n , is set by the expansion rate of the universe and the nuclear interchange reaction rates. Since the nuclear interchange reaction rates can all be computed from the neutron lifetime, calculation of the freezeout temperature depends on the measurement of the neutron lifetime. More precisely, the freezeout temperature $T_{freezeout} \propto \tau_n^{1/3}$ [16]. Hence the neutron to proton ratio at the time of freezeout,

$(x_n)_{freezeout}$, depends on the neutron lifetime as

$$(x_n)_{freezeout} = e^{-QA\tau_n^{-1/3}} \sim 1/6 \quad (1.25)$$

where $1/A$ is the proportionality constant relating $T_{freezeout}$ to $\tau_n^{1/3}$, which depends on the expansion rate of the universe and the phase space for neutron-proton interchange reactions. From the time at which freezeout occurs until the time at which nucleosynthesis begins (several hundred seconds), x_n decreases slowly due to neutron beta decay (the only significant interchange reaction still in play). Hence, the neutron lifetime comes in a second time in the determination of x_n when nucleosynthesis begins. The neutron to proton ratio at the commencement of nucleosynthesis, $(x_n)_{nucl}$, is thus

$$(x_n)_{nucl} = e^{-QA\tau_n^{-1/3}} e^{-t_{nucl}\tau_n^{-1}} \sim 1/7 \quad (1.26)$$

where $t_{nucl} \sim 150$ s is the time at which nucleosynthesis begins (determined by temperature dependence of the abundance of photons with energy greater than the deuterium photodissociation threshold of ~ 1 MeV) [18].

Once nucleosynthesis begins, essentially all of the neutrons present are bound into ^4He nuclei, since that is the most tightly bound of all of the light nuclei. The vast majority of the excess protons remain unbound. Small amounts of deuterium and ^3He are formed (about 10^{-5} mass fraction each), as well as a miniscule amount of ^7Li (about 10^{-10} mass fraction). The ^4He mass fraction, Y can be estimated as

$$Y = \frac{4n_{He4}}{n_{Nucleons}} = \frac{4(n_n/2)}{n_n + n_p} = \frac{2x_n}{1 + x_n} \quad (1.27)$$

where x_n is just the neutron to proton number ratio at the time of nucleosynthesis.

Precise calculations of Y are performed using Monte Carlo simulations which take into account higher order effects not described herein. When Y is calculated from the BBN Monte Carlo, the contribution to the fractional error $\delta Y/Y$ due to the

uncertainty in the neutron lifetime is [19]:

$$\frac{\delta Y}{Y} \simeq 0.8 \frac{\delta \tau_n}{\tau_n}. \quad (1.28)$$

Thus the uncertainty in the mass fraction depends linearly on the uncertainty in the neutron lifetime. A fit of the calculated values of Y to the free parameters in the calculation yields [20]:

$$Y = 0.228 + \underbrace{0.010 \ln(\eta 10^{10})}_{0.011-0.014} + \underbrace{0.012(N_\nu - 3)}_{0 \text{ or } 0.012} + \underbrace{0.185 \frac{\tau_n - \bar{\tau}_n}{\bar{\tau}_n}}_{< \pm 0.0014 (95\% \text{C.L.})} \quad (1.29)$$

where N_ν is the number of neutrino flavors, η is the baryon to photon ratio of the universe, and $\bar{\tau}_n = 889$ s was the mean measured value of the neutron lifetime (in 1991). Hence we can see again the linear dependence of Y on τ_n . Currently, the most precise theoretical calculation of Y is [17, 19]

$$Y = 0.2462 \pm 0.0004 (\text{expt}) \pm < 0.0002 (\text{theory}) \quad (1.30)$$

where the experimental error is due to the uncertainty in the measurement of the neutron lifetime. Improvements in the precision of the neutron lifetime measurement could decrease the experimental error by slightly more than a factor of two before the next limiting uncertainty, that of the reaction rate for $d(d,n)^3\text{He}$, would dominate [19].

The usefulness of the BBN theory comes in comparisons between its predictions and the observed primordial elemental abundances. The typical approach is to fit the BBN theory to the measured D, ^3He , and ^7Li abundances by adjusting η [20]. The primordial abundances of ^3He is determined by comparing the elemental abundances of older meteorites (carbonaceous chondrites) with newer, gas-rich, meteorites while the primordial abundance of ^7Li is determined from observations of extremely metal-poor stars [20]. The abundance of deuterium is determined by studying high-redshift

clouds of hydrogen [17]. The observed abundances of these three elements constrain [20] the baryon to photon ratio to the range

$$2.8 \leq 10^{10}\eta \leq 4.0 \quad (1.31)$$

with deuterium currently providing the most stringent experimental limits [17]. Using the deuterium-inferred baryon density, the helium baryon number fraction can be calculated, taking the neutron lifetime and the number of neutrino species as additional inputs. The helium 4 baryon number fraction, Y , can be computed, with equal contributions to the error from the measurement uncertainty of the neutron lifetime (as described above) and the experimental uncertainty in the determination of η from deuterium observations [17]. The uncertainty due to deuterium observations is expected to be reduced by an order of magnitude in the next few years due to the Sloan Digital Sky Survey [17].

Comparison of predicted and observed values of the primordial ${}^4\text{He}$ abundance can be used to constrain N_ν . Currently, the errors on such a comparison are dominated by systematic errors in the measurement of Y at about the 1% level [17]. The experimental value of Y is mostly derived from determinations of the helium abundance in metal-poor extragalactic regions [21]. The determinations of Y reported by different groups differ from each other by more than the quoted errors, for example $Y = 0.234 \pm 0.002$ [22] and $Y = 0.244 \pm 0.002$ [23]. The main effort in improving the observational value of Y will thus center for the near term on better understanding of the host of systematic errors which complicate the determination of Y from observed hydrogen and helium line intensities [21].

If Y can ultimately be determined to the 10^{-3} level, which is unlikely for the foreseeable future, then comparison between the calculated and measured values of Y will provide a powerful test of the BBN theory and a strong constraint on the

number of light neutrino species. If the fractional error on the observed value of Y is decreased below the fractional error on the neutron lifetime, and the expected gains in the precision of η determination from deuterium observations is realized, then the error on the comparison between observed and calculated values of Y would scale linearly with the fractional error on τ_n .

1.4 Previous Measurements of τ_n

To measure the neutron lifetime, one can either observe the disappearance of neutrons or the appearance of the decay products, primarily the proton or electron. Both strategies have been used in the past and as a result all previous measurements of the neutron lifetime fall into two broad categories: storage experiments and beam experiments.

1.4.1 Beam Experiments

One approach to measuring the neutron lifetime is to count the number of neutron decays which occur in part of a neutron beam. Neutron beams are widely available at nuclear research reactors, and prior to the development of sources of much lower energy, storable, neutrons in the 1970s, this technique was the only one used for measuring the neutron lifetime. In beam measurements, the neutron lifetime is estimated from the differential decay law

$$\frac{dN_d}{dt} = -\frac{dN_n}{dt} = \frac{N_n}{\tau_d} \quad (1.32)$$

where N_n is the number of neutrons, and N_p is the number of decay events. Thus the technique requires an absolute measurement of two parameters: the number of neutrons within a volume of the beam, and the rate of neutron decays within that




Figure 1.3: A schematic of the beam lifetime apparatuses from [24] and [27]. Neutron decays within a penning trap result in trapped protons, which are intermittently counted. The neutron flux is measured by a thin detector which weights the flux by the inverse of the neutron velocity.

fiducial volume. In practice, a beam of neutrons contains neutrons of many velocities, so that the time dt in equation (1.32) must be calculated as l_{eff}/\bar{v} , where \bar{v} is the average neutron velocity, and l_{eff} is the effective length of the fiducial volume. Then, the neutron lifetime is determined from

$$\tau_n = \frac{l_{eff}}{\bar{v}} \frac{N_n}{N_d} \quad (1.33)$$

A number of approaches have been used to precisely define the fiducial volume and perform the two counting measurements. The two most recent beam measurements of the neutron lifetime have used the combination of a Penning trap to trap decay protons, and a $1/v$ weighted neutron detector for flux measurement [24, 25, 26, 27]. A schematic of the apparatus used is shown in figure 1.3. In these experiments, the neutron flux is measured by detecting alpha particles emitted from a thin foil, coated with one of the strong neutron absorbers ^{10}B or ^6Li . The foil absorbs only a small fraction of the entire neutron beam which passes through it. Since the capture cross section of the absorber varies with neutron velocity, v , as $1/v$, the number of number

of detected alpha particles, N_α can be related to N_n/\bar{v} by

$$\frac{N_n}{\bar{v}} = \frac{N_{alpha}}{\epsilon_0 v_0} \quad (1.34)$$

where ϵ_0 is the absolute efficiency for detecting a thermal neutron with velocity v_0 . The detection efficiency ϵ_0 can be determined by precise metrology of the capture foil and alpha detectors. Substituting equation (1.34) into equation (1.33), we can calculate the neutron lifetime as:

$$\tau_n = \frac{l_{eff}}{v_0} \frac{N_\alpha}{\epsilon_0 N_d}. \quad (1.35)$$

The other great difficulty in the experiment, the definition of the fiducial volume, is addressed by use of a segmented Penning trap to capture decay protons. By changing which electrodes on the trap are biased, and arranging for the electrodes to have identical size and separation, the effective length of the trap can be varied by a known amount, assuming that end effects are the same for each electrode. Then by measuring the decay rate at different lengths, the decay rate per unit length can be extracted. At each length, the number of decays is measured by counting the number of trapped protons produced by neutron decays approximately once per second.

The most recent lifetime results using the beam method are $\tau_n = 889.2 \pm 4.8$ s [25]² and 885.3 ± 4.0 s [27]. This method of measuring the neutron lifetime is currently limited by the systematic error associated with the absolute neutron flux measurement. To this end, precise calibration of the $1/v$ detector used in [27] by a cryogenic neutron calorimeter has been undertaken [28]. With a recalibration of that detector, the error on the result in [27] may be reduced to 2 s [29]. It is unlikely, however, that a beam measurement with errors much better than ± 1 s is feasible.

²The result in [25] is a reanalysis of the result in [24].

1.4.2 Storage Experiments

Since the 1970s, a new class of experiments has been performed, relying on the storage of very low energy neutrons in bottles. Neutrons have been confined by materials, the earth's gravitational field, and magnets. In all cases, a consistent number of neutrons is repeatedly placed in the bottle and after some length of time, the number remaining is measured. By measuring the variation in the number of neutrons remaining as a function of storage time, the neutron lifetime can be determined.

Ultracold neutrons (UCN) are defined as neutrons with wavelengths long enough that they can totally externally reflect from material walls. When neutrons with wavelengths long compared to interatomic spacings interact with a material wall, they see an effective potential which is a volume average of the delta-function nuclear potentials from the nuclei in the material. This potential barrier has the form

$$V = \sum_i \frac{2\pi\hbar^2}{m} N_i a_i \quad (1.36)$$

where m is the neutron mass, and N_i and a_i are the number density and bound coherent scattering length of the i th isotope in the material. Effective potentials as high as 335 neV (^{58}Ni) exist for UCN, while there are many materials with potentials on the order of 100 neV [30].

While UCN may be confined in material bottles, there is a finite probability per bounce that the UCN will be lost from the bottle. This loss can be due to neutron absorption or by inelastic scattering of the UCN to an energy higher than the confining potential. In the simplest model ³, these wall losses lead to an additional loss rate,

³Since the wall loss rate is a function of UCN energy (at least because the rate of wall collisions depends on the energy), wall losses add a non-exponential component to the observed decay curve.

$1/\tau_{loss}$, which will contribute to the measured loss rate, $1/\tau_{meas}$ as

$$\frac{1}{\tau_{meas}} = \frac{1}{\tau_n} + \frac{1}{\tau_{loss}} \quad (1.37)$$

In order to determine the true neutron beta decay lifetime, τ_n it is thus necessary either to reduce τ_n/τ_{wall} to below the desired experimental accuracy or to correct for τ_{wall} . Although extensive work has gone into minimizing the wall loss rate, and wall lifetimes on the order of 10 hours have been reported [31], in order to reduce errors below the 1% level, it has been necessary in all bottle experiments to employ some technique for compensating for the effect of wall losses. In different experimental setups, different compensation techniques have been used, but all require extrapolating to the case of no wall interactions. Two such extrapolations are: variation of the surface area to volume ratio of the bottle and extrapolation to the case of infinite volume [32], and use of a gravitationally sealed UCN spectrometer to measure the storage time as a function of UCN energy and extrapolating to the case of zero energy (zero wall collision rate) [31].

As an additional approach the two most recently reported bottle measurements of the neutron lifetime use neutron counters surrounding the UCN bottle to measure the rate of inelastically scattered neutrons [33, 34]. In the most recent bottle result, detection of inelastically upscattered UCN is combined with a double bottle arrangement in which τ_{meas} can be determined for two different neutron collision rates [34]. The quoted error for this most recent result is ± 0.9 s, dominated by statistical uncertainties, with the largest component of the systematic error due to the uncertainty in the relative detection efficiency for upscattered neutrons from the two bottles (this ratio is computed using a Monte Carlo simulation) [34].

In order to get away from the problems of wall losses, several past experiments have attempted to confine neutrons using magnetic fields. Confining UCN using a

combination of magnetic fields and gravity was originally proposed by Vladimirskii in 1961 [35] and later demonstrated [36]. The weak magnetic fields used for confinement in this apparatus, resulted in only about one UCN stored per loading cycle [36, 37]. The NESTOR experiment [38], on the other hand, used a superconducting magnetic hexapole storage ring to store circulating neutrons with large tangential velocities but small velocity in the two dimensions perpendicular to the storage ring. In NESTOR, wall losses are eliminated but a new limiting systematic was introduced. Betatron oscillations coupled momentum from the tangential to perpendicular directions, resulting in an additional loss rate $1/\tau_{loss}$ akin to wall losses. This systematic limited the achievable accuracy of NESTOR to ± 10 s [38]. An attempt was also made to magnetically trap UCN in a spherical hexapole trap [39]. In that work, the UCN were produced using the superthermal process (see below) and were stored for a time in superfluid helium before being dumped out to a neutron detector. Due to the high operating temperature of the experiment (1.2 K), the UCN upscattering lifetime in the superfluid helium was too short to permit observation of trapped UCN.

1.4.3 Current Status of τ_n

The current accepted value of the neutron lifetime is 885.7 ± 0.8 s and is computed as the weighted average of the six measurements with errors less than 10 s [5]. This average value does not include the most recent measurement using the beam technique [27]. Both the beam and bottle techniques which have been previously used are at or near the experimental accuracy allowed by systematics. In order to decrease by more than order of unity the uncertainty with which the neutron lifetime is known, a new technique is required.

1.5 τ_n from Magnetically Trapped UCN

A new technique for measuring the neutron lifetime, using magnetically trapped UCN was originally proposed in 1994 [40]. The experimental centerpiece of this technique is a cell filled with superfluid ^4He cooled to a temperature of less than 250 mK and surrounded by a three-dimensional magnetic trap. UCN with sufficiently low energy and in the low-field seeking spin state are confined by the interaction of their magnetic moment with the magnetic trap's field gradient. Since the magnetic trap is a conservative potential, any neutron with sufficient energy to enter the trap will continue through it and exit, unless it dissipates some energy while inside the trap. The dissipation mechanism used in our case also serves as the production mechanism for the UCN. This mechanism, known as the “superthermal process” is the downscattering of neutrons with wavelength near 8.9 \AA from superfluid helium by emission of a single phonon. As long as the helium bath is cooled sufficiently, the inverse upscattering process can be suppressed, and the UCN produced within the trap have no further interaction with the helium. The magnetic trap is loaded by passing a beam of 8.9 \AA neutrons through it for several neutron lifetimes, during which the number of the neutrons in the trap slowly builds up. After the trap is loaded, the neutron beam is turned off and neutron decays may be observed. The helium bath acts as a scintillator, converting the energy from the high energy electron produced in each decay event into light at extreme ultraviolet (EUV) wavelengths. This light is converted to visible light using an organic fluor. The visible light is then piped out of the apparatus to room temperature detectors. In this manner, a pulse of light is generated when each neutron decays and recording the time of each such pulse allows continuous measurement of the neutron decay rate as a function of time. Fitting this decay curve allows determination of the neutron lifetime.

Initial demonstration of this technique, including verification of the UCN production, magnetic trapping, and detection of neutron decays using helium scintillation, has been previously reported [41, 42, 43].

1.5.1 Magnetic Trapping

An alternative to using material walls to confine ultracold neutrons is to use a magnetic trap. A static magnetic trap is simply a configuration of coils which produces a three-dimensional magnetic field minimum. Particles with a magnetic dipole moment, $\vec{\mu}$, interact with a magnetic field, \vec{B} according to the dipole interaction

$$H = -\vec{\mu} \cdot \vec{B} \quad . \quad (1.38)$$

The neutron has a magnetic dipole moment which is anti-aligned with its spin, so that

$$\vec{\mu} = -\mu_n \vec{\sigma}_n \quad (1.39)$$

where $\mu_n = 1.91 \mu_N = 0.7 \text{ mK/T}$ ⁴ and the neutron's spin is $\vec{s}_n = (\hbar/2)\vec{\sigma}_n$. When the neutron's spin is aligned with the magnetic field and $\vec{\sigma}_n \cdot \vec{B} = |\vec{B}|$, the magnetic interaction potential reduces to $H = \mu_n |\vec{B}|$. In this “low-field seeking” spin state, the potential energy of the neutron increases with the magnitude of the local magnetic field and the neutron can be confined in a magnetic field minimum. If the neutron's spin is anti-aligned with the magnetic field then it will seek to minimize its interaction energy by moving towards regions of higher magnetic field.

Although the magnetic field is changing direction within the trap, the neutron's spin will remain aligned with the local magnetic field if an adiabaticity condition is

⁴The use of temperature units to denote energies is used throughout. In all cases $E = k_b T$ where k_b is Boltzmann's constant. For neutrons, 1 mK is equivalent to 86 neV.

maintained. The neutron spin precesses around the local magnetic field at the Larmor frequency $\omega_L = 2\mu_n B/\hbar$. As long as the change in the direction of the magnetic field seen by the neutron is small compared to this frequency, that is

$$\frac{d\vec{B}/dt}{|\vec{B}|} \ll \frac{2\mu_n B}{\hbar} \quad (1.40)$$

the neutron will remain in a well-defined spin state with respect to the local field. In practice this condition is always maintained so long as the magnetic field in the trap never approaches zero [42].

Since Maxwell's equations forbid the existence of a magnetic field maximum in free space in the absence of time-varying fields, and given the experimental difficulty of producing large magnetic fields at RF frequencies, static magnetic field minima are typically used to trap low-field seeking particles. While any configuration of magnetic fields that produces a field minimum in free space may be used as a magnetic trap, certain geometries have practical advantages. One important characteristic of magnetic traps is their "trap depth" which is the difference in magnetic field between the field minimum and the field at the edge of the trap (as defined by the physical walls). The maximum energy neutrons which may be trapped in a trap of depth B_{trap} are those with energy $\mu_n B_{trap}$. As will be seen below, the number of trapped neutrons is a strong function of the trap depth.

Due to the much smaller magnetic moment of the neutron compared with the typical paramagnetic atom ($\mu_n \sim 10^{-3}\mu_B$, where atomic moments are of the order μ_B , the Bohr magneton), neutrons must be cooled to much lower temperatures in order to load traps. For a trap depth of order 1 T, neutron temperatures of order 1 mK are required.

The magnetic trap used in this work is constructed in the Ioffe configuration (shown conceptually in figure 1.4) [44]. In this geometry, a magnetic quadrupole is




Figure 1.4: Conceptual current arrangement in an Ioffe configuration magnetic trap is shown.

formed by four infinite current bars distributed evenly around a central axis. This quadrupole produces a magnetic field which increases linearly with distance from this central axis, producing a radial magnetic trap. The ends of the trap are pinched off by two current loops. The two current loops have current flowing in the same direction (Helmholtz configuration) so that there are no zero-field regions within the magnetic trap. The resultant three dimensional magnetic trap has a sausage shape.

The Ioffe configuration has several advantages for this work. The large trap volume, especially due to the large achievable trap bore, couples well with the production mechanism, resulting in the possibility of a large number of trapped neutrons. Furthermore, the lack of a zero-field region in the trap makes possible the reduction of the loss rate due to Majorana spin-flips to less than 10^{-5} of the beta decay rate [40, 42].

1.5.2 UCN Production

Since a magnetic trap is a conservative potential, any particle with sufficient energy to enter the trap will have sufficient energy to exit it. Loading the trap thus requires a mechanism to dissipate energy while inside the trapping region. Dissipation mechanisms used in conjunction with magnetic trapping of atoms and molecules

include inelastic scattering from cold surfaces [45, 46], gases [47], or laser cooling [48]. None of these is practical for neutrons. Laser trapping is impossible since neutrons cannot be optically excited. Small scattering cross sections and low densities preclude cooling with a gas. Cooling neutrons with a liquid or a solid is generally infeasible due to the large absorption cross section of almost all materials. Liquid and solid forms of hydrogen, deuterium, and compounds containing them are typically used as moderators for thermal and cold neutrons. These moderators rely on limiting the number of scattering events to optimize the probability of thermalization versus absorption. In the case of a liquid or solid moderator filling a magnetic trap the moderator will continue to interact with the neutrons after they have been cooled, and so even relatively small capture or inelastic scattering cross sections would lead to continual losses.

An exceptional material, with regards to neutron interactions, is superfluid ^4He . Having the most tightly bound nucleus of any of the light isotopes, ^4He has an absorption cross section which is rigorously zero for low energy neutrons. Furthermore, at temperatures below the lambda point, 2.17 K, ^4He becomes a superfluid. In this state, the wavefunctions of individual helium atoms overlap and a macroscopic wavefunction must be used to describe the entire liquid. Thus neutrons no longer scatter from individual atoms in the superfluid, but rather from the superfluid as a whole. Quasiparticle excitations of the superfluid, among them phonons and rotons, can carry both energy and momentum. Inelastic scattering of neutrons results in the creation (or destruction) of such excitations. The momentum versus energy dispersion curves of the neutron and of the excitations in superfluid helium are shown in figure 1.5.

The unusual dispersion curve of superfluid ^4He is a key feature which allows the “superthermal process” (so named because it achieves greater phase space densities




Figure 1.5: The dispersion curves of the free neutron and the excitations in superfluid helium are shown. The curves cross only at two points: $q = 0$ and $q=q^*$ which corresponds to a wavelength of 8.9 \AA or an energy of 12 K . The measurements of the superfluid helium dispersion curve shown are from [49].

than thermal cooling would allow). The superthermal process, first proposed by Golub and Pendlebury [50, 51], is an attractive mechanism for producing UCN which can be adapted to be used inside of a magnetic trap. Unlike thermal moderation processes, in which neutrons come into equilibrium with a cold moderator through multiple inelastic scattering events, the superthermal process is a nonequilibrium process whereby a single scattering event converts warmer neutrons to colder ones, after which the neutrons no longer interact with the scattering medium, at least on any relevant timescale. The scattering interaction which produces the superthermal

process is the scattering of neutrons in superfluid helium with the production of a single phonon.

The dispersion curves of the free neutron and the phonons in superfluid ^4He cross only at zero energy and 12 K (see figure 1.5). Since both energy and momentum must be conserved, the only scattering processes which involve a neutron and the creation or destruction of a single phonon are those in which a 12 K (8.9 \AA) neutron scatters to rest and a 12 K phonon is created or in which a 12 K phonon is destroyed, promoting a neutron from near rest to 12 K. Of course, multi-phonon processes allow scattering of neutrons from all energies to all energies. Under the single phonon downscattering process, neutrons with energy near 12 K can scatter in the helium to produce UCN with arbitrarily low energy. Conversely, UCN in the helium bath, with energies of order 1 mK, can only be upscattered by absorption of a 12 K phonon.

The superthermal production of UCN has been demonstrated experimentally by several groups [41, 52, 53]. The theory of superthermal production rates was verified directly in our demonstration of magnetic trapping [41, 42, 43]. Further details on the theory of the superthermal process and its application to the calculation of UCN production rates may be found in section C.2.9.

The interaction of neutrons with a single phonon can be described as a two level system with an energy difference of 12 K between the two allowed neutron energy levels. Given a long enough time, a sample of neutrons within the helium must come into thermal equilibrium with the surrounding helium bath. The upscattering rate, can be made arbitrarily long by cooling the helium bath, decreasing the population of 12 K phonons by the Boltzmann factor, $e^{-12\text{K}/T}$ where T is the temperature of the helium bath.

At low enough temperatures, where single phonon upscattering becomes negligi-

ble, multi-phonon scattering can no longer be ignored. The rate of the second order, two-phonon process, is theoretically predicted to vary with temperature as T^7 [54]. The total upscattering rate of UCN in superfluid helium has been measured for helium temperatures down to 750 mK, where the measured upscattering rate is 10^{-3} s^{-1} [52]. Extrapolating that measurement to lower temperatures, it is expected that at our operating temperature of 250 mK the total UCN upscattering rate is approximately 2500 times less than the beta decay rate. At a temperature of 150 mK, the loss rate is approximately 10^{-5} of the beta decay rate. Hence, by sufficiently cooling the helium bath, the modification to the observed trap lifetime due to UCN upscattering can be reduced to a part in 10^5 .

Another possible UCN loss mechanism in liquid helium is neutron absorption. Although ^4He has zero absorption cross section, ^3He has a large thermal neutron absorption cross section of 5330 barns. Natural abundance helium has only a very small concentration of ^3He ($x_3 = ^3\text{He}/^4\text{He} = (1 - 2) \times 10^{-7}$ for commercial helium liquid and gas from natural gas sources and $x_3 = 1.3 \times 10^{-6}$ for atmospheric helium) [55]. Even at these concentrations, the lifetime of trapped UCN in such helium would be a small fraction of a second.

The loss rate due to absorption of UCN on ^3He , Γ_{abs} is

$$\Gamma_{abs} = n_3 \sigma_{UCN} v_{UCN} \quad (1.41)$$

where n_3 is the number density of ^3He inside the trapping region, σ_{UCN} is the neutron capture cross section for UCN and v_{UCN} is the UCN velocity. Since the capture cross section has a $1/v$ dependence, the product $\sigma_{UCN} v_{UCN}$ is equivalent to $\sigma_{th} v_{th}$, where the thermal neutron cross section for ^3He , $\sigma_{th} = 5330$ barns, and thermal neutron velocity, $v_{th} = 2200 \text{ m s}^{-1}$, are used instead. The number density of ^3He can be computed from its relative abundance and the known mass density of cold ($< 1 \text{ K}$),

superfluid helium $\rho_{he} = 0.14 \text{ g cm}^{-3}$ as

$$n_3 = \rho_{he} x_3 N_A / A \quad (1.42)$$

where $N_A = 6.022 \times 10^{23}$ is Avogadro's number and $A = 4.02602$ amu is the atomic mass of helium. The relative loss rates due to absorption of neutrons on ^3He and beta decay, $\Gamma_\beta = 1/\tau_n$, are then

$$\frac{\Gamma_{abs}}{\Gamma_\beta} = 2.2 \times 10^{10} x_3 \quad (1.43)$$

Therefore, in order to measure the neutron lifetime it is necessary to use isotopically purified ^4He , with a value of x_3 much less than the natural abundance. Isotopically purified ^4He has been produced with an indirectly measured⁵ $x_3 < 5 \times 10^{-16}$. The “ultrapure helium” (isotopically purified ^4He) used in this experiment is manufactured using a continuous flow apparatus for which the best indirect limit on the isotopic purity is $x_3 < 5 \times 10^{-13}$ [57]. We are currently undertaking a direct measurement of x_3 using accelerator mass spectroscopy, with which a limit on the order of 10^{-16} should be possible [58].

By filling a magnetic trap with superfluid ^4He , and scattering a beam of 8.9 \AA (12 K) neutrons from the helium, UCN with energies less than the trap depth can be produced. Those in the low field seeking spin state will then be trapped. The band of neutron momenta, k which contribute to UCN production is quite narrow:

$$k = k^* \pm k_{UCN} \quad (1.44)$$

⁵At the time the isotopic ratio measurements were performed, the maximum available sensitivity for measuring x_3 was 4×10^{-10} , by mass spectroscopy. In order to measure lower concentrations, the heat flush process used for isotopic purification was performed in reverse to increase the ^3He concentration in the sample. An upper limit on isotopic ratio x_3 was then obtained by measuring the reconcentrated sample and the ratio of the original, purified sample was inferred [56].




Figure 1.6: An overview of the scintillation process in liquid helium is shown.

where k^* is the crossing point of the dispersion curves and k_{UCN} is the maximum momentum of UCN under consideration (this is set by the trap depth in our case, where $\mu_n B_{trap} = (\hbar k)^2/2m$) [30]. For a neutron trap depth of 1 T, the required input wavelengths for superthermal production are 8.90 ± 0.07 Å. Thus only a narrow band of neutron wavelengths is needed for superthermal production, and a monochromatic neutron beam may be used to load the trap (see Chapter 3).

1.5.3 Detection of Neutron Decay Events

Neutron decays are detected by a combination of helium scintillation and wavelength shifting. An overview of the scintillation process is shown in figure 1.6. As the neutron's decay electron recoils through the helium bath, it creates along its track a large number of helium ions. These ions quickly recombine to form metastable dimers, He_2^* . The metastable dimer is formed in both the singlet and triplet states, with approximately half in each state [59]. The singlet He_2^* molecules radiatively decay in less than 10 ns, producing a pulse of light in the extreme ultraviolet (EUV) ranging in wavelength from 70 to 90 nm (the Hopfield Continuum) [60]. It is this prompt photon pulse that is used in detecting neutron decays. The prompt photon yield from helium

scintillation has been measured to be approximately 23 photons/keV [61]. Those He_2^* molecules formed in triplet states also decay radiatively, but the transition is highly forbidden and has a lifetime of 13 ± 2 s [59]. The light from the decay of triplet states forms an extremely weak phosphorescent tail which is not used for detection.

In order to use room temperature detectors, the EUV light is converted to the visible and piped out of the trapping region and the cryostat. Direct detection of EUV light using low temperature detectors inside of the cell is possible in principle, however activation, space, and cryogenic concerns prevent it in practice. Similarly, EUV light cannot easily be extracted from the cell given the extremely limited choice of window materials (EUV light will pass only through aluminum windows of a few hundred Å thickness). EUV light is thus frequency downconverted to the visible using an organic fluor.

A variety of fluors were studied and the one with the highest fluorescence efficiency (FE, the ratio of photons emitted to photons absorbed) was chosen [62]. This fluor is 1,1,4,4,-tetraphenyl buta-1,3-diene (TPB). TPB absorbs photons over a broad wavelength region from the soft UV to X-rays and emits blue light with a spectrum peaked at 440 nm and a width of approximately 50 nm [63]. The FE of an evaporated TPB film was measured to be 3.7 ± 0.1 relative to Sodium Salicylate [62], which has been independently measured to have an absolute FE of 0.37 [64]. Combining these two measurements, the absolute FE of TPB is 1.4. Thus for an average neutron decay of 350 keV a total of about 5600 blue photons should be present in the prompt pulse.

Detection of this large pulse of blue lights is accomplished through a series of lightguides and photodetectors (see section 4.4.1). The TPB fluor is evaporated onto a white diffuse reflector. Blue light which diffuses to the end of the cell passes through an optically clear beamstop, a series of lightguides and windows, and an optical

splitter before being coupled to two photomultipliers. The two photomultipliers are operated in coincidence mode in order to reject single photon backgrounds. The principles and characteristics of our detection system are described in greater detail elsewhere [65].

1.5.4 Continuous Decay Rate Measurement

The experiment is operated in two stages. In the first stage, the cell is exposed to the 8.9 Å neutron beam. Some of those neutrons downscatter to produce UCN. Those with sufficiently low energy and in the low-field seeking spin state are magnetically trapped. The neutron beam remains on for a period of several neutron lifetimes, during which the UCN density within the trap builds up. The second stage begins when the neutron beam is shut off and the observation of decay events begins. Scintillation events are detected and recorded as a function of time. If no backgrounds were present, the observed decay rate, $\dot{N}(t)$ would have the form:

$$\dot{N}(t) = \frac{N_0}{\tau_n} e^{-t/\tau_n} \quad (1.45)$$

where N_0 is the number of trapped neutrons at time $t = 0$ when the neutron beam is turned off. Fitting the observed decay rate to this model then yields the neutron lifetime.

The observed scintillation rate contains backgrounds as well as neutron decay events. Our detection system is designed to detect the passage of high energy electrons through the superfluid helium filling the trapping region. While it is our objective to detect electrons produced in the beta decay of trapped neutrons, the electrons (and indeed charged particles of any kind) may produce EUV scintillation light by the mechanism described in section 1.5.3. The decay of activated materials and the Compton scattering of electrons by high energy gamma rays may produce energetic

electrons, and hence scintillation, within the superfluid helium. The passage of muons (produced by cosmic ray showers in the upper atmosphere) through the liquid helium also produces scintillation light. Regardless of its source, all EUV scintillation light from within the superfluid helium is downconverted to the visible and piped to the photomultiplier tubes.

The helium is not the only part of the detection system which is sensitive to the passage of high energy charged particles (such as muons or Compton scattered electrons). The acrylic components of the detection system may also detect these particles by scintillation. Visible light may also be produced by luminescence of certain materials (see section 2.4).

Some backgrounds are time-independent, due to detection of randomly occurring events which produce ionizing radiation such as gamma rays and cosmic ray muons (see section 6.2.1). Time-dependent backgrounds, due to the introduction of the neutron beam into the apparatus, include neutron activation and neutron induced luminescence (see chapter 2 and section 6.2.2).

A background subtraction technique is employed in order to compensate for the presence of backgrounds, both time-dependent and time-independent. Two types of runs are taken, “positive runs” and “negative runs.” In the positive runs, the magnetic trapping field remains at its maximum value through both the loading and observation stages of the run. In these runs UCN are produced, stored, and observed as described above. In negative runs, however, the magnetic field is lowered to zero during the loading phase, and raised to its maximum value during the observation phase. In these runs, UCN produced during the loading phase are not magnetically stored and thus interact with the walls and are quickly lost. It is assumed that the backgrounds are independent of whether the magnetic field is on or off during the

loading phase⁶. Under this assumption the positive runs are a measurement of the neutron decay signal plus backgrounds and the negative runs are a measurement of the backgrounds only. Then the neutron decay signal can be extracted by subtracting the countrate of negative runs from that of positive runs.

1.5.5 Context of this work and Overall Prospects

Using a smaller version of the current apparatus, magnetic trapping of ultracold neutrons was demonstrated in 1999 [41, 42, 43]. In that work, approximately 500 UCN were trapped per loading cycle and their decay was observed with a 31% detection efficiency. The neutron lifetime estimated from two months running time was 660^{+290}_{-170} s [42], which is consistent with the known neutron beta-decay lifetime. The error in estimating the neutron lifetime was primarily due to the fact that the signal to background ratio was approximately 1/20.

The work described herein is the result of an effort to improve upon our previous trapping result. Various improvements in the apparatus and experimental technique were made in order to increase the signal to background ratio, increase the signal, and improve the statistical sensitivity to the neutron lifetime. A large number of changes were made, which are described in detail in the following chapters. A careful assessment of the neutron activation and neutron-induced luminescence properties of all materials exposed to the neutron beam was performed (see Chapter 2). In our previous work [41] the entire cold neutron beam, only a small fraction of which is in the narrow wavelength region which contributes to single phonon UCN production, was introduced into the apparatus. In an effort to increase signal to background, a

⁶The validity of this assumption is called into question by the results of the measurements described below (see section 6.3).

novel 8.9 Å monochromator was developed (see Chapter 3 and Appendix C). Changes made to the cryostat and its contents (see Chapter 4) include a larger magnetic trap and cell and a new superfluid helium-filled heat link connecting the cell to the dilution refrigerator. The detection system was also substantially overhauled (see Chapter 4). Changes include moving to larger light collection optics and detectors, the substitution of an optically clear neutron beamstop, and the switch from embedding the TPB in a clear plastic matrix optically coupled to a plastic tube, to evaporating the TPB onto a GoreTex diffuse reflector. A new data acquisition system and the software for analyzing the data are discussed in Chapter 5. The results from running the trapping experiment for about one month, and a discussion of their implications, are presented in Chapter 6. The upgrade to the apparatus and the results of running are also presented in another thesis being prepared concurrently [65]. The conclusions which can be drawn from the results, with special emphasis on possible future directions for the experiment, are found in Chapter 7.

Ultimately, any technique for measuring the neutron lifetime will be limited by systematic as well as statistical errors. The limiting systematic errors of bottle and beam lifetime measurements, wall interactions and flux measurement, respectively, are not relevant when measuring the neutron lifetime using magnetically trapped neutrons. The use of a magnetic trap prevents interactions with material walls, and the continuous measurement of the neutron decay rate makes an independent measurement of the neutron flux unnecessary. The limiting systematic errors for the magnetic trapping technique are expected to be non-beta-decay loss mechanisms for trapped UCN. All known loss mechanisms, including capture on ^3He , single- and multi-phonon upscattering, neutron depolarization (Majorana flips), and marginal trapping, are expected to result in loss rates less than 10^{-5} times the beta-decay rate

[40, 42]. Another proposed systematic error, the modification of the free neutron lifetime due to nuclear interactions with the surrounding helium, is calculated to have even less of an effect [42]. In principle, the measurement of the neutron lifetime using magnetically trapped UCN should remain statistics limited until the fractional error reaches 10^{-5} . As described below, the imperfect subtraction of time-dependent backgrounds would introduce a systematic error into the measurement. A change in the experimental technique (see section 7.3) may be necessary to eliminate this error.

Chapter 2

Materials Selection

2.1 Overview

Time dependent backgrounds in the experiment are the result of the interaction of neutrons with materials in the apparatus. During the filling phase of the experiment, large numbers of neutrons must be introduced into the apparatus in order to produce UCN. These neutrons must interact with various materials in addition to the ultrapure ^4He . If such an interaction results in the storage of energy in a metastable state (such as a radioactive isotope or a color center), then the release of that energy will have a time dependence after the neutron beam is shut off. Such time dependent backgrounds may mimic the neutron decay signal. Even when background subtraction is employed, one must worry about materials-related time-dependent backgrounds which may be imperfectly subtracted. Such imperfect background subtraction may be due to magnetic focusing effects which cause more neutrons to be incident upon some surfaces, depending on whether the magnet is on or off, or to changes in the decay rate of energy-storing material defects in magnetic fields (see for example [66]).

For these reasons, great care has been taken in selecting the materials which inter-

act with the neutron beam in order to minimize the known sources of time-dependent background (activation and neutron induced luminescence). In this chapter, some of the considerations for selecting the materials used, as well as measurements of their activation and luminescence properties will be described.

To the greatest extent possible, all neutrons introduced into the apparatus are captured in neutron shielding, which is fabricated from materials which have large neutron absorption cross sections. Neutron shielding can be used to prevent neutrons from interacting with the structural components of the cell, cryostat, and magnet, whose functions require them to be made from materials which activate upon neutron capture. In selecting neutron shielding, one must choose materials which are sufficiently strong neutron absorbers that they capture all of the neutrons which impinge upon them and which do not themselves activate. The functions of certain parts of the apparatus, such as windows for introducing neutrons into the apparatus, and the downconverting fluor and diffuse reflector portions of the detection system, necessitates that they be unshielded from neutrons. These materials must then be selected with extreme care to minimize activation. Minimizing activation requires selecting materials which are fabricated with high purity from elements with small neutron absorption cross sections for any isotopes which may result in activation. Additionally, any materials which are viewed directly by the detection system must be made from materials which do not luminesce after exposure to neutrons. For materials which must transmit the neutron beam, a low scattering cross section is also desirable, which generally precludes the use of hydrogenous materials, due to hydrogen's 80 barn scattering cross section.

The isotopes which may be considered for neutron shielding are shown in table 2.1. For shielding, one may consider any stable isotope with a large neutron absorption

Table 2.1: Stable isotopes with large ($\gtrsim 1000$ barns) neutron absorption cross sections are shown. For each isotope, the isotopic abundance Y , the thermal neutron capture cross section (σ_{th}), and dominant capture reaction are shown. All data are from [67].

Isotope	Y	σ_{th} [barns]	Capture Reaction
^3He	1.4×10^{-6}	5330	$^3\text{He} + \text{n} \rightarrow ^3\text{H} + \text{p}$
^6Li	0.075	941	$^6\text{Li} + \text{n} \rightarrow ^3\text{H} + \alpha$
^{10}B	0.199	3840	$^{10}\text{B} + \text{n} \rightarrow ^7\text{Li} + \alpha$
^{113}Cd	0.122	20600	$^{113}\text{Cd} + \text{n} \rightarrow ^{114}\text{Cd} + \gamma$
^{155}Gd	0.148	61000	$^{155}\text{Gd} + \text{n} \rightarrow ^{156}\text{Gd} + \gamma$
^{157}Gd	0.157	255000	$^{157}\text{Gd} + \text{n} \rightarrow ^{158}\text{Gd} + \gamma$
^{164}Dy	0.282	2700	$^{164}\text{Dy} + \text{n} \rightarrow ^{165}\text{Dy} + \gamma$

cross section. Each of the isotopes listed has advantages and disadvantages depending upon the application. In this experimental work, the primary consideration in choosing a shielding material is that the shielding material not become activated after absorbing neutrons. In this regard, ^{10}B is the ideal shielding isotope for our work. For all of the other isotopes considered, either one of the products of the reaction is radioactive (^3H and ^{165}Dy) or else another isotope of the same element activates to an undesirable extent (Cd and Gd). Neutron capture on ^{10}B results in a stable isotope of lithium, while the only other stable isotope of boron, ^{11}B , has only a 5 mbarn cross section for absorption and produces the isotope ^{12}B which decays with a lifetime of 20 ms, which is negligible on the time scale of our work. Therefore, all of the neutron shielding inside of the cryostat is made of boron compounds, which we have studied extensively. Some studies have also been performed on lithium based shielding.

The function of the various parts of the apparatus which are directly exposed

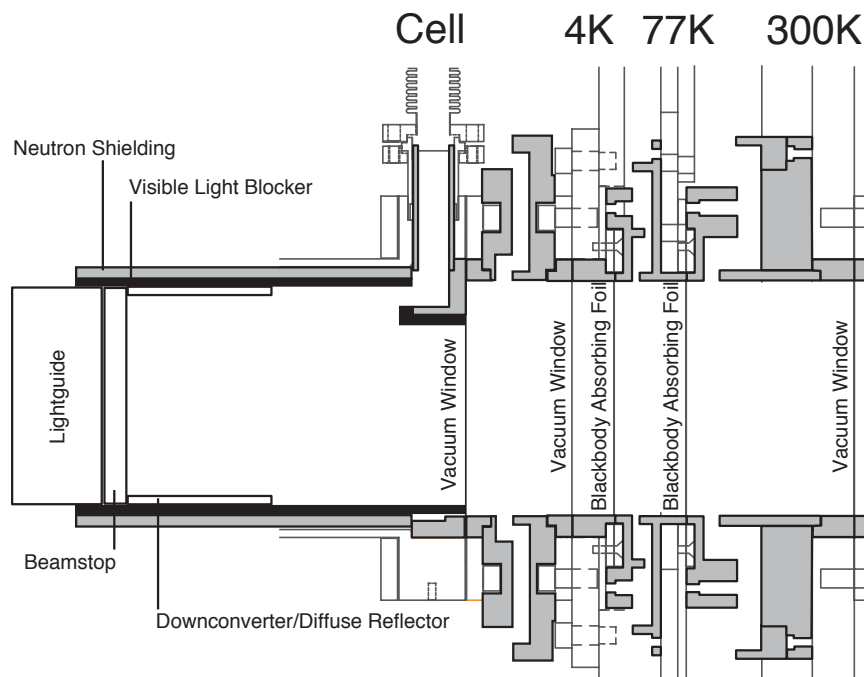


Figure 2.1: A schematic of the materials interacting with the neutron beam is shown (not to scale). The details of the materials chosen and their functional properties are described in Chapter 4.

to the neutron beam are shown in figure 2.1. This chapter contains descriptions of measurements of the the activation and luminescence properties of possible materials suitable for the functions shown.

2.2 Neutron Activation

The activation properties of any material can be thought of as deriving from two sources: first, the activation of the material’s main elemental components and second, the activation of any “impurities” contained in the material. Certain elements are of greater concern for activation than others. Elements which will lead to more activation are those with larger abundances of isotopes with larger neutron absorption

cross sections which, upon capture of a neutron, result in unstable radioactive isotopes of intermediate lifetimes (seconds to hours). Based upon these considerations, certain isotopes can be considered “safe” for neutron activation, including: beryllium, carbon, nitrogen, oxygen, and fluorine. For each of these materials, the neutron absorption cross section of the only isotope which can activate is relatively small, and the product isotope has either a very long (Be, C) or very short (N, O, F) lifetime. In addition, for nitrogen and oxygen, the only isotope of concern has less than one percent natural abundance. By selecting materials whose main elemental constituents are one of these elements or boron, the sources of problematic neutron activation can be reduced, to first order, to impurities. A quantification of the amount of activation which would be caused by different elemental impurities is found in Appendix A.

Information about impurities in materials comes from several sources. To the extent possible, materials were selected with the lowest possible impurity levels specified by their manufacturer. Based upon these specification, expected amounts of activation can be easily calculated using the information in Appendix A. As an example of this type of analysis, see the manufacturer’s compositional analysis for the purity of lot used to fabricate our Beryllium windows (used as blackbody absorbing foils, see section 4.2.4), as shown in table 2.2. In order to estimate the activation background due to the known impurities in the Beryllium foils, we should also multiply by the solid angle for decay products to reach the detection region ($\sim 5 \times 10^{-3}$) and the detection efficiency for each product which is a function of its type (gamma or beta) and its energy. In order to independently verify the elemental analyses specified by manufacturers, we performed several different types of activation tests.

Gamma spectroscopy is a key method we have used for activation testing. Materials can be activated with an intense neutron source (either a cold neutron beam or the

Table 2.2: The compositional analysis of Beryllium lot used to manufacture the windows used in the apparatus [68]. The windows in question are from Grade IF-1, the highest purity Beryllium foil commercially available. The manufacturer is Brush Wellman Engineered Materials, Fremont, CA. For the isotopes shown, the initial rate of radioactive disintegrations \dot{N} due to its products with lifetime τ are calculated for two windows with total mass 1.5 g, using Appendix A. Activation is calculated for an exposure time of 2100 s in an 8.9 Å beam of $10^7 \text{ n cm}^{-2} \text{ s}^{-1}$.

Element	Mass [%]	Isotope	τ [s]	\dot{N} [s^{-1}]
Be	99.90	9	6.8×10^{13}	1.2×10^{-6}
C	0.02	14	2.6×10^{11}	9.3×10^{-11}
Mg	0.0005	26	818	3.6×10^{-2}
Al	< 0.01	27	195	< 39
Si	0.0060	30	1.4×10^4	4.6×10^{-2}
Ca	0.0020	48	755	4.3×10^{-2}
Ti	0.0005	50	499	0.83
Cr	0.0015	54	303	0.11
Mn	0.0010	55	1.3×10^4	16
Fe	0.0145	54	1.2×10^8	2.7×10^{-4}
		58	5.6×10^6	1.6×10^{-4}
Co	0.0003	59	906	44
Ni	0.0065	64	1.3×10^4	0.12
Mo	0.0008	100	1.3×10^3	5.6×10^{-2}
Ag	0.0003	107	207	23
		109	36	837
Pb	< 0.0001	208	1.7×10^4	< 5.5×10^{-6}

reactor core). By observing with a low-background Germanium counter the gamma spectrum due to the sample, the amount of certain impurities can be inferred. This technique only works to identify impurities whose activation products decay with the emission of gamma rays. The results of activation testing by gamma spectroscopy following in-core neutron irradiation are given in Appendix E of [42] for many of the materials used in this work. Additional tests of that kind were performed in choosing materials for this work.

Another technique used for characterizing activation is specifically aimed at de-

tecting beta activity. Samples are irradiated in the cold neutron beam, then covered in a thin aluminum-coated mylar foil (enough to block visible light but not betas), and placed next to a scintillator which is coupled to a PMT. By observing the time variation of the scintillation rate, the amount and isotope of impurities can sometimes be identified. This technique was used to identify the second order activation mechanism producing ^{13}N in our Boron/Lithium loaded glass beamstop (described below and in [65]). Both gamma spectroscopy and beta activity scintillator measurements of candidate beamstop were used to select Boron Oxide as the least activating beamstop [65].

2.3 Second Order Activation Processes

A total of $\sim 10^{12}$ cold neutrons enter the neutron trapping apparatus per loading cycle. The large difference between the total number of cold neutrons and the number of trapped neutrons expected ($\sim 10^4$) makes the experiment sensitive to rather unlikely activation processes. Among these are secondary activation reactions, in which the products from an intended nuclear reaction (such as the capture of neutrons on shielding material) themselves cause activation.

The class of secondary reactions was first recognized as a problem for this experiment when the activation of ^{18}F was observed in a glass containing lithium and oxygen (see below). The very reaction which makes Lithium a good neutron absorber, neutron capture on ^6Li , produces 2.73 MeV tritons which can then react with the dominant natural oxygen isotope by the reaction $^{16}\text{O}(\text{T},\text{n})^{18}\text{F}$ [69]. The decay by positron emission of ^{18}F with a lifetime of 2.64 hours [67] results in the production of 511 keV gammas. This reaction renders shielding materials containing both lithium and oxygen unusable for the experiment.

Another second order reaction may make all boron shielding compounds problematic. In boron compounds, the dominant neutron capture reaction is $^{10}\text{B}(\text{n},\alpha)^7\text{Li}$, which produces a 1.47 MeV alpha approximately 94% of the time [70]. This alpha in turn may interact with the boron again by the reaction $^{10,11}\text{B}(\alpha,\text{n})^{13,14}\text{N}$. The total cross section for this process in natural isotopic abundance boron has been measured to be $2 \times 10^{-26} \text{ cm}^2$ for 1.5 MeV alphas [71]. It has been estimated that the yield of neutrons from ^{11}B is about 10 times higher than that from ^{10}B [72]. One can thus estimate that ^{13}N is produced in boron oxide at a rate of about 20 atoms/second in a neutron beam of 10^9 neutrons/s. Since ^{13}N decays by positron emission with a lifetime of 863 s [67], it is a particularly troubling background for neutron lifetime experiments. In studies of beamstop materials (see below), no conclusion could be drawn about the presence or absence of this second order reaction.

Further experiments are necessary to determine the overall efficiency of producing ^{13}N by exposing boron shielding material to a cold neutron beam. Additional second order reactions, as yet unidentified, may make a significant contribution to the observed activation backgrounds. The difficulty of identifying these reactions may be extreme. Furthermore, the additional constraints placed on the selection of materials exposed to the direct neutron beam may be impossible to satisfy fully. Such considerations add weight to the importance of modifying the experimental design to include separate regions for production and detection of UCNs (see section 7.3).

2.4 Neutron Induced Luminescence

The neutron induced luminescence of neutron shielding and scintillation detection materials was studied at 4 K. Samples were irradiated with a cold neutron beam. After the beam was turned off, luminescence in the visible wavelengths was observed as

a function of time using photomultipliers, while the samples were held at a constant temperature. Strong luminescence was observed for the shielding materials boron nitride and lithium fluoride. A measurable, but substantially smaller luminescence was observed from boron oxide. Other materials tested, including boron carbide, acrylic, GoreTex with an evaporated coating of the downconverting fluor tetraphenyl butadiene, and a boron and lithium loaded glass, displayed no detectable luminescence. The boron/lithium glass was determined to activate, by the secondary reaction $^{16}\text{O}(\text{T},\text{n})^{18}\text{F}$.

2.4.1 Introduction

The luminescence of materials after exposure to neutrons is a well known phenomenon. In particular, thermoluminescence of strong neutron absorbers, such as lithium fluoride (LiF), which are exposed to neutrons at one temperature and then heated to induce the luminescent release of energy stored in defects, is an important and widely used technique in neutron dosimetry [73].

Neutron induced luminescence represents a serious background for neutron experiments which look for a time-varying light signal. Two examples of such experiments are a measurement of the neutron lifetime [40] and a search for a neutron electric dipole moment [74]. In both experiments, a beam of cold neutrons is introduced into a bath of liquid helium at low temperature ($T < 0.5$ K). By way of the “superthermal process” [30], some of the cold neutrons downscatter to produce ultracold neutrons (UCN) while the rest of the cold beam either exits the apparatus or is absorbed by shielding materials. After UCN have been produced and accumulated, the cold beam is shut off and a time-dependent optical signal is observed. Scintillation light is generated when high energy particles resulting from either neutron decay or neutron

capture on ^3He recoil through a bath of liquid helium. The detection system used is sensitive to light over a broad spectral region from the extreme ultraviolet to the visible. A decaying luminescent signal due to the exposure of materials within the detection region to the intense cold neutron beam would introduce a systematic error in either measurement.

The work reported here was motivated by the observation of exactly such luminescence in an apparatus used to demonstrate magnetic trapping of neutrons [41, 43]. The luminescence background encountered was initially orders of magnitude more intense than the desired signal. Elsewhere we have reported the stabilization of the luminescence from boron nitride, which was used as a shielding material, in large magnetic fields ($\sim 1\text{ T}$) and at low temperatures ($< 0.5\text{ K}$) [66].

The luminescence properties of many materials, including some of those used in the magnetic trapping apparatus, have been studied after excitation from various sources, such as ionizing radiation, neutrons, UV light, and ion bombardment. There is a particularly rich literature for LiF and BN. Due to its use as a dosimeter, the thermoluminescence of neutron and γ irradiated LiF has been widely studied at and above room temperature (one representative example of many is [75]). The luminescent characteristics of BN similarly command a wide literature, including: neutron-induced luminescence at and above room temperature [76], UV and x-ray induced luminescence at temperatures down to 77 K [77, 78], and measurement of a continuous emission spectrum across the visible wavelengths when bombarded by 1.8 MeV α particles (near the 1.47 MeV produced by neutron capture on ^{10}B) [79]. The luminescence of other materials of interest has been studied using other stimuli than neutrons, such as: thermoluminescence of x-ray irradiated acrylic at temperatures down to 77 K [80], and γ -induced thermoluminescence of polytetrafluoroethylene (PTFE, the basic

constituent of GoreTex) from 77 K to 400 K [81]. In general these measurements have not been made below 77 K, nor has neutron-induced luminescence been studied for all of the materials in question.

In order to suppress the luminescent background in the neutron trapping experiment and in other proposed experiments, we have studied the luminescence properties of materials essential to the shielding and detection systems at liquid helium temperature. Because intense neutron beams must pass through the detection regions of these experiments, neutron-absorbing shielding must be used to prevent activation of the structural materials of the apparatus. Decays from activated structural materials produce betas and/or gammas, both of which the detection system is sensitive to. The detection system requires lightguides to transport visible light out of the detection region. Acrylic lightguides are typically used for this purpose.

Here we report studies of the luminescence of materials used in the optical detection and shielding systems of fundamental physics experiments with cold neutrons. Materials were studied at liquid helium temperature. The intensity and time dependence of the neutron induced luminescence is compared for the neutron shielding materials boron nitride, boron carbide, lithium fluoride, boron oxide, and boron/lithium glass, and the materials acrylic (Poly-methyl-methacrylate, PMMA), GoreTex with an evaporated coating of tetraphenyl butadiene (TPB), and graphite each of which is useful for other, non-shielding, purposes.

2.4.2 Apparatus and measurement technique

The apparatus used is shown in figure 2.2. Neutrons enter the apparatus through a 0.25 mm thick PFA teflon (Dupont Corporation) vacuum window which transmits neutrons with low scatter [82]. The teflon window is covered with 0.05 mm thick




Figure 2.2: The apparatus used for luminescence measurements in the single PMT configuration.

aluminum foil to keep room light out of the apparatus. Neutrons also pass through a $50\text{ }\mu\text{m}$ thick aluminum foil window at 77 K which is used to shield blackbody radiation. The sample material is clamped between two pieces of aluminum using aluminum screws. The sample holder is rotated about the vertical axis so that its front face makes a 45 degree angle with both the neutron beam and the light collection system. The neutron beam is collimated outside of the apparatus so that it is 7 mm in diameter at the sample. Exit foils and windows through which the beam passes for the case of non-beam-absorbing samples are identical to the entrance geometry.

All of the samples tested except for acrylic and GoreTex contract less than aluminum when cooled to low temperature so the differential thermal contraction in-

creases the clamping force to ensure good thermal contact as the sample is cooled. The goreTex sample, which is 1 mm thick, contracts by about $18\text{ }\mu\text{m}$ more than the aluminum holder, when cooled to 4 K. Since the goreTex is spongy the sample was compressed by more than this amount at room temperature so that after cooling the sample is still compressed. Likewise, acrylic contracts more than aluminum, but can be pre-compressed to largely, if not entirely, compensate for this contraction¹.

The sample holder is bolted to the bottom of the helium bath. The temperature of the sample holder is measured using a silicon diode thermometer². The sample holder routinely reaches an operating temperature of $< 5\text{ K}$. The samples are clamped on all sides and have an exposed area of 6.7 cm^2 . Based on the known thermal conductivities of the materials studied and the estimated blackbody heating from the 77 K shield, the maximum temperature at the center of any of the samples should be no more than 15 K.

Some of the light emitted by the sample is transmitted through a 3.7 cm diameter acrylic lightguide. One end of the lightguide is positioned 1 cm away from the face of the sample and is thermally anchored at 77 K. Thermal anchoring is achieved by clamping one end of the lightguide to ten 0.1 mm thick by 2 cm wide by 10 cm long aluminum foils which are connected between the 77 K bottom flange and the nitrogen shield at the other end. A diode thermometer measures the temperature of the aluminum clamp at the lightguide. A period of at least 3.5 h passes between first filling the nitrogen shield with liquid nitrogen and taking low temperature luminescence data

¹With an elastic modulus of 3 GPa, a compressive force of ~ 5000 pounds would be necessary to overcome the 8×10^{-3} differential contraction between acrylic and aluminum. The compressive force is provided by four 8-32 aluminum screws whose tensile strengths are of order 1000 pounds each.

²Lakeshore Cryotronics DT-470-CU-13. Lakeshore Cryotronics, Westerville, OH, (<http://www.lakeshore.com>).

to allow the temperature of the light guide to equilibrate. The minimum observed temperature of the clamp is 80 K and the clamp temperature remains below 130 K throughout all low temperature observations reported here.

The lightguide passes through a room temperature vacuum feedthrough and is optically coupled to either one or two Burle 8850 photomultiplier tubes (PMTs). When only one PMT is used it is coupled directly to the lightguide. When two PMTs are used they are each coupled to an acrylic “Y” piece which splits the light from the lightguide. Elsewhere we have measured that this splitter transmits 40% of the total light to each PMT. The PMTs are positively biased and operated in pulse counting mode. The Burle 8850 has superior gain dispersion due to a GaAs first dynode, and hence one-, two- and three- photoelectron peaks can be clearly distinguished from its pulseheight spectrum. Prior to recording time-dependent luminescence data, pulse height spectra were recorded for each PMT observing the constant background signal (see below) using a charge-coupled preamplifier, a shaping amplifier and a multichannel analyzer.

Time dependent signals were detected by counting the rate of PMT pulses directly from the anode (that is without any preamplifier). A 200 MHz voltage amplifier was used to buffer the anode signal from each PMT and fan it out to three discriminators³. The thresholds of the three discriminators for each PMT were set to obtain rates that were the same as the observed rate for thresholds of one, two, and three photoelectrons respectively in the pulse-height spectrum. While the best discrimination between different numbers of photoelectrons is achieved by using a threshold based on the area (charge) of the PMT pulse (as the MCA does), these voltage thresholds give reasonable photoelectron discrimination and permit counting of much higher rates

³Previous measurements have verified the linearity of this amplifier and the fact that it does not distort the integrated area (charge) of the PMT pulses.

than the charge-coupled preamp would allow. Output rates from each of the six discriminators is recorded using 100 MHz CAMAC counters which are read out every 1 second. The system has fixed readout dead time of 3%. The rate of coincidence between the two photomultipliers for each threshold level was also recorded using a 40 ns coincidence window.

Measurements of some of the samples were made using a simplified version of this data acquisition system which used only one PMT. In this case only one discriminator was used and its threshold was set so that the total observed rate was equal to the rate of all pulses of one photoelectron or greater measured using the MCA. The discriminator output a pulse of 1 μ s width which was fed into a Canberra Accuspec FMS multichannel scaler rather than the CAMAC-based system. We separately demonstrated using a function generator to feed pulses into the MCS that it is linear up to countrates⁴ of $7 \times 10^5 \text{ s}^{-1}$.

A thin-walled aluminum light shield is mounted on a linear vacuum feedthrough. The light shield is made from a commercial 474 ml aluminum beverage can by slicing off the top of the can and sanding the exterior of the can to remove all labelling.⁵ The can is epoxied to a thick aluminum disk which fits within the shaped recess at the base of the can (see figure 2.2). A 10 cm long section of 3.2 mm outer diameter, 0.25 mm wall thickness stainless steel tube provides a thermal disconnect between the can and the translational-rotational feedthrough. The can is heat sunk to the 77 K shield by three strips of 0.5 cm long beryllium-copper spring-clips. A thermometer is mounted on the aluminum disk at the bottom of the can. During operation, the

⁴The Canberra MCS has two signal inputs: “analog” and “digital.” Counting TTL pulses on the digital input allows higher maximum count rates. The $7 \times 10^5 \text{ s}^{-1}$ countrate quoted was measured using the digital input.

⁵Coors Light “Tall Boy,” graciously emptied by S. Dzhosyuk. Coors Brewing Company, Golden, CO, (<http://www.coors.com>).

temperature of the can remains below 150 K.

By means of the translational-rotational feedthrough, the light shield can be raised or lowered. In the raised (“can up”) position the light shield blocks all direct paths of light from the luminescent sample to the acrylic light guide. In the lowered (“can down”) position the can does not block any direct line of sight between the sample and the acrylic light guide. In the can down position, the aluminum can is out of the direct neutron beam. Scatter from the teflon beam windows, acrylic light guide, and the luminescent sample likely contribute to activation of the light shield (see section 2.4.4 below). The temperature of the luminescence sample does not change as the can is raised and lowered.

This work was performed at the NG-0 cold neutron beam station at the NIST Center for Neutron Research. A total of 10^9 neutrons/s enter the apparatus when the neutron beam is on. The neutron guide used is curved, reducing the gamma flux on the sample. Nevertheless, a constant gamma background of ~ 1 mR/h is present at the sample location and the acrylic lightguide. The neutron beam is collimated by two apertures, both made from ^6Li -loaded plastic. Beam collimation is chosen so that neutrons entering the apparatus hit the sample, except for those which scatter from the teflon entrance window or are absorbed by the thin 77 K shield. The apparatus is aligned so that the center of the neutron beam hits the center of the sample holder opening.

2.4.3 Samples Tested

The luminescent properties of various materials which are important for shielding neutrons, or for detecting the decay products of neutrons in the magnetic trapping experiment [41] were studied.

Five neutron shielding materials were studied: boron nitride (BN), boron carbide (B_4C), boron oxide (B_2O_3), lithium fluoride (LiF), and a boron/lithium glass. For all of the neutron shielding materials studied, the samples were at least 6 mm thick and black to the cold neutron beam ($< 10^{-4}$ transmission). The boron nitride studied was Grade AX-05 hot pressed hexagonal boron nitride⁶ which is a white material similar to graphite. Boron carbide is a black, hot-pressed ceramic. The sample tested was made from boron carbide powder isotopically enriched in ^{11}B (89% ^{11}B by number)⁷. The LiF sample is an optical quality window, 30 mm in diameter and 5 mm thick⁸. The boron/lithium glass was specially fabricated to minimize neutron activation while being black to the neutron beam. Its approximate composition (by weight) is 70% B_2O_3 , 27% Li_2O , and 3% ZrO_2 [65, 83].

In addition to the neutron absorbing materials, several other materials were also studied. The organic downconverter 1,1,4,4-tetraphenyl buta-1,3-diene (TPB) is evaporated onto the fluoropolymer GoreTex. A sample with approximately 200 mg/cm² of TPB evaporated onto 1 mm thick GoreTex was tested⁹. A sample of Ultraviolet Transmitting (UVT) acrylic was also tested¹⁰.

⁶Grade AX-05 is the highest purity boron nitride which we know to be commercially available. It is produced by Saint-Gobain Advanced Ceramics Corp., Amherst, NY, (<http://www.bn.saint-gobain.com/>).

⁷Boron carbide was purchased from Boron Isotopes, a division of Eagle Picher Technologies, LLC, Quapaw, OK, (<http://www.epi-tech.com>). The ^{11}B enriched sample was chosen due to its higher specified purity than natural abundance boron carbide.

⁸The lithium fluoride window was purchased from I.S.P. Optics, Irvington, NY, (<http://www.ispoptics.com/>).

⁹Scintillation grade ($> 99.5\%$) TPB is purchased from Fluka, a division of Sigma-Aldrich. GoreTex GR gasket material, which is manufactured from expanded PTFE, is obtained from W. L. Gore Associates.

¹⁰UVT acrylic is manufactured without the additives that are placed in standard

A scheme to block luminescence light was tested using a sample consisting of BN enclosed in graphite, a material which is transparent to neutrons. For this sample, a 25 mm diameter, 4.6 mm thick disk of BN was prepared. This disk was inserted into a matching well machined into a graphite disk of diameter 38 mm and 6.4 mm thickness, to form solid BN/graphite disk. The side of this composite disk with exposed BN was then covered with a 0.5 mm thick sheet of graphite, thus completely enclosing the BN. The side of the sample with the 0.5 mm graphite sheet covering the BN was the side exposed to neutrons and to the detector. All of the graphite used in fabricating this composite structure was high purity, baked graphite, grade AXF-5Q1¹¹.

2.4.4 Results and Discussion

Samples are exposed to neutrons for either 300 s or 1500 s (only when no luminescence was observed in a 300 s exposure). While the neutron beam is on, the PMT is turned off and the light shield is down (sample unblocked). Immediately after turning the neutron beam off, the PMTs are biased and observation begins. Two runs are taken with each sample. In one case (“positive”), the aluminum can remains down during observation so that light is collected from luminescence of the sample as well as all other background sources. Another, “negative”, run is taken where the aluminum can is raised immediately after the neutron beam is turned off. In this case, only light due to various background sources is collected. Taking the difference of these

commercial acrylic to absorb ultraviolet (and short-wavelength blue) light. UVT acrylic tube and rod are manufactured by Polycast Technology, Stamford, CT, (<http://www.digitalscribe.com/Polycast/>). We purchased acrylic from the distributor Boedeker Plastics, Shiner, TX, (<http://www.boedeker.com/>).

¹¹Graphite of grade AXF-5Q1 was purchased from Poco Graphite, Decatur, TX, (<http://www.poco.com/>). This grade of graphite has been baked for 3 weeks at 500 C to reduce impurities to < 5 ppm, according to the manufacturer’s specifications.

two runs yields the signal due only to visible light emitted by the sample. Samples are observed until the luminescence rate returns to the background.

As is described above, measurements were performed on two occasions, with different data acquisition systems. Initial measurements were made using two photomultipliers and slightly more complicated electronics. In the two PMT setup, with no sample present, the apparatus was exposed to the neutron beam for 300 s. The observed signal on one of the PMTs¹² after the beam was turned off is shown in figure 2.3 for both the can up and can down positions. In both the can up and down runs, the signal, which is on the order of several ks^{-1} initially, decays in a manner consistent with a single exponential with the aluminum-28 lifetime of 195 s. The beta-decay of aluminum produces a beta with energy up to 2.9 MeV and a 1.8 MeV gamma. Either betas or gammas may produce a signal in the acrylic lightguide. The small, time-independent difference between the can up and can down rates shows that the change in the activation signal due to the can position is small. The fact that the activation signal is slightly larger in the can up configuration is consistent with our model, since the solid angle for decay products from the activated aluminum can interacting with the acrylic lightguide is increased compared with the can down configuration. As these empty run measurements show no long-time signal, then any such signal observed must be due to whichever sample is under observation.

The difference between can down and can up signals for the BN and LiF samples are shown in figure 2.4. For both BN and LiF, large, long-time signals are observed in the can down runs. With the can up, in both cases the signal observed is consistent with aluminum activation and a constant background. Hence, in the difference signal,

¹²In the 2 PMT setup, the two PMTs always observed virtually identical signals. For ease of presentation, the signals from one of the two PMTs (denoted “PMT E”) are always shown unless otherwise noted.




Figure 2.3: The signal observed on a single PMT after the neutron beam has been on for 300 s with no sample. Data are shown for the can up and can down signals. The first 594 s of the can down data have been shifted down by 364.81 s^{-1} to correct for a discontinuous change in the gamma background rate, due to a sample change on a nearby apparatus. The difference of the two signals is also shown.

a long-time signal is still observed. This signal must be due to luminescence of the samples which is blocked by the aluminum can in the can up configuration. Attempts to fit the BN and LiF difference signals to various simple models (exponential + $1/t$ + constant, constant + 2 exponentials, etc.) all resulted in large values of χ^2 (> 4 per degree of freedom) and residuals with visible structure. Hence an exact model for the time dependence of the luminescence was not obtained.

One important question regarding the luminescence background is whether or not it produces any correlated multi-photon events. Since it is due to a solid state process [84] it is expected that the luminescence should consist only of uncorrelated single photoelectron events. Figure 2.5 shows the coincidence signal measured for the BN




Figure 2.4: The difference between can down and can up signals observed on a single PMT (in the 2 PMT setup) after the neutron beam has been on for 300 s is shown for BN and LiF samples.

can down run, with single photoelectron thresholds on each PMT. The accidental coincidence rate is also shown (this is computed as the product of each PMT's individual single photoelectron rate, times the coincidence window of 4×10^{-8} s plus the measured constant offset). The accidental and measured coincidence rates are in good agreement at very short times (when the count rate is high) and at long times (when the measured coincidence rate is low). From about 200 to 1000 s after the beam is turned off, the measured coincidence is higher than the accidental coincidence rate. An analysis of the difference between the measured and accidental coincidence rate shows that the difference signal fits well to a single exponential consistent with the aluminum lifetime. Electrons decelerating in the acrylic and producing scintillation light can produce a short time burst of many photons. Hence, it is not surprising to

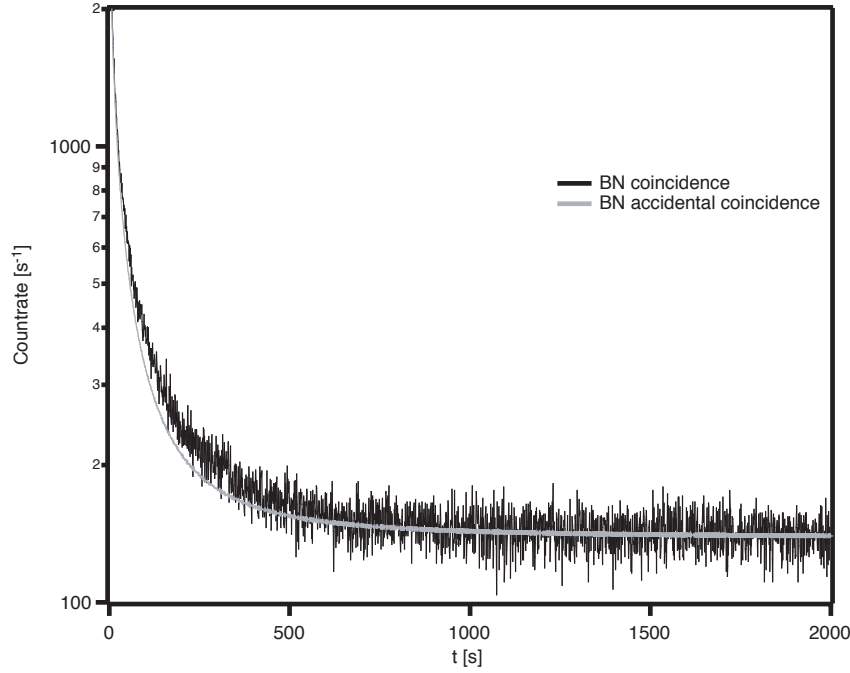


Figure 2.5: The can down BN luminescence signal is shown, where coincidence is required between the two PMTs, each with a its threshold set to detect single photoelectron events. The computed rate of accidental coincidences is shown for comparison. The accidental coincidence rate is the product of the two single PMT rates times the coincidence window 4×10^{-8} s plus the constant offset determined as the average from $t=2000$ to $t=6000$ s of the measured coincidence rate.

detect the aluminum activation signal in the coincidence data, which should select out only the multi-photon events plus accidental coincidences. Within the experimental sensitivity, there is no detectable luminescence signal due to coincidence. Starting about 1000 s after the beam is shut off, the coincidence data is flat to within the noise. By comparing the noise at 1000 s with the measured luminescence rate at 1000 s, one can put an upper limit on the fraction of the luminescence events due to multiple photons of $< 10^{-3}$.

One possible method for ameliorating backgrounds due to luminescence is to introduce a light-blocking material between the luminescent material and the photon

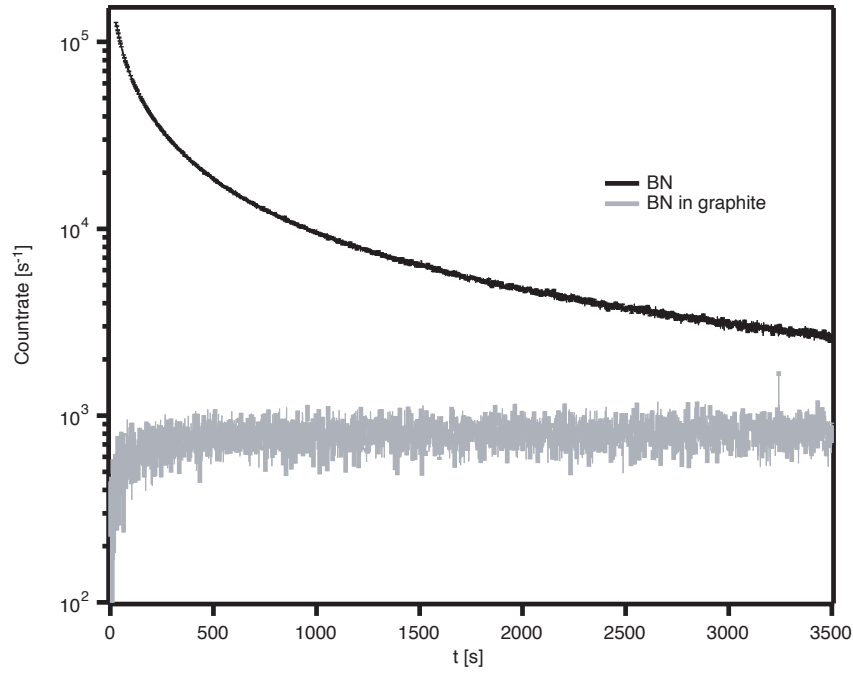


Figure 2.6: The difference between can down and can up signals observed on a single PMT (in the 2 PMT setup) after the neutron beam has been on for 300 s is shown for BN and BN embedded in graphite. The BN/graphite difference curve has been shifted up by 2000 s^{-1} to allow the use of a logarithmic y scale. The uncorrected BN/graphite difference curve has a negative constant offset because the gamma background increased between the can down and can up runs when a sample on an adjacent instrument was changed.

detector. So long as this light blocker does not luminesce or activate on its own, then the overall background should be reduced. A test of this scheme was performed by fabricating a sample consisting of boron nitride embedded in graphite. A comparison of the difference signals for boron nitride and BN embedded in graphite is shown in figure 2.6. In both cases, the samples have been irradiated for 300 s before counting, and the signal shown is that of a single PMT in the two PMT configuration. For the BN/graphite sample, the observed decay signal is consistent with aluminum activation and a constant background for both the can up and can down cases. The difference signal for the BN/graphite sample fits to a negative amplitude aluminum

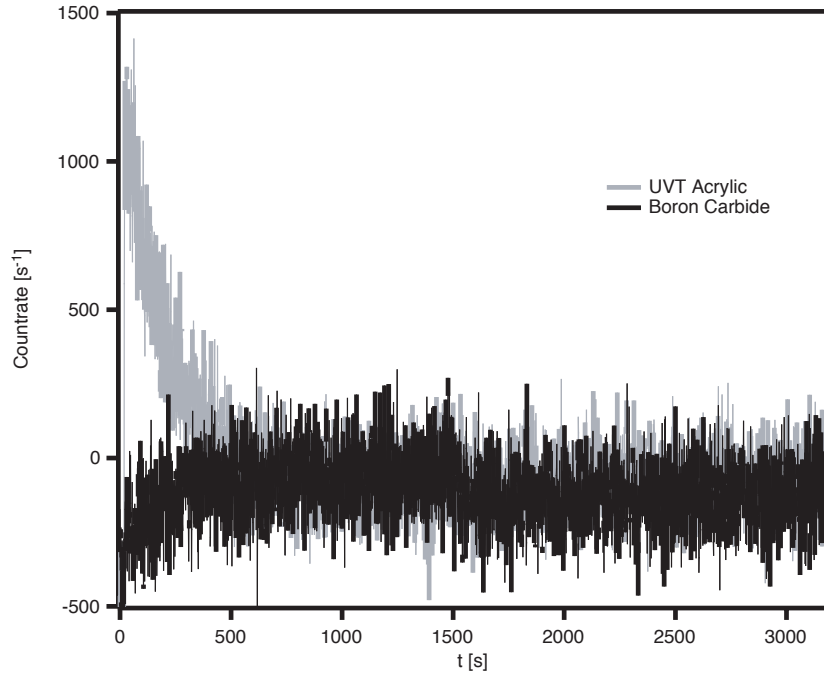


Figure 2.7: The difference between can down and can up signals observed on a single PMT (in the 2 PMT setup) after the neutron beam has been on for 1500 s is shown for boron carbide and UVT acrylic. The slight discontinuity in the boron carbide signal at $t=1500$ s is due to a change in the gamma background caused by an adjacent instrument which changed samples during the can down run.

signal, as in the case of no sample present. Hence we can conclude that the graphite itself does not luminesce and that 0.5 mm thick graphite is sufficient to block all of the luminescence light from BN. The suppression factor for the luminescence after embedding BN in graphite is at least 10^4 .

The difference between can down and can up signals for boron carbide and UVT acrylic are shown in figure 2.7. Both materials were irradiated for 1500 s before counting (after no luminescence was observed following a 300 s exposure). The difference signal for the UVT acrylic fits to a positive amplitude exponential decaying with a lifetime consistent with that of aluminum. The ratio of the difference amplitude to the can up or can down signal (approximately 0.1) is very similar to the ratio of




Figure 2.8: The difference between can down and can up signals for LiF is compared for two experimental setups. For the “2 PMT” setup, the signal shown is that of one of the two PMTs. The occasional drops in the signal in the one PMT setup are due to an intermittent problem in the electronics.

acrylic masses of the sample to the lightguide. Hence, the difference signal in UVT acrylic is consistent simply with the detection by the acrylic of gammas and betas from aluminum activation. The boron carbide difference signal is a negative amplitude aluminum decay exponential, just as in the empty sample and BN/graphite runs. Hence we can conclude that there is no evidence, within our experimental sensitivity, for luminescence of boron carbide or UVT acrylic at liquid helium temperature.

The luminescence of additional materials was studied with a simplified experimental setup in which only a single PMT was coupled to the apparatus’ lightguide. A comparison of the luminescence from LiF made using the two experimental setups is shown in figure 2.8. For both setups, the difference between the can down and can up signals is substantially similar. For both setups, the LiF was exposed to the neutron




Figure 2.9: The can down and can up signals are shown for TPB evaporated on GoreTex. The sample was exposed to neutrons for 1500 s before observation began. The difference signal (can down - can up) is also shown.

beam for 300 s. The overall similarity in the amplitude and shape of the difference signals for the two setups indicates that comparison between results taken in the two setups is valid, especially in terms of the ability to detect the presence or absence of luminescence.

Using the single PMT setup, the luminescence of several additional materials was measured. The can down, can up, and difference signals for TPB evaporated on GoreTex are shown in figure 2.9. Both the can down and can up signals fit well to a single aluminum exponential plus a constant offset. Since the amplitude of the aluminum signal, as well as the constant offset, varies between the two runs, the small residual difference signal displays a time dependence. Since the two raw signals fit so well to aluminum activation, this residual time dependence is attributed to the




Figure 2.10: The can down and can up signals are shown for the boron/lithium glass. The sample was exposed to neutrons for 300 s before observation began. The difference signal (can down - can up) is also shown.

imperfect subtraction of the activation background, due to the changing efficiency for detection of activated aluminum between the can down and can up runs.

Similarly, the can down, can up and difference signals for the boron/lithium glass are shown in figure 2.10. Neither the can up nor can down signals for boron/lithium glass fit to a single aluminum exponential. The time dependence can be modelled as a single exponential of a longer (300-500 s) lifetime. The difference signal can be fit to a single exponential with a very long lifetime (consistent within 2σ with a constant) and low amplitude ($< 100 \text{ s}^{-1}$). This long signal could be due to luminescence at a level $< 10^{-4}$ of LiF or BN. Most likely the small residual difference signal is due to an imperfect subtraction of the can down and can up signals (as for the TPB/graphite sample).

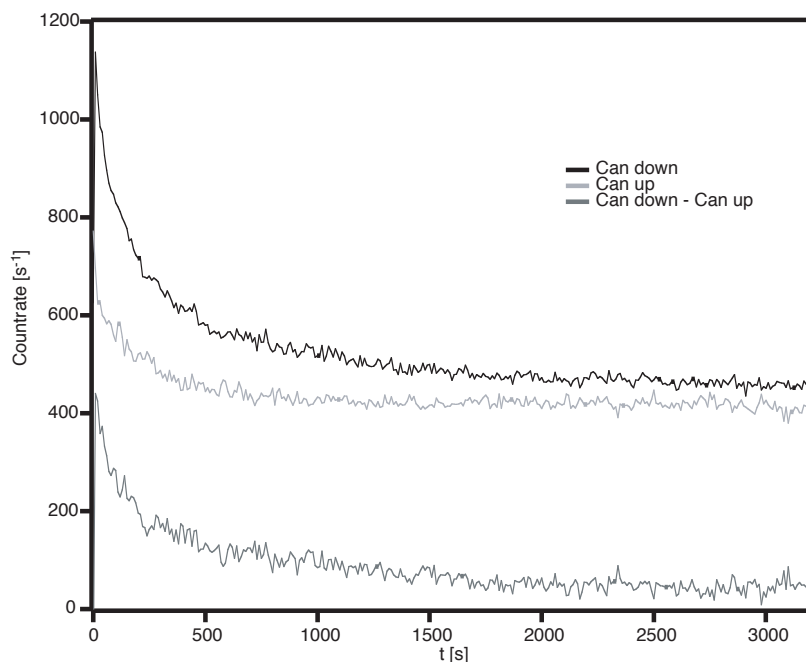


Figure 2.11: The can down and can up signals are shown for boron oxide. The sample was exposed to neutrons for 1500 s before observation began. The difference signal (can down - can up) is also shown.

Even in the absence of luminescence, the boron/lithium glass suffers from an important, inherent activation background. As described above (see section 2.3), any compound containing lithium and oxygen produces ^{18}F when bombarded by neutrons. The reaction pathway includes two reactions, $^6\text{Li}(\text{n},\text{T})^4\text{He}$, followed by $^{16}\text{O}(\text{T},\text{n})^{18}\text{F}$. The production of ^{18}F in neutron-bombarded boron/lithium glass was confirmed. Using a detector sensitive only to gammas and not to visible light, a decaying exponential signal in good agreement with the ^{18}F lifetime of 2.64 hours was measured following exposure of the boron/lithium glass to neutrons (see [65] for additional details).

The final material studied was boron oxide. The can down, can up, and difference signals for boron oxide are shown in figure 2.11. While the can up signal fits to the

model of only aluminum activation and a constant background, this model is not a good fit to the can down or the difference data. The difference signal fits about equally well to a double exponential or a single exponential plus $1/t^{power}$. Hence, we cannot tell from the shape of the curve alone whether the observed time dependent difference signal in boron oxide is due to luminescence or activation. If the signal is due entirely to luminescence, then the relative intensity of luminescence in boron oxide versus BN or LiF is about 10^{-3} .

As is the case with the boron/lithium glass, a second order activation process of concern exists for boron oxide. Boron oxide with total impurities < 1 ppm is obtainable (see section 4.4.1) and should be relatively free from direct activation by neutron capture. As with all boron compounds (see section 2.3), boron oxide may produce ^{13}N by the second order reaction $^{10}\text{B}(\alpha, n)^{13}\text{N}$. In this case, the alphas are produced by neutron capture on ^{10}B . It is estimated that 20 atoms/second of ^{13}N are produced when boron oxide is exposed to a neutron beam of 10^9 neutrons/s. The signature of such an activation would be the observation of ^{13}N decay, by positron emission, with a lifetime of 863 s [67].

As noted above, curve fitting to the boron oxide difference signal yields equally good fits to a double exponential or an exponential plus $1/t^{power}$. Two possible curve fits are

$$y = (36 \pm 2) + (171 \pm 7)e^{-t/(879^{+62}_{-54})} + (260 \pm 63)e^{-t/(88^{+22}_{-14})} \quad (2.1)$$

which has $\chi^2 = 1.93$ or

$$y = (-23 \pm 34) + (73 \pm 24)e^{-t/(738^{+123}_{-92})} + (1149 \pm 144)t^{-0.36 \pm 0.08} \quad (2.2)$$

which has $\chi^2 = 1.98$. In either case the longer exponential is consistent with the activation of ^{13}N . In order to definitively identify the activation of ^{13}N a separate measurement should be made where a boron oxide sample is irradiated, then wrapped

in a light absorber and placed on a scintillator. Such a setup would be insensitive to luminescence but quite sensitive to the pair production gammas produced when ^{13}N decays.

2.4.5 Conclusions

For experiments which detect visible light from regions which have been exposed to large numbers of neutrons, materials selection is extremely important. In addition to the well known concern of neutron activation, which in combination with any scintillation mechanism can result in visible light, neutron-induced luminescence may be a significant source of background. In this work, we have studied a variety of materials which are useful in experiments which detect light produced by neutron beta decay or neutron capture on ^3He [40, 74]. None of the weak neutron absorbers studied (acrylic, graphite, TPB evaporated on GoreTex) displayed any detectable neutron-induced luminescence.

Strong luminescence, of the order of 10^5 s^{-1} initially after capture of $\sim 10^{11}$ cold neutrons, was observed for the neutron absorbers BN and LiF. By comparing the rates of luminescence on two PMTs to the rate of their coincidence, a limit on fraction of luminescence events producing more than one correlated photon was set as $< 10^{-3}$. Although BN is an appealing shielding material due to its availability in high purity form and its easy machinability, its strong luminescence is a concern. An alternative shielding material, boron carbide, did not display any measurable luminescence. Since boron carbide is extremely difficult to machine, it may be preferable in many application to use BN as a neutron absorber but to shield the detectors from BN's luminescence light using a thin ($\sim 0.5 \text{ mm}$) layer of graphite as a light absorber.

Of the optically transparent neutron absorbers, lithium fluoride, boron oxide, and

boron/lithium Glass, strong luminescence was observed only for lithium fluoride. The boron oxide displayed a time varying signal consistent either with luminescence with $\sim 10^{-3}$ intensity relative to that of lithium fluoride or with activation. In either case, the signal is consistent with the activation of ^{13}N by the reaction $^{10}\text{B}(\alpha, n)^{13}\text{N}$. The boron/lithium glass did not display any measurable luminescence. All three of these materials are subject to some activation concerns. The fluorine in LiF becomes activated by the $^{19}\text{F}(n, \gamma)^{20}\text{F}$ reaction. Since ^{20}F decays with a lifetime of 16 s, which is much shorter than the neutron beta-decay lifetime of 886 s, it may be sufficient in many applications where LiF must be used simply to wait for any activated fluorine to decay. The production of ^{18}F in the boron/lithium glass and ^{13}N in boron oxide may be much more problematic due to their longer lifetimes. Due to the strong ^{18}F decay signal observed in the boron/lithium glass, and the intense luminescence observed in LiF, boron oxide was chosen as the beamstop material for the neutron trapping measurements made hereafter.

Chapter 3

An 8.9 Å Neutron Monochromator

3.1 Introduction

The production of UCN by single phonon downscattering of cold neutrons from superfluid helium (the “superthermal process” [30]) requires input neutrons only in a narrow wavelength band around 8.9 Å (see sections 1.5.2 and C.2.9). For the deepest magnetic trap conceivable for this work (see section 7.7), with a trap depth of 4.4 T, corresponding to a UCN potential of 273 neV (3.1 mK), the production band is 8.9 ± 0.1 Å. Hereafter, references of 8.9 Å neutrons implicitly refer to neutrons in the wavelength band 8.9 ± 0.1 Å.

Only a small fraction ($< 1\%$) of the spectrum of the cold neutron beam at NG-6, where the UCN trapping experiment is performed, lies within the relevant wavelength band for production. Thus, for the neutron lifetime measurement, all other cold neutrons contribute mainly to backgrounds¹. The signal to background ratio of the

¹For the dominant, single-phonon downscattering process, only neutrons with wavelengths near 8.9 Å can produce UCN. Higher order, multi-phonon downscattering processes can produce UCN from neutrons of other wavelengths ($\lambda \gtrsim 6$ Å). It has been estimated that the total non-8.9 Å production is less than or equal to

experiment can be substantially improved by filtering the cold neutron spectrum so that the only neutrons entering the trapping apparatus are those which can produce UCN by single phonon downscattering.

Optimally, the number of neutrons in the production band entering the UCN trapping apparatus would be maximized, while all other neutrons would somehow be eliminated before entering the apparatus. Although such an ideal situation cannot be achieved, several existing technologies allow some manner of neutron wavelength selection. Mechanical velocity selectors use a spinning helical neutron absorber to constrain transmitted neutron wavelengths [85, 86, 87]. Transmission of neutrons through polycrystalline materials with Bragg edges can introduce a wavelength long-pass filter (see for example [88, 89] and section 3.8.1 below). Supermirrors or crystal monochromators may diffract neutrons in certain wavelength band away from the primary beam.

The relative merits of the different possible neutron filtering techniques depend both on the efficiency of eliminating neutrons outside the production band, and maximizing the number of useable 8.9 Å neutrons. These two concerns are influenced not only by the transmission as a function of wavelength of the filter, but also by the effects of neutron divergence. Cold neutrons are delivered to the UCN trapping experiment by the NG-6 beamline, which uses a ^{58}Ni coated cold neutron guide (see section 4.1). The 8.9 Å neutrons exiting this guide have a divergence of $\pm 1^\circ$, and pass through a defining aperture of 6 cm. Thus, minimization of the distance between the end of the neutron guide and the apparatus is crucial to the success of any filter, as

the 8.9 Å production [137]. Since there are more than 100 times as many neutrons participating in the multi-phonon downscattering process, it is still the case that the signal to background ratio can be substantially improved by restricting the neutrons entering the apparatus to those participating in the single phonon downscattering process.

is the minimization of any additional divergence which the filter may introduce.

Each of the possible wavelength filtration technologies has advantages and disadvantages. Mechanical velocity selectors are commercially available, but at a price of $\sim 10^5$ USD, with a minimum bandwidth of $\Delta\lambda/\lambda \sim 0.1$, and with less than unity transmission (typically about 0.8) of in-band neutrons [85]. Furthermore, velocity filters increase the total path length neutrons must travel between the neutron guide and the trapping region by ~ 0.5 m, and hence decrease the divergence-attenuated flux available for UCN production. Polycrystalline filters are low complexity but offer only limited reductions in the neutron flux (the best widely-available filter material for this work would probably be polycrystalline bismuth, which transmits only neutrons with $\lambda > 6.5$ Å [88]). Monochromators and supermirrors have the advantage that by physically separating 8.9 Å neutrons from the rest of the neutron beam, they offer the opportunity for the two parts of the spectrum to be simultaneously used by separate experiments. Supermirrors are costly, and at 8.9 Å provide only very small angular separations ($\sim 5^\circ$) between the white and 8.9 Å beams. Hence the total path length necessary to separate the two parts of the spectrum is unacceptably long due to divergence-induced flux attenuation.

Long wavelength monochromators with sufficiently high reflectivity may offer the ideal combination of characteristics. While widely used at shorter wavelengths, monochromators are not commonly constructed for 8.9 Å neutrons. The development of an 8.9 Å monochromator is the focus of this chapter.

3.2 Monochromator Design

3.2.1 Diffraction Theory

Monochromators operate by diffracting neutrons within a desired wavelength band out of a primary neutron beam. The basic operation of monochromators can be well understood from simple diffraction principles. The Bragg condition defines the grazing angle θ_B , at which incident waves of wavelength λ will diffract from a lattice with lattice spacing d as:

$$n\lambda = 2d \sin \theta_B \quad (3.1)$$

where n , called the order of diffraction, is any positive integer. This formulation arises readily from conservation of momentum in the kinematic theory of diffraction. For single perfect crystals, the line width for diffraction is extremely narrow ($\Delta\lambda/\lambda \sim 10^{-5}$, the Darwin width, determined by the effect of multiple scattering, as considered in the dynamical theory of neutron diffraction, see for example [90]) around the Bragg peak is diffracted. Due to this extremely narrow bandwidth for diffraction, perfect crystals have wide application as wavelength analyzers, but are not so attractive as monochromators.

In most monochromator applications, the desired output wavelength spectrum is broad enough that an alternative to a perfect crystal must be used. In order to diffract a broader wavelength band than a perfect crystal will allow, one must use a “mosaic crystal.” The mosaic model for crystal structure, originally proposed by C. G. Darwin [91, 92], describes a real crystal as being made up of many microscopic mosaic blocks. Each of these mosaic blocks is a perfect crystal with an orientation that varies slightly from the overall crystal’s orientation. The distribution of mosaic block orientations

has, for most real crystals, a cylindrically symmetric Gaussian distribution function:

$$W(\Delta) = \frac{1}{\sqrt{2\pi}\eta} e^{-\Delta^2/2\eta^2} \quad (3.2)$$

where Δ is the angle between the normal to the mosaic block and the normal to the entire crystal, and $\gamma \equiv \eta/0.424$ is the full width half maximum of the gaussian distribution.

Neutron diffraction experiments are typically used to characterize mosaic crystals. Two particular types of diffraction measurements are of particular importance: “rocking curves” and “ $\theta - 2\theta$ curves.” Both of these types of measurements can be made with double-axis neutron spectrometers, in which a well-collimated, monochromatic, neutron beam is incident on a sample. The grazing angle (θ) of the neutron beam with respect to the sample can be continuously varied. Diffracted neutrons are detected by a detector at an angle ω away from the incident neutron beam.

In a $\theta - 2\theta$ measurement, the distribution of d spacings within a crystal can be determined. A $\theta - 2\theta$ curve is measured by varying θ and ω near some Bragg peak until the signal measured by the detector is maximal. When the signal is maximal, $\omega = 2\theta$, as required by the Bragg condition. Then ω is incremented by $\Delta\omega = 2\Delta\theta$ for each $\Delta\theta$ by which θ is incremented. In this way, the Bragg condition is continually met. The counts recorded by the detector are recorded as a function of θ . Peaks in the $\theta - 2\theta$ curve indicate the presence of d and n combinations which satisfy the Bragg condition (3.1).

Similarly, the mosaic angle, γ , can be determined by measuring the “rocking curve” of a crystal. In this measurement, the diffraction peak of fixed order n is studied for a sample with known d -spacing. The detector is fixed at the angle $\omega = 2\theta_B$, where θ_B is the (known) Bragg angle for the diffraction peak of interest. By varying the angle $\theta = \theta_B + \Delta$ near the Bragg angle, the rocking curve maps out the distribution

$W(\Delta)$. The diffracted intensity as a function of rocking angle, Δ , maps out the relative number of microcrystals within the mosaic crystal which have an angle Δ between their normal and the crystal's overall orientation. Hence the full width half maximum of the rocking curve is also the full width half maximum of the Gaussian distribution W .

Further details of the theory of neutron diffraction are contained in appendix C.

3.2.2 Constraints on d and γ

In order to efficiently transport 8.9 Å neutrons to the trapping apparatus, a neutron monochromator must satisfy certain constraints on its d -spacing and mosaic. The acceptable range of d spacings is constrained by the geometry of the NG-6 beamline where the UCN trapping experiment is performed. The beam layout (see figure 4.2) is compatible with θ_B in the range of 25° to 65°. This is set by the requirement that the neutron trapping apparatus not block access to any of the other beamlines. By constraining λ and θ_B in the Bragg equation, we restrict the product d/n to the range 4.9 – 10.5 Å. Since obtaining high quality crystals with large d spacings is difficult, this effectively puts a lower limit on the usable d spacing of 4.9 Å.

The optimal mosaic of an ideal 8.9 Å monochromator should maximize the number of 8.9 Å neutrons produced in the UCN trapping apparatus. The NG-6 beam contains 8.9 Å neutrons whose momentum vectors make an angle of up to $\pm 1^\circ$ with the beamline. In order to efficiently diffract all of these neutrons, a monochromator should have a large mosaic ($\gtrsim 1^\circ$), so that a significant number of microcrystals within the monochromator will satisfy the Bragg condition for all incident neutron angles up to 1° away from the Bragg angle. On the other hand, as the mosaic angle is increased, the divergence of the diffracted beam increases. Since the start of the

trapping region must be at least 1.5 m away from the monochromator, divergence of the monochromatic beam is a crucial effect in determining the total UCN production. Typically, the desire to maximize the reflected phase space will tend to encourage the use of monochromators with larger mosaics, while the desire to minimize divergence will push towards smaller mosaics. Optimization of these competing parameters is considered using a Monte Carlo simulation (see appendix C). It turns out that for our situation, the optimal monochromator mosaic is about 1° . For a monochromator with a mosaic of 2° , the number of trapped neutrons is only about 15% less than for and otherwise identical monochromator with a mosaic of 1° .

3.3 Long Wavelength Monochromator Materials

As described in the previous section, the geometry of the NG-6 beamline introduces a constraint of $d > 4.9 \text{ \AA}$ for any 8.9 \AA monochromator material of interest. The dearth of naturally occurring materials with such a large d -spacing is primarily responsible for the fact that long wavelength ($> 7 \text{ \AA}$) monochromators are not in wide use. Herein we examine two materials for possible use as long wavelength neutron monochromators: mica and intercalated graphite. We are not aware of any other materials with suitable d -spacing which are easily obtained or fabricated.

3.3.1 Mica

Various types of mica crystals have long been considered, and used, as long wavelength neutron reflectors. Micas are crystals with large d -spacing in which a central alkali (typically potassium) atom separates two layers of aluminosilicate compounds [93]. Depending on the type of mica, the exact nature and structure of the aluminosilicate components may vary. The naturally-occurring muscovite mica

($\text{KAl}_3\text{Si}_3\text{O}_{12}\text{H}_2$) was used for long wavelength neutron studies in the late 1950s [94]. Following the development of techniques for growing synthetic micas, the perfluorinated mica, fluorophlogopite ($\text{KMg}_3(\text{Si}_3\text{Al})\text{O}_{10}\text{F}_2$), also became a possibility. Fluorophlogopite, known industrially as Thermica, has a larger structure factor, and hence higher reflectivity, than the naturally occurring micas muscovite or phlogopite, due to the substitution of Mg for Al in some of the crystal sites, and the substitution of fluorine for hydroxide [93].

Goland and colleagues measured the lattice spacing of fluorophlogopite to be $d = 9.963 \text{ \AA}$ and first proposed its use as a monochromator in 1959 [94]. Later, a prototype ultracold neutron source was built based upon doppler-shifted Bragg scattering from a rotating fluorophlogopite crystal [95, 96, 97]. Additional studies of mica reflectivity were performed as part of the construction of a second generation rotor at Los Alamos [93]. Other studies of fluorophlogopite neutron scattering properties have also been performed [98, 99] and neutron scattering instruments using mica neutron optical elements have been built [100]. Many previous authors have noted the very narrow ($\sim 0.3^\circ$) intrinsic mosaic of mica [93, 94, 98]. Increasing the effective mosaic of mica by producing a fan of many thin mica layers has been tested by several authors [95, 98].

3.3.2 Graphite Intercalation Compounds

Graphite intercalation compounds form a large class of materials in which atoms or molecules are inserted (intercalated) between the layers of hexagonal graphite. Graphite intercalation was first reported by C. Shafthäutl in 1841 [101]. Although interest in Graphite Intercalation Compounds (GICs) has existed for over 150 years, it accelerated dramatically in the 1970s as GICs became widely used tools for study-

ing the two dimensional physics and chemistry of the intercalants [102]. Studies of the structure, dynamics (lattice vibrations and diffusion of intercalants), electronic properties (superconductivity, anisotropic electrical conductivity, etc.) and transport properties have been widely pursued (see [103] for a review of the field).

Various graphite intercalation compounds have been studied for possible use as long-wavelength monochromators. A group at the Institute Laue-Langevin (ILL) studied stage 1 compounds (see below) of potassium, rubidium, and cesium intercalated graphite [104]. Based upon that work, a stage 1 potassium intercalated graphite monochromator was built and used for 7.56 Å neutron monochromatization on the DB21 low resolution biological diffractometer at the ILL [105]. Alkali intercalated graphites have the advantage of large structure factors, and hence high expected reflectivity for long wavelength neutrons. Their disadvantage is that due to the high reactivity of the alkali metals, alkali intercalation compounds must never be exposed to oxygen or water [106, 107]. The graphite-SbCl₅ intercalation compounds have been studied as an air-stable alternative [108]. The graphite-AsF₅ compounds have also been studied for possible use as long-wavelength monochromators, since they can be fabricated with very narrow intrinsic mosaic ($\sim 0.8^\circ$) [109].

Our efforts in studying graphite intercalation compounds as monochromator materials were restricted to potassium intercalated graphite. We chose to work with alkali intercalation compounds due to their high reflectivity and the large body of information about them in the literature. It was expected that potassium, rubidium, and cesium intercalation compounds would perform similarly as monochromators [104]. Potassium was chosen from among these since it is the easiest of the three to intercalate [110]. Also, the neutron absorption cross section of cesium (29 barns) is much larger than that of potassium (2.1 barns) [67], leading to greater attenuation of the




Figure 3.1: Structure of un-intercalated graphite. Carbon atoms (black circles) form covalently-bonded planar hexagons. Adjacent layers are displaced to allow hexagonal close packing. Figure courtesy of H. Zabel [111].

transmitted and reflected beams, and increasing the potential for radiation damage of the monochromator.

The structure of un-intercalated graphite is shown in figure 3.1. Graphite forms a crystal with a layered structure. Within each layer carbon atoms form strong covalent bonds with each other and form a 2D hexagonal lattice. Alternating layers are offset from each other so that the overall structure forms a hexagonally close-packed arrangement of carbon atoms from stacking the two layer configurations in an ABAB form. The bonds between adjacent layers are extremely weak (due only to dipole-dipole interactions).

Intercalated graphites are formed by taking advantage of the weak bonding between layers to insert the intercalant atom or molecule between carbon layers. The introduction of the intercalant pushes the graphite layers apart and shifts the carbon planes so that the carbon atoms of adjacent layers line up. Graphite intercalation




Figure 3.2: Staging of graphite intercalation compounds. A stage n compound has n layers of carbon atoms between each layer of intercalant atoms. The two layers of carbon atoms adjacent to the intercalant layer are equivalent. Adjacent graphite planes have the normal ABAB stacking. The distance between adjacent intercalant layers is indicated for each stage. Figure courtesy of H. Zabel [111].

compounds are formed in discrete stages, with distinct periodicity to the arrangement of intercalant atoms within the graphite host matrix. These stages are characterized by the stage number n , where n is the number of layers of carbon atoms between each layer of intercalant atoms (see figure 3.2).

The stage composition of an intercalated graphite sample depends upon the sample preparation method (see section 3.4). For a monochromator, only samples consisting of a single stage are desirable (to minimize the number of Bragg peaks in the diffracted beam and to maximize reflectivity of 8.9 Å neutrons). For potassium-graphite intercalation compounds, only stages 1 and 2 can be prepared with high purity; all higher stages only exist in mixtures of multiple stages [112].

3.3.3 Materials Selection

Based upon a series of our own measurements, we selected stage 2 potassium intercalated highly-oriented pyrolytic graphite (IHOPG) as our monochromator material. The details of our preliminary studies of mica, stage 1 IHOPG, and stage 2 IHOPG are found in Appendix D. Therein, we measured the mosaic and 8.9 Å reflectivity of samples of the three potential monochromator materials.

Based upon the scattering parameters in the literature and the preliminary measurements described in Appendix D, a Monte Carlo simulation of the scattering from three hypothetical monochromators was constructed (see appendix C for a general description of the Monte Carlo, and section C.2.3 for a discussion of the selection of the scattering parameters). Using this Monte Carlo simulation, estimates were made of the number of UCN which were expected to be trapped by the apparatus described in Chapter 4. The simulated output spectra of the three hypothetical monochromators, relative to the NG-6 input spectrum is shown in figure 3.3

Of the three possible monochromator materials, the fluorophlogopite was ruled out since its mosaic was narrower and its peak reflectivity lower than the stage 2 IHOPG, which has nearly the same d -spacing. Hence for the two, almost equivalent, diffraction geometries, the stage 2 IHOPG offered many more reflected 8.9 Å neutrons than the fluorophlogopite (as shown in figure 3.3).

The choice between the stage 1 and stage 2 IHOPG relies upon more subtle considerations. The stage 1 IHOPG was measured to have a higher peak reflectivity than the stage 2 IHOPG (70% vs. 50%). Nevertheless, a comparable number of trapped UCN are to be expected from each monochromator. This is due to the relative divergences of the two monochromatic beams. Since the stage 2 monochromator has a smaller Bragg angle than the stage 1 monochromator (30.6°, as opposed to 56.3°),




Figure 3.3: The simulated neutron spectra output by three potential monochromator materials is compared to the spectrum of NG-6. The count rates have been scaled so that they are relative to the 8.9 Å flux of NG-6. The Monte Carlo is performed as described in Appendix C.

the divergence of the diffracted beam will be less from the stage 2 monochromator. As discussed above, minimizing the divergence is crucial to efficient UCN production. From the Monte Carlo simulations, it was determined that the expected lower reflectivity of the stage 2 IHOPG relative to the stage 1 IHOPG is made up for by its lower divergence. As will be seen below, the actual peak reflectivity of the stage 2 IHOPG monochromator was much higher than initially expected, making its advantage over the stage 1 IHOPG even clearer.

Since the two IHOPG monochromators were expected to produce about the same number of trapped UCN, the decision of which material to use was based on several secondary factors. First, the geometry of a monochromatic beam, with a take-off angle of $2\theta_B = 61.2^\circ$, was a better fit for the layout of NG-6 (see figure 4.2). Second,

we expected the fabrication of IHOPG samples with small mosaic to be easier for stage 2 samples than stage 1 samples, based on the mosaics measured for samples Z1 and Z3, the experience of our collaborators [110], and the literature [113]. Therefore, the conclusion was to fabricate an 8.9 Å monochromator from a stage 2 potassium-graphite intercalation compound.

3.4 Fabrication of IHOPG

The standard technique used for intercalation is the “two-bulb technique” due to A. Herold [114]. In this technique, an evacuated glass ampule with two connected, but physically distant, bulbs is prepared. The graphite to be intercalated and the intercalant material are each located within one of the two bulbs. The ampule is then heated in a two zone furnace so that the graphite and the intercalant equilibrate at different temperatures.

The temperature of the intercalant determines the vapor pressure of a gas of intercalant atoms or molecules which fills the ampule. Intercalation proceeds as the system approaches a thermodynamic equilibrium, in which the chemical potential of the intercalant atoms “dissolved” within the graphite equals the chemical potential of the intercalant vapor [111]. The chemical potential of the intercalated compound is a function of the temperature of the graphite bulb, T_G , while the chemical potential of the vapor is a function of the intercalant temperature, T_K . Hence the equilibrium ratio of intercalant atoms to carbon atoms in the intercalated graphite is a function of the temperature difference, $T_G - T_K$. The phenomenon of discrete stages, which appears structurally as described above, is due to the existence of temperature plateaus in the ratio of carbon/intercalant.

The nominal temperature dependence of stage formation in potassium interca-




Figure 3.4: Nominal temperature dependence of stage formation in potassium intercalated graphite. Three stage plateaus are shown: KC_8 is stage 1, KC_{24} is stage 2 and KC_{36} is stage 3. Figure courtesy of H. Zabel [110].



Figure 3.5: The design of a glass ampoule used for graphite intercalation is shown. The graphite and potassium are contained in the two bulbs.

lated graphite is shown in figure 3.4 [110]. The precise temperature differences corresponding to the plateaus for each stage depend on the detailed geometry of the intercalating ampoule and the graphite sample. When preparing stage 2 potassium GICs it is known that the highest stage purity (minimal contamination of stages 1 and 3) is obtained close to the transition between stages 2 and 1 [112]. From this constraint and figure 3.4, the expected target value for $T_G - T_K \sim 115$ K.

A typical intercalation cell used in this work is shown in figure 3.5. The cells were




Figure 3.6: The vacuum system used for loading potassium into intercalation cells is shown.

blown and loaded at the NIST Optical Shop. Since each cell is custom made, no two are identical. Nominal dimensions were provided to the glass blower, however, and the variation from cell to cell was not more than a few cm in any given dimension. The cell is fabricated from lengths of standard diameter glass tubing. The nominal dimensions of the cell are: the graphite bulb is 12.7 cm long; the tube connecting the two bulbs is 19.1 cm long; the potassium bulb is 5 cm long. The inner diameter of the tube containing the HOPG sample is just large enough so that the graphite can move freely after intercalation (when its thickness has expanded by a factor of 1.6). The total assembly must fit inside of a 6.4 cm bore furnace (see below).

A multistep process was used for encapsulating the graphite and loading the potassium into the cell [115]. The glass cell, as shown in figure 3.5, is blown with the HOPG substrate sealed inside. A tube connects the potassium bulb to an auxiliary vacuum system, shown in figure 3.6. The system is initially prepared with an empty “auxiliary potassium bulb,” as shown. While flowing argon gas through the vacuum system, the top of the auxiliary potassium bulb is opened, a commercially obtained² 5 g ampule of potassium is broken open and dropped into the auxiliary potassium bulb, and the

²99.95% purity K, Alfa Aesar, Ward Hill, MA, (<http://www.alfa.com>).

top of the auxiliary bulb is re-sealed. By attaching the potassium to the system in this manner, the time of exposure of the potassium to air is minimized. Nevertheless, an oxide layer forms on the exposed surface of the potassium.

After the potassium is attached to the vacuum system, the system is prepared for loading the potassium into the cell. Using a diffusion pump with a liquid nitrogen cold trap, the system is pumped to a base pressure of $< 5 \times 10^{-5}$ torr and then flushed with 1 atmosphere of argon gas. The pumping and flushing procedure is repeated a total of 5 times. The system is then pumped to a base pressure of about 2×10^{-6} torr. While pumping at this low pressure, the entire cell is flashed with a torch to bake off any water vapor which may have adhered to the glass or HOPG. After pumping for about 30 minutes, the cell is then backfilled to 500 torr with argon gas. A torch is then lightly applied to the auxiliary potassium bulb. After about 3 minutes of heating the potassium becomes warm enough to liquify. The potassium then runs down the inclined tube and fills the potassium reservoir of the cell. The potassium oxide formed while loading the potassium ampule into the vacuum system freezes out along the walls of the loading tube and does not reach the cell reservoir. After loading the cell is pumped to a pressure of $\sim 10^{-6}$ torr for a period of an hour before the cell is pinched off from the vacuum system.

The intercalation cell is heated in a two-zone furnace, as shown in figure 3.7. The furnace used is a 6.4 cm bore, two-zone furnace³. A thermocouple measures the temperature of each zone and is used as an input to an on-off temperature controller⁴. The temperature controllers turn on and off variacs, which supply current to the

³Model 10372Z, 6.4 cm bore, 115 V, 8.5 A/zone. Manufactured by Marshall Furnace, which is now a subsidiary of Thermcraft, Inc., Winston-Salem, NC (<http://www.thermcraftinc.com/marshall.html>).

⁴Model 2104-TO120, Chromalox, Pittsburg, PA (<http://www.chromalox.com>).




Figure 3.7: The setup used for intercalation is shown. The glass ampoule is placed in a two-zone furnace. Glass wool is used to block convection currents and reduce the thermal conductivity between zones.

furnace's heaters. The voltage from each of the variacs is adjusted so that when the zones of the furnace are at their stable operating temperatures the variacs will be on about 75% of the time. This reduces temperature fluctuations of each zone to about ± 2 K. Glass wool is also packed into the ends of the furnace and between the two zones in order to minimize convection and maximize the temperature stability of each zone. The intercalation cells are loaded into the furnace after the furnace is heated to the desired temperatures.

Since the exact dimensions of each cell, and its placement within the furnace, varies from cell to cell, the temperature of the cell's bulbs may vary from run to run even if the temperature controllers are programmed to produce the same temperature in each zone of the furnace. The position dependence of the temperature within the furnace is shown in figure 3.8. As can be seen, the temperature is not completely uniform within each zone. Therefore, the programmed temperature is corrected using auxiliary thermocouples. A thermocouple is attached to each bulb of




Figure 3.8: The position dependence of the temperature within the two-zone furnace is shown. The furnace settings are $T_G = 320\text{ }^{\circ}\text{C}$ and $T_K = 200\text{ }^{\circ}\text{C}$. The temperature was measured by moving a thermocouple along the center axis of the furnace. Glass wool sealed the two ends of the furnace, which contained nothing else.

the intercalation cell. The junction of the thermocouple is positioned so that it is at the center of the respective sample (HOPG or potassium metal). The programmed temperatures are then adjusted during the first few hours of intercalation until the measured temperature of the center of each bulb reaches the desired value.

The optimal parameters for intercalation were determined iteratively by intercalating samples, measuring their stage purity, intercalating further, and adjusting the parameters for further intercalations (see section 3.5.1). It was determined that samples with high stage 2 purity (no measurable stage 1 or stage 3 diffraction peaks) could be fabricated by heating the sample with $T_K = (206 \pm 2)\text{ }^{\circ}\text{C}$, $T_G = 318 \pm 2\text{ }^{\circ}\text{C}$, and $T_G - T_K = 112 \pm 3\text{ }^{\circ}\text{C}$ (this last constraint is the most important). Three samples were intercalated with these settings and all displayed high stage 2 purity after 4–5 days of intercalation. At the conclusion of intercalation, the cell is removed from the furnace. The potassium side of the cell is removed first and allowed to cool while the IHOPG remains in the furnace. Once the potassium solidifies the rest of the cell is

removed. This procedure prevents potassium vapor from solidifying on the IHOPG⁵.

3.5 Sample Characterization

During the fabrication process, IHOPG samples were characterized by thermal neutron diffraction. Using the triple-axis spectrometer⁶ on thermal beam port BT-7 at the NCNR, rocking curve and $\theta - 2\theta$ curves were measured. An advantage of using thermal neutron diffraction is that thermal neutrons penetrate reasonably well through the borosilicate glass ampules used for intercalation. Hence diffraction measurements could be made while the samples were still in their intercalation ampules.

3.5.1 Stage Purity

In order to determine the stage purity of the IHOPG samples we were fabricating, we measured $\theta - 2\theta$ curves (as described above) for each sample. As described above, such a measurement is used to determine the distribution of d within a crystal. Since

⁵On a few samples where this procedure was not followed, miniscule amounts of potassium condensed on the IHOPG. While they were not directly observable, their presence was inferred by the physical intercalation which took place. Spots on the surface of the sample turned gold, indicating the presence of stage 1 IHOPG. The stage 1 contamination was only on the surface, however, and could not be detected by neutron diffraction.

⁶Although the BT-7 instrument is a triple-axis spectrometer, it was used in double axis mode. A triple axis spectrometer has independent rotation stages for setting the sample angle (θ), the detector angle (ω), and the rotation of an analyzer crystal at the detector position. Incident neutrons have a fixed wavelength, $\lambda = 2.4649 \text{ \AA}$ in this case, and direction, so their wavevector \vec{k}_i is completely determined. By varying θ and ω , the momentum transfer $\Delta\vec{q} = \vec{k}_f - \vec{k}_i$ of the scattering process can be constrained. The analyzer crystal allows wavelength selection (or equivalently, $|\vec{k}_f|$ selection) before neutrons are detected. Since we were only interested in elastic scattering from the monochromator, the analyzer crystal was not used, and the instrument functioned as a double-axis spectrometer.

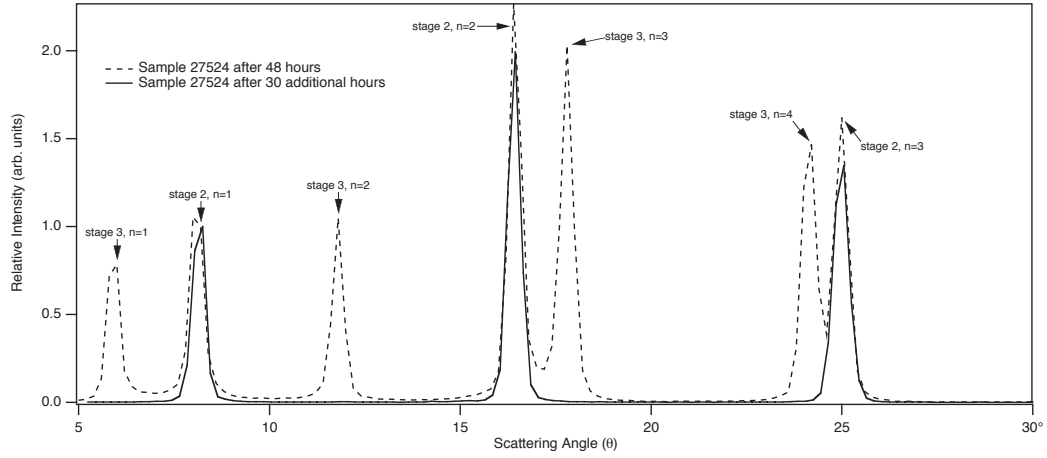


Figure 3.9: $\theta - 2\theta$ diffraction curves are shown for sample #27524 after two heating periods. The sample was initially intercalated for 48 hours with $T_K = 202$ K and $T_G = 323$ K. After the initial heating phase, stage 3 peaks are of roughly equal intensity to the stage 2 peaks. The sample was then intercalated for an additional 30 hours with $T_K = 205$ K and $T_G = 316$ K. Following the second intercalation, stage 3 diffraction peaks are no longer evident. For both curves, the intensity has been normalized so that the relative intensity of the stage 2, $n = 1$ peak is 1.0.

the values of d are known for each of the stages of potassium-intercalated graphite, we are effectively scanning for the presence of different stages.

This diagnostic technique was used to determine the optimal temperatures for intercalation, and the necessary length of intercalation. For example, the $\theta - 2\theta$ curve for sample 27524 is shown in figure 3.9 after two successive periods of intercalation. After initially intercalating the sample for 48 hours with $T_K = 202$ K and $T_G = 323$ K, both stage 2 and stage 3 diffraction peaks are clearly evident. Based upon this finding, it was necessary to intercalate the sample for an additional period, and with a smaller temperature difference $T_G - T_K$. After intercalating for an additional 30 hours with $T_K = 205$ K and $T_G = 316$ K, no stage 3 diffraction peaks can be seen. In this manner, high stage purity was ensured for all of the samples used in the final monochromator. Furthermore, the appropriate temperatures for intercalation were determined, so that

the last few samples that were intercalated reached high stage 2 purity after only a single heating period with the parameters discussed above (see section 3.4).

After the stage purity of a sample was verified, it was removed from its intercalation ampule and stored in a Conflat nipple with a glass window. From this storage cell, the IHOPG sample could be transferred to the compression rig, pinwheel sample-holder, or monochromator holder (see below for descriptions of each) for measurements. All sample transfers were performed in a glove box⁷ filled with high purity helium. The glove box atmosphere was maintained at < 5 ppm O_2 and < 5 ppm H_2O by a getter cylinder.

3.5.2 Sample Mosaic

The mosaic of each sample was determined (using rocking curve measurements) for two sample orientations: with the long axis of the sample vertical (“vertical mosaic”) or with the short axis of the sample vertical (“horizontal mosaic”). The mosaic angles of the samples were measured in a variety of sample holders. While still in the glass, the vertical mosaic of most samples was measured. The horizontal mosaic could not be measured with the samples still sealed in their intercalation ampules, since the sample position was unstable in this orientation. The horizontal and vertical mosaics of multiple samples could be quickly measured using a “pinwheel sampleholder” (see figure 3.10). The pinwheel sample-holder holds 12 samples, only one of which is in the neutron beam at any time. By simply rotating the sample holder, the sample to be measured is selected. The only drawback of this sample holder is that for vertical

⁷The glove box used was a VAC Model HE-493, Vacuum Atmospheres Company, Hawthorne, CA (<http://www.vac-atm.com/>). The system was equipped with moisture and oxygen analyzers (models AM-2031 and AO-316-C, respectively), and a DRI-TRAIN purification/regeneration cylinder.




Figure 3.10: The pinwheel sample holder is shown. Samples are loaded into the sample holder in the inert atmosphere of a glove box. The inert atmosphere is maintained within the sample holder by the rubber o-ring near the outer edge.

mosaic measurements it was necessary to study the stage 2, $n = 2$ reflection. The incident angle for the stage 2, $n = 1$ reflection ($\theta = 8.1^\circ$ for $\lambda = 2.4649 \text{ \AA}$ neutrons) is small enough that the beam could not be prevented from hitting the two samples adjacent to the target sample.

In studying the mosaics of the early samples, an unexpected characteristic was discovered. While the horizontal mosaics of the samples were simple Gaussian peaks, as expected, the vertical mosaic of almost every sample showed a double peak (see figure 3.11) The double peak indicates that there are two groups of crystal planes within the sample, each of which has a comparable mosaic, but which are separated from each other by a fixed angle (about 1.5°). A photograph of the surface of one sample showing this double peaked mosaic is shown in figure 3.12. From the variation in brightness of the sample's surface, the same physical condition which causes the two peaks in the rocking curve can be seen. The sample is much brighter in the top half than in the bottom. This is because the sample's shiny surface acts like a mirror.

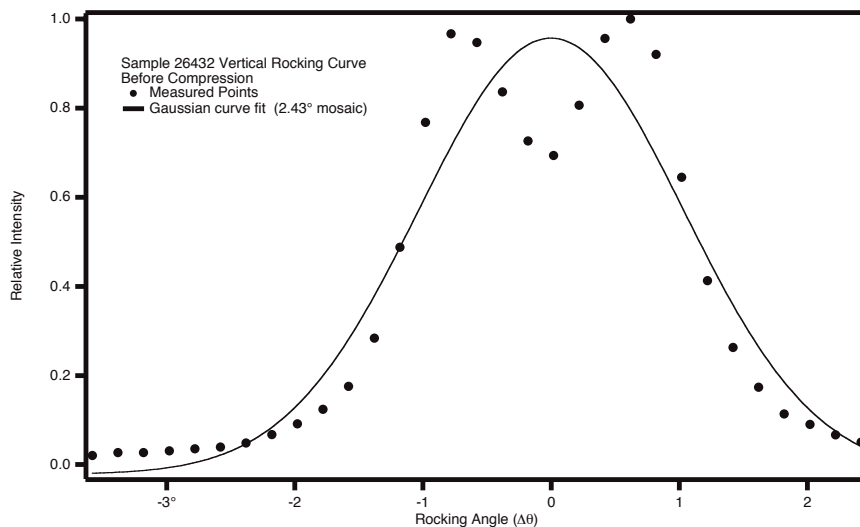


Figure 3.11: The rocking curve of sample 26432 in the vertical orientation is shown after intercalation, but before compression.




Figure 3.12: A photograph of one side of an IHOPG sample showing the bend described in the text is shown. The reflection of the illuminating light from one side of the “trough” is much brighter than the other, indicating that the light and dark areas on the sample’s surface are not parallel.




Figure 3.13: The rig used for compressing the samples is shown. The construction of the rig is described in the text.

The upper half of the sample is reflecting the incident light towards the camera, while the lower half reflects the light elsewhere. Hence at least the surfaces of the upper and lower halves are not parallel. Both the visible appearance and the rocking curve are consistent with a distortion of the sample so that the IHOPG, viewed along the long axis, forms a “v” shape. This “v” distortion of the samples is evidently caused by the stresses of the intercalation process, and was seen to some extent in each sample.

Since the distortion of the samples was macroscopically visible, the idea of compressing the samples to remove the distortion was suggested by H. Zabel [110]. A compression rig (shown in figure 3.13) was constructed from standard vacuum components. An aluminum ISO-80 to ISO-63 reducer is sealed at each end to a stainless steel blanking flange. A bellows valve is attached to the ISO-80 blankoff. In place of the valve’s seal is a metal plate with an indentation, in which an IHOPG sample is rested. By “closing” the valve, the sample is compressed against the ISO-63 blanking flange. The valve was tightened by hand, and during compression it was noted that

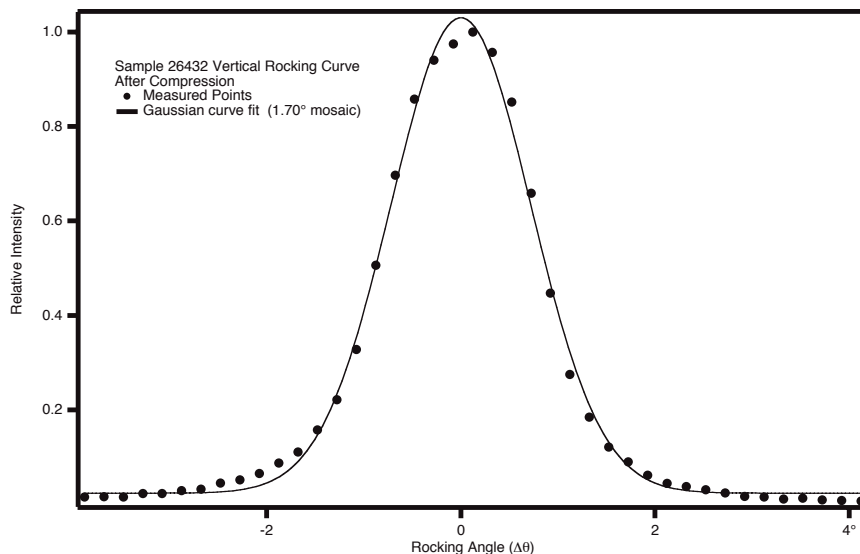


Figure 3.14: The rocking curve of sample 26432 in the vertical orientation is shown after compression.

the resistance to tightening increased discontinuously at a certain point. Some samples were compressed with the ISO-63 blanking flange replaced with a Pyrex Window. In these cases, it was observed that the sharp decrease in resistance to compression corresponded with a change in the area of optical contact between the sample and the window. Before the sharp increase in resistance, the sample only made optical contact with the window at the edges or the center (depending on which side of the “v” was towards the window). After the sharp change in resistance, optical contact was made between the window and the entire surface of the sample. Samples were typically left under compression for at least 24 hours. Based upon the sharp change in resistance and optical contact, it is believed that the samples plastically deformed at a certain compression, and that further, static compression, had no additional effect.

Once the compression was removed, the mosaics were remeasured. The vertical mosaic of sample 26432 after compression is shown in figure 3.14. Whereas before compression two peaks were evident in the rocking curve, after compression only a

single peak is evident. Therefore we conclude that compression removed the distortion of IHOPG samples and restored a single mosaic. All of the samples produced were compressed once, even if only a single peak was seen in the first rocking curve measurement. Compression made no significant change in the mosaics of the samples which already had a single peak in their rocking curves. For some of the samples, compression reduced the separation of the two peaks in the rocking curve but did not eliminate it.

3.6 Assembled Monochromator

3.6.1 Sample Characteristics

The nine highest quality IHOPG samples produced were used in the final assembled monochromator. The samples were chosen for high stage 2 purity and smallest mosaics. The measured characteristics of the samples used are shown in table 3.1, along with the characteristics of 6 samples which were not used. As can be seen, all but one of the samples used in the monochromator has minimal stage 3 contamination. Also, only one of the samples used has a double-peaked rocking curve. The samples which were not used all have large stage 3 contamination, multiple peaks in the vertical rocking curve, or poor reflectivity.

3.6.2 Monochromator Holder

The monochromator is assembled from 9 tiled crystals. The crystals are laid out 3 by 3, so that the total monochromator is 6 cm tall and 15 cm wide. Since the Bragg angle at 8.9 Å is 30.6°, the monochromator presents a width of 7.63 cm to the incident neutron beam. The layout of the samples used within the monochromator

Table 3.1: The measured characteristics of 15 IHOPG samples are shown. The top nine samples in the table are those used in the final monochromator assembly. For each sample, the horizontal (β_h) and vertical (β_v) mosaics are shown. The number of peaks observed in the vertical rocking curve after compression is also indicated (N_{peak}). The relative thermal neutron reflectivity, R_{th} is computed as the number of counts at the Bragg peak for each sample in the pinwheel geometry, normalized to the most intensely reflecting sample. The measured 8.9 Å reflectivity, $R_{8.9}$, is shown for the samples used in the final monochromator (see section 3.7). The stage purity of each sample is indicated by $P_{3/2}$ which is the ratio of the peak heights of the stage 3, $n = 3$ and stage 2, $n = 1$ diffraction peaks on the $\theta - 2\theta$ curve. The $n = 3$ peak is the most intense of all of the stage 3 peaks.

Sample #	$\beta_h[^\circ]$	$\beta_v[^\circ]$	N_{peak}	R_{th}	$R_{8.9}$	$P_{3/2}$
26432	1.49	1.70	1	0.78	0.86	0.00
27524	1.24	1.91	1	0.95	0.91	0.00
27523	1.47	2.13	2	0.82	0.79	0.00
26460	1.50	1.90	1	0.80	0.85	0.01
27549	1.12	1.47	1	1.00	0.85	0.00
27537	1.31	1.43	1	0.92	0.73	0.00
27580	1.78	1.53	1	0.76	0.91	0.00
26472	1.42	1.70	1	0.88	0.90	0.00
26520	1.71	2.04	1	0.90	0.81	0.84
26519	2.27	1.94	1	0.88	0.75	0.49
27539	1.68	2.00	2	0.80		0.00
26469	1.91	2.09	2	0.66	0.86	0.06
27536	1.57	2.31	1	0.45	0.78	0.22
27533	1.80	2.31	1	0.36	0.39	3.14
27538	1.56	2.41	1	0.22		2.00




Figure 3.15: The placement of individual IHOPG samples within the monochromator assembly is shown.

is shown in figure 3.15.

The IHOPG samples which make up the monochromator must remain in an inert atmosphere at all times. A vacuum-tight aluminum monochromator holder, shown in figure 3.16 was thus constructed. The monochromator pieces are pressed against a flat aluminum window by aluminum bars. Each bar holds down three monochromator pieces. A layer of 1 mm GoreTex GR gasket material is placed between the aluminum bars and the monochromator pieces to distribute the bar's force across all three samples. The aluminum window has thickness 3.0 mm. The sample holder is made from aluminum to minimize activation and to maximize transmission through the windows. The transmission of 8.9 Å neutrons through the aluminum window is 0.99 (total transmission for both passes which a reflected neutron must make). The front piece of the monochromator assembly makes an indium seal to the back piece, which has a 1.9 mm aluminum window. A total of 46 8-32 titanium bolts are used to

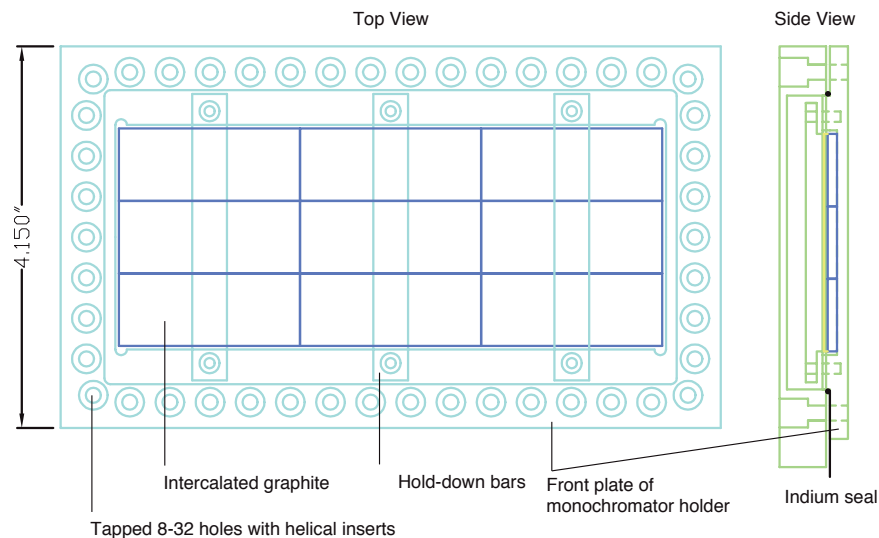


Figure 3.16: The monochromator holder is shown.

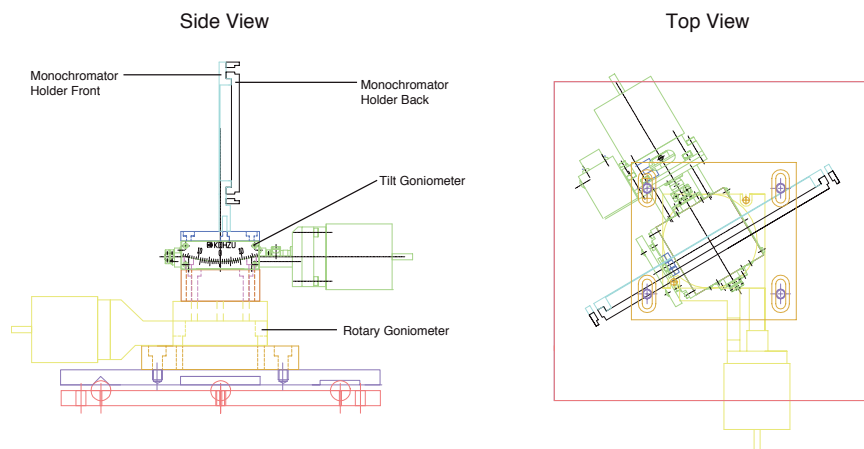


Figure 3.17: The monochromator assembly is shown. The monochromator is mounted on tilt and rotary goniometers.

make the indium seal. The tapped holes are fitted with stainless steel helical inserts. The integrity of the indium seal is indicated by the consistency between the 8.9 Å reflectivity measurements made more than 6 months apart (see section 3.9.1).

The monochromator holder is mounted on two goniometers as shown in figure 3.17. A rotational goniometer⁸ driven by a stepper motor allows computer control of the

sample angle, θ . The tilt of the monochromator is adjustable using a stepper-motor-driven tilt goniometer⁹.

3.7 Time of Flight Measurements

3.7.1 General Considerations

Consider a chopper of radius r [cm] with N slits of width w [cm], a rotational frequency f [Hz], a distance l [m] between the chopper and the detector of diameter d [cm], a beam divergence of θ_d . The time of flight t [s] for a neutron of wavelength λ [Å] will be

$$t = \frac{l}{v_n} = \frac{l\lambda}{4000}. \quad (3.3)$$

The resolution of the time of flight measurement is determined by the opening time of slit. The opening time of the slit is set by

$$\Delta t = \frac{w}{2\pi r f}. \quad (3.4)$$

So the fractional resolution of the chopper is

$$\frac{\Delta\lambda}{\lambda} = \frac{\Delta t}{t} = \frac{4000w}{2\pi r f l \lambda}. \quad (3.5)$$

In order to maximize the resolution of the chopper the product lf should be maximized. Opposing the goal of maximum resolution is the additional constraint that consecutive frames do not overlap. This leads to a maximum value for the product

⁸Model M-037 rotation stage, manufactured by Physik Instrumente, Karlsruhe, Germany (<http://www.physikinstrumente.com>).

⁹Model SA05-12 50 × 50 mm brass crossed roller guide tilt goniometer, Kohzu Precision Co., Ltd., Tokyo, Japan (<http://www.kagaku.com/kohzu/english.html>).

lf as a function of the maximum and minimum wavelengths present in the neutron spectrum

$$\frac{l(\lambda_{max} - \lambda_{min})}{4000} \leq \frac{1}{fN}. \quad (3.6)$$

Additional considerations in chopper design are that the signal should be maximized and that the detector must be able to capture all of the neutrons. For a beam divergence of $\pm\theta_d$ and a detector of diameter d this sets the maximum l according to

$$l \leq \frac{d}{2 \sin \theta_d}. \quad (3.7)$$

The signal goes roughly as

$$\frac{w^2 \Delta t}{T} = \frac{w^3 N}{2\pi r}. \quad (3.8)$$

Thus, in order to maximize the signal w should be maximized.

3.7.2 NG-6 Setup

The setup used for making time of flight measurements at NG-6 is shown in figure 3.18. The chopper sits between the end of the NG-6 shielding and an additional shield wall (steel frame filled with steel shot and polyethylene). The chopper disk is made of ${}^6\text{LiF}$ -loaded plastic glued to a phenolic disk. The radius of the disk at the aperture position is 15 cm. The disk has 4 apertures, 2 of which were covered up. The defining aperture, upbeam from the chopper, has diameter 3.8 mm, and is slightly larger than the slit width. The chopper frequency was 60 Hz. The total TOF path length $l = 1.9$ m.

We thus have a chopper with $w = 0.38$ cm, $r = 15$ cm, $N = 2$ and a rotational frequency of $f = 60.0$ Hz. Using a 12.7 cm diameter detector and considering a maximum beam divergence of 1.5 degrees (the maximum acceptance of the apparatus when placed with $z_{helium} = 120$ cm, see Appendix C) we can calculate the rest of the

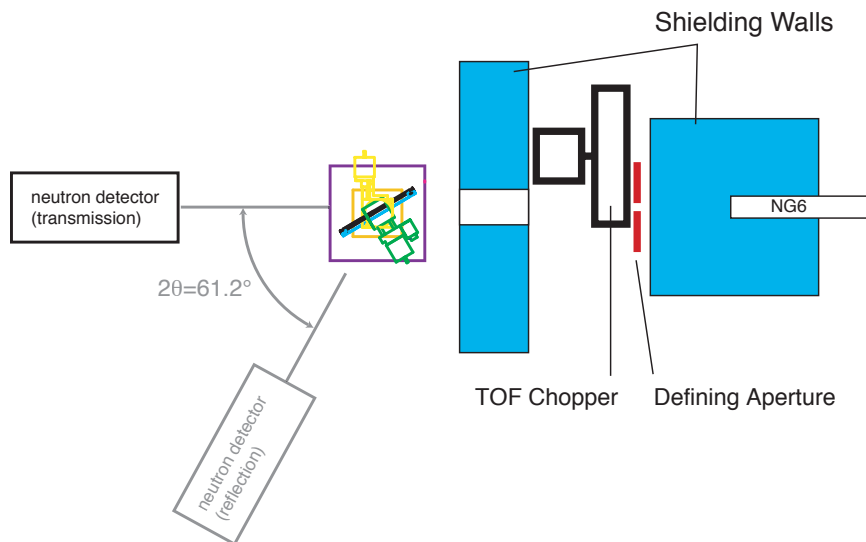


Figure 3.18: The time of flight setup used at NG-6 for making wavelength dependent diffraction measurements is shown.

relevant parameters. From (3.7) we get a maximum path length l of 2.4 m. With $\lambda_{max} - \lambda_{min} = 10 \text{ \AA}$, (3.6) sets a maximum length of 3.3 m. With $l = 1.9 \text{ m}$ we can achieve a resolution (from (3.5)) of 1.6% at 8.9 \AA .

Neutrons are detected using a scintillator coupled to a PMT. The scintillator is a 12.7 cm diameter piece of $^6\text{Li}/\text{ZnS}(\text{Ag})$ scintillator¹⁰. The scintillator is coupled to a Burle 8854, 12.7 cm bialkali PMT. The PMT is biased at +1670 V. The anode signal is passed through a $\times 10$ voltage preamplifier, a NIM discriminator, and a gate and delay box (which converts NIM to TTL pulses and outputs a long enough signal so that only one count is observed per neutron capture), before being counted by a multi-channel scaler. The MCS is triggered by a photodiode which watches an LED on the other side of the chopping disk. The photodiode puts out a TTL pulse, triggering the MCS, when a hole passes between the LED and the photodiode. There

¹⁰Bicron BC-702 Thermal Neutron Scintillator. Bicron, a division of Saint-Gobain Industrial Ceramics, Newbury, OH (<http://www.bicron.com/bc702.htm>). The BC-702 scintillator is based on that described in [116]

is only one such hole on the disk. It is at a slightly larger radius than the neutron slits, and at roughly the same angle as one of the slits. The exact phase relationship between the trigger and the slit opening time is not known and is not assumed in analyzing the data.

In order for the data obtained using this setup to be interpreted in a useful manner, one must understand the linearity and the efficiency of the detection system. The first concern is that the detector must be linear, yielding a number of counts proportional to the number of neutrons captured, though perhaps with the efficiency a function of wavelength. Nonlinearity of the detection system can be caused by saturation of the PMT, if the countrate is too high for the given gain and depletion of the cathode or the dynode chain occurs. In order to guard against PMT saturation the PMT is operated at low gain (low bias voltage).

Linearity is also sensitive to the choice of output width from the gate and delay. When the PMT's output pulses were observed on an oscilloscope it was noted that there are two characteristic pulse shapes: 25 ns wide pulses and $\sim 1.5 \mu\text{s}$ wide pulses. It is likely that the narrower pulses are due to gammas and dark counts while the long pulses are due to neutron capture in the scintillator. Since the discriminator triggers whenever there is a level crossing below (for negative pulses) the threshold, the discriminator may output multiple pulses for each neutron capture, due to fluctuations in the falling PMT signal causing multiple level crossings¹¹. Therefore if the gate and delay output is set too short, it may output multiple pulses per neutron capture. Conversely, if the gate and delay output is too wide then pulse pileup may occur at

¹¹Multiple level crossings would not necessarily make the system non-linear, unless there is a difference in the time profile of the scintillator output as a function of neutron wavelength. Such a difference in the pulse shape is not inconceivable given the large change in the neutron absorption length of the scintillator, due to the $1/v$ dependence of the neutron capture cross section of lithium.




Figure 3.19: The neutron wavelength spectrum of NG-6, measured using a time of flight technique, is shown. The spectrum is measured with three different output widths from the gate and delay which outputs pulses to the multi-channel scaler (see text). Spectra are shown with and without an attenuator. The attenuated spectra are shown scaled up by a factor of 5.6.

shorter wavelengths where the countrate is high. This would introduce a nonlinearity in the system.

In order to study linearity effects, TOF spectra of the neutron beam were taken under identical conditions, varying only the output width of gate and delay circuit. Two sets of measurements were made, with and without a 1 cm thick book attenuating the neutron beam. The six TOF spectra determined in this manner are shown in figure 3.19. The spectra taken with the attenuator in are shown scaled up by a factor of 5.6. While the spectra are in good agreement at long wavelengths ($\lambda > 6 \text{ \AA}$), there are large variations between the spectra measured at shorter wavelengths. Hence for long wavelength measurements, such as determining the 8.9 \AA reflectivity

of the monochromator, any of the output widths would be suitable. In order to make meaningful comparisons between the relative fluxes at different wavelengths, however, the output pulse width is clearly important. As indicated above, when the output pulse width is too long ($1.3\ \mu\text{s}$), pulse pileup causes undercounting at short wavelengths. With the attenuator in place, pulse pileup no longer occurs with the $1.3\ \mu\text{s}$ pulsewidth. Hence, all of the attenuated spectra are in good agreement at all wavelengths, and, when scaled up, agree with the 283 ns spectrum taken without an attenuator. We can thus conclude that with a 283 ns output pulse width, the detector is in the linear regime for the countrates observed at all wavelengths in the spectrum. Unless otherwise noted below, TOF spectra were thus measured using a gate and delay output pulse width of 250 ns.

3.7.3 TOF Results

With the monochromator assembled as described in section 3.6, the $8.9\ \text{\AA}$ reflectivity was measured as a function of position across the face of the monochromator. The layout of samples within the monochromator is shown in figure 3.15. The TOF chopper was mounted on vertical and horizontal translational stages so that the 3.8 mm aperture on the front of the chopper could be scanned across the beam. Since the TOF setup has only the single defining aperture, neutrons of all wavelengths and divergences which pass through the aperture are collected by the detector. Since the wavelength and divergences vary as a function of aperture position, the reflectivity measured for different positions on the monochromator depends not only on the inherent properties of the monochromator samples, but also on its placement in the holder, and hence the distribution of neutron momenta which reach it.

The $8.9\ \text{\AA}$ reflectivity was determined for each aperture position by measuring




Figure 3.20: Time of flight spectra measured with the time of flight aperture in line with the center of the monochromator and the center of the neutron beam. The neutron spectrum transmitted through the monochromator is measured on ($\theta = 30.6074^\circ$) and off ($\theta = 45.6074^\circ$) the 8.9 Å Bragg condition. The gate and delay output width is 1.5 μs .

the TOF spectrum with the detector in the transmission geometry. Two TOF spectra were made, one with the monochromator aligned so that the Bragg condition is satisfied for 8.9 Å neutrons ($\theta = 30.6074^\circ$), and another with the monochromator rotated off-Bragg ($\theta = 45.6074^\circ$). A typical pair of TOF spectra, in this case with the TOF aperture in line with the neutron beam and the center of the monochromator, are shown in figure 3.20. The set of measurements of 8.9 Å reflectivity as a function of wavelength were made with a gate and delay output pulse width of 1.5 μs . By comparison with figure 3.19, one can see that the detector is saturating at low wavelengths. The system should still be fully linear at 9 Å, however, so the measured reflectivities are unaffected. Background counts are subtracted from each

run by subtracting the average counts per channel over the wavelength band from 16 to 17 Å, where the neutron spectrum is observed to be flat (and should be zero).

The conversion between neutron arrival time and wavelength is performed using only information within the measured spectra themselves. Three reference points are identified: the position of a dip in the spectrum at 4.8 Å (due to an up-beam monochromator belonging to another instrument), the peak of the primary Bragg reflection (determined as the peak of a Gaussian curve fit to the difference between on-Bragg and off-Bragg spectra), and the minimum point in the second order ($\lambda/2$) Bragg peak. The starting wavelength and wavelength per channel can then be uniquely determined by requiring the 4.8 Å dip to be at 4.8 Å and the wavelength of the primary Bragg peak to be twice that of the second-order Bragg peak.

For each pair of TOF spectra, the monochromator's reflectivity is determined over the wavelength band 8.8 to 9.0 Å, the wavelength band which contributes to UCN production for the largest imaginable trap depth. The reflectivity at the center of the monochromator is $(85.4 \pm 0.6)\%$. The reflectivities measured with the beam centered on other samples is shown in table 3.1. The beam is roughly 2 cm in diameter at the monochromator position. Hence the measured reflectivities ascribed to each sample may depend somewhat on the reflectivities of adjacent samples.

Another measurement of the reflectivity of the center of the monochromator is shown in figure 3.21. These data were taken with the output width of the gate and delay set to 700 ns. The absence of saturation of the detector at this output width can be seen by comparing the spectra in figures 3.23 and 3.24, taken with 700 ns and 250 ns (confirmed to be saturation-free previously) pulse widths, respectively. The dips in the transmission-mode measurement correspond exactly in size with the peaks in the reflection-mode measurement, indicating that the system is not saturated, and that




Figure 3.21: Time of flight spectra for the center of the monochromator are shown in transmission and reflection. The neutron beam spectrum in the absence of the monochromator is shown for comparison. The output width of the gate and delay is 700 ns.

the dips in the transmission curve are due to diffraction of neutrons. A comparison between the transmitted and empty spectra indicates that the total transmission of the monochromator assembly is 0.81.

3.8 Neutron Wavelength Filtering

In addition to reflecting 8.9 Å neutrons, the monochromator also reflects neutrons which satisfy the Bragg condition at higher order ($n \geq 2$). As can be seen in figure 3.20, the number of neutrons reflected in the $\lambda/2$ and $\lambda/3$ (with respect to $\lambda = 8.9$ Å) peaks is significant. In order to achieve the maximum possible improvement of signal/background, it is desirable to filter out these higher order wavelength

components from the monochromatic beam. Several techniques are available for such filtering. They are, in fact, the same techniques which were originally considered for separating the 8.9 Å neutrons from the white spectrum. Once the 8.9 Å monochromator is in use, however, the constraints on additional filters are much less. We only require that we filter out neutrons near 4.5 Å and below about 3 Å.

3.8.1 Polycrystalline Filters

The simplest technique which may be used for neutron wavelength filtering is diffraction from polycrystalline blocks. Such blocks contain many small crystals, with lattice spacings d_{cutoff} at random orientations. For a sufficiently thick crystal, any neutron with wavelength λ such that

$$\lambda \leq 2d_{cutoff} \quad (3.9)$$

will encounter a crystal satisfying the Bragg condition and will be diffracted out of the beam. A variety of materials have been considered for use as polycrystalline neutron filters [88, 89].

Polycrystalline bismuth, with a Bragg cutoff of 6.5 Å [88], would be an ideal wavelength filter for removing higher Bragg reflections from the monochromatic beam without reducing the 8.9 Å flux. Bismuth was not tested in these measurements, but should be considered in the future.

The transmission spectra through cold, 5 cm thick blocks of beryllium and graphite are shown in figure 3.22. Both blocks were cooled by immersion in liquid nitrogen for 5 minutes prior to measuring the transmission spectra, in order to minimize Debye scattering. The polycrystalline beryllium block shows the characteristic Bragg cutoff at 3.95 Å and demonstrates > 95% transmission at 8.9 Å. The graphite filter block¹²,

¹²G330S Purified Reactor Grade Graphite, MWI Inc., Rochester, NY,




Figure 3.22: The neutron spectra transmitted through 5 cm thick blocks of beryllium and reactor-grade graphite are shown. The beam spectrum without any filter blocks is shown for comparison. The spectra were measured using neutron time of flight spectrometry. The filter blocks were cooled by dipping them in liquid nitrogen for 5 minutes before measuring the transmission. The output width of the gate and delay is 250 ns.

on the other hand, has very poor long-wavelength transmission, in addition to showing the expected Bragg cutoff at 6.7 Å. Although carbon has very low absorption and incoherent scattering, the long-wavelength attenuation of this filter is believed to be due to small-angle scattering. Small angle scattering can occur at the index of refraction boundaries between powder grains and air. Since the graphite block is somewhat porous, such small angle scattering is not surprising.

3.8.2 Graphite Monochromators

The polycrystalline beryllium filter serves to sharply attenuate the $\lambda/3$ component of the monochromatic beam, but does nothing to reduce the $\lambda/2$ component. Since the polycrystalline graphite block had poor long wavelength transmission and no polycrystalline bismuth was available, an alternative approach was used to reduce the 4.45 Å contamination.

A $\lambda/2$ filter was constructed using ZYH grade HOPG¹³ as a 4.45 Å monochromator. Time of flight spectra with the monochromator and $\lambda/2$ filter are shown in figure 3.23. With the monochromator rotated off-Bragg ($\theta_1 = 45.6^\circ$), the high reflectivity can be seen by comparing the curves with the $\lambda/2$ filter on ($\theta_2 = 41.6^\circ$) and off ($\theta_2 = 90^\circ$) the Bragg condition. The total reflectivity for 4.45 Å neutrons only increases slightly when the primary monochromator is rotated onto the Bragg condition. Thus we can see that the HOPG monochromator makes a substantial reduction in the $\lambda/2$ wavelength component in the monochromatic beam.

3.9 Beam Characterization

The setup of the monochromatic beam, including the placement of the primary monochromator, $\lambda/2$ filter, and filter cryostat (containing the 5 cm of polycrystalline beryllium and 10 cm of single crystal bismuth, for γ and fast-neutron shielding), is shown in figure 4.2. With the monochromatic beamline set up in this fashion, a variety of measurements were made to characterize it. In this section each of these measurements will be presented and all of the measurements will be checked against each other for consistency. In section C.2.1 the consistency of the beam Monte Carlo

¹³The piece of HOPG used was 10 mm thick grade ZYH, with a quoted mosaic of 2.2 °. Grade ZYH HOPG is generally specified as having a mosaic of $(3.5 \pm 1.5)^\circ$.




Figure 3.23: Time of flight spectra of the 8.9 Å monochromator and $\lambda/2$ filter. The 8.9 Å monochromator is rotated to angle θ_1 and the $\lambda/2$ filter is rotated to angle θ_2 . A 5 cm block of 77 K beryllium filters the neutron beam prior to the monochromators. The output pulse width of the gate and delay is 700 ns.

simulations are compared with the measurements detailed herein.

3.9.1 TOF Measurements

With the monochromatic beamline assembled as described above, TOF spectra were measured with the detector in the transmission and reflection geometries and with the monochromator on and off the 8.9 Å Bragg condition. The measured TOF spectra are shown in figure 3.24. As can be seen, the number of neutrons “missing” from the dip in the transmitted spectrum due to reflection of 8.9 Å neutrons out of the beam is equal to those found in the reflected neutron peak. Hence, the measured reflectivity of 85% at the center of the beam is an accurate indication of the reflectivity




Figure 3.24: Time of flight spectra measured with the time of flight aperture in line with the center of the monochromator and the center of the neutron beam. The HOPG $\lambda/2$ filter is aligned to remove second order reflections from the 8.9 Å monochromator. Transmission through the $\lambda/2$ filter is shown with the 8.9 Å monochromator on and off the bragg condition. The TOF spectrum with the detector in the reflected beam is also shown. The output pulse width of the gate and delay is 250 ns.

(absorption is negligible). The $\lambda/2$ filter can also be seen at work in the reflected spectrum. Of the neutrons in the reflected spectrum, 40.1% are found in the primary reflected peak while 41.1% are found in the $\lambda/2$ peak and 17.5% are found in the $\lambda/3$ peak. The fraction of the cold neutron beam reflected into the monochromatic beam¹⁴ is 0.019 ± 0.001 .

¹⁴The total reflectivity of the monochromator is calculated as follows: the total number of counts reflected is divided by the ratio of the total number of counts transmitted through the monochromator to the monochromator's measured total transmission. The total transmission was determined above to be 0.81 without the beryllium filter in the beam. The total transmission above 3.95 Å (determined in the same method as the total transmission) is 0.85. Hence we take the total transmission with




Figure 3.25: Time of flight spectra taken at the end of running the trapping experiment (8/25/01) are shown.

In order to determine whether the monochromator reflectivity is degrading over time¹⁵, a time of flight spectrum was taken seven months after assembly of the monochromator (see figure 3.25). The measured reflectivity seven months after the monochromator's assembly is 96%. That this measured reflectivity is higher than that measured upon initial assembly (see figure 3.21) is likely due to the different geometry in this measurement. The TOF chopper was placed downstream of the monochromator, so those neutrons for which the reflectivity measurement is made are more highly collimated, which should increase the apparent reflectivity. Given the high measured reflectivity in this setup, we can conclude that no major degradation of the

the beryllium filter in place to be 0.83 ± 0.02 .

¹⁵Degradation of the monochromator over time could be caused either by radiation damage or water vapor introduced through a leak in the monochromator holder.

monochromator has occurred over seven months.

3.9.2 Flux Measurements

The neutron beam flux was measured in several positions using a calibrated fission chamber. The capture flux of the NG-6 white beam, with 5 cm of polycrystalline beryllium and 10 cm of single crystal bismuth in the filter cryostat (see above) was measured to be $(7.6 \pm 1.9) \times 10^8 \text{ n cm}^{-2} \text{ s}^{-1}$ at the exit of the NG-6 shielding, where the dominant error contribution is from uncertainty in the position and alignment of the fission chamber. The flux at the center of NG-6 at the same position was measured more carefully in the past to be $(9.0 \pm 0.5) \times 10^8 \text{ n cm}^{-2} \text{ s}^{-1}$ with 10 cm of bismuth and 10 cm of beryllium in the filter cryostat [42].

The capture flux of the monochromatic beam was measured as a function of position. The same fission chamber used to measure the flux in the white beam was used for this measurement. The fission chamber uses a deposit with an active region 1 cm in diameter. The fission chamber was scanned in two dimensions perpendicular to the monochromatic beam. The distance between the monochromator and the fission chamber was 89 cm. The measured flux as a function of position is shown in figure 3.26. From these measurements it appears that the apparatus is about 0.5 cm above the centroid of the monochromatic beam. This deviation is acceptable compared to the diameter of the neutron trap (8.6 cm), but could be made smaller through use of a more careful alignment technique (see section 4.1.1).

From the measured white beam flux of $7.6 \times 10^8 \text{ n cm}^{-2} \text{ s}^{-1}$ and the measured total monochromator reflectivity of 1.9%, one can calculate the expected flux of the monochromatic beam to be $1.4 \times 10^7 \text{ n cm}^{-2} \text{ s}^{-1}$. Correcting for the divergence of the neutron beam over the 2.21 m pathlength between the positions where the white




Figure 3.26: The monochromatic beam flux was measured using a fission chamber with a 1 cm diameter active area. The measured flux is shown as a function of the position of the center of the aperture, relative to the center of the 300 K end flange of the trapping apparatus.

and monochromatic flux are measured, the expected monochromatic flux decreases¹⁶ to $7 \times 10^6 \text{ n cm}^{-2} \text{ s}^{-1}$. The consistency of the expected and measured monochromatic fluxes indicate that the TOF data and the flux measurements are in rough agreement. Further consistency checks of the flux measurements, TOF measurement, and Monte Carlo simulations of beam production are found in section C.2.1.

3.10 Conclusions

An 8.9 Å monochromator has been constructed using stage 2 potassium-intercalated graphite. The monochromator, tiled from nine pieces, has a total size of 6 cm by 15 cm. The individual monochromator pieces have high stage purity and mosaics varying from 1° to 2°. The monochromator reflects more than 80% of the incident 8.9 Å neutrons, while reflecting less than 2% of the total cold neutron beam. The signal/neutron-induced-background in the magnetic trapping experiment is thus improved by a factor of 40 through use of the monochromator.

¹⁶The divergence can be roughly corrected by considering the white beam to have 6 cm diameter at the exit of the NG-6 shielding. Given the expected beam divergence of 1°, the 0.5 cm radius aperture of the fission chamber can accept neutrons from a 4.4 cm radius neutron beam 221 cm away. The effective reduction of the flux due to divergence is then the ratio of the actual beam to the potentially accepted beam: $(3/4.4)^2 = 0.5$

Chapter 4

Experimental Apparatus

The experimental apparatus used for magnetically trapping UCN is described in this chapter. The apparatus as it was in 1999, when the magnetic trapping of UCN was first demonstrated, is described in detail in [42]. The changes that have been made to the apparatus since then are described below. In section 4.1 the layout of the neutron beam, including the monochromator (see chapter 3), is presented. The changes which have been made to the cryogenic apparatus are described in section 4.2. These changes were largely necessitated by the installation of a much larger bore magnet, the design and testing of which is reported in section 4.3. The increased size of the magnet bore allowed a redesign of the detection system (see section 4.4) and the various shielding, neutron and otherwise (see section 4.5).

4.1 Neutron Beam

The neutron trapping experiment is performed at the NIST Center for Neutron Research (NCNR). The NCNR houses a 20 MW research reactor which produces the neutrons used in the experiment. Neutrons are released by the fission of enriched




Figure 4.1: The reactor, coldsource, and cold neutron guides at the NCNR [117].

^{235}U and moderated by a room temperature bath of heavy water (D_2O). Small beam ports for various instruments (primarily diffractometers) allow some of these thermal neutrons to escape the reactor where they are used for experiments. A larger hole in the reactor shielding (see figure 4.1) allows a cold neutron source to sit within the high thermal neutron flux of the reactor core.

The cold neutron source is a thin spherical shell of liquid hydrogen maintained at a temperature of approximately 20 K. Neutrons passing through the cold source to a series of neutron guides are partially thermalized, resulting in a neutron spectrum which is approximately Maxwellian with a characteristic temperature of 34 K (see section C.2.1). During late 2001 and early 2002, an upgrade to the NCNR is being performed in which the cold source shown in figure 4.1 is being replaced. The new cold source, which has an ellipsoidal rather than spherical moderator is expected to increase the 8.9 Å flux by a factor of 1.82 [118].

As shown in figure 4.1, some of the cold neutrons exiting the cold source enter




Figure 4.2: The layout of the end station at NG-6 is shown. The monochromator and apparatus are shown in their locations during 2001. By March of 2002 the monochromator will be permanently installed as shown in figure 7.2.

neutron guides (NG). These guides transport cold neutrons with high efficiency to the various experimental stations, tens of meters away. Cold neutron guides rely on total external reflection of grazing incidence neutrons from high-potential materials. A coating of ^{58}Ni is used for the cold neutron guides at the NCNR¹. The neutron trapping experiment is operated on neutron guide NG-6, the end station of which is designated for fundamental physics research. The end station and experimental layout as it was throughout calendar year 2001 is shown in figure 4.2. The neutron guides are housed within large shielding structures consisting of steel shells filled

¹Isotopically purified ^{58}Ni has the highest neutron potential (330 neV), and hence largest critical angle ($2.0 \text{ mrad}/\text{\AA}$) for total external reflection of any isotope. Larger critical angles can be obtained using multilayer materials (“supermirrors”) which rely on interference between alternating layers of materials with different neutron potentials.

with steel shot embedded in polyethylene wax. These shields are designed to absorb gammas, fast neutrons, and slow neutrons coming from the reactor core and guides. The guide NG-6 is 6 cm wide and 15 cm tall. Some neutrons from the bottom half of the guide are reflected by a pyrolytic graphite crystal to form a 4.96 Å monochromatic beamline, NG-6M, while the rest are dumped into a beamstop within the shield walls.

Neutrons from the upper half of the guide pass through a 6 cm diameter collimator, a shutter, and a filter cryostat before exiting the NG-6 shielding. The collimator is made of tungsten, which is very effective for gamma and fast neutron shielding, with a borated aluminum cover, which absorbs slow neutrons. The main shutter is a 5 cm thick steel plate with Lithium glass neutron absorber mounted on the reactor side, and stops gammas and neutrons from the reactor. In addition to the shutter, the cold neutron beam can be blocked by a computer-controlled beam blocker. The blocker is a paddle of lithium glass which can be driven in or out of the beam and allows the opening and closing of the neutron beam to be precisely timed (see section 5.2.1). The filter cryostat contains 5 cm of polycrystalline Bismuth (for gamma attenuation) and 5 cm of polycrystalline Beryllium (for attenuating neutrons of wavelength less than 3.95 Å, see section 3.8). The filter materials are cooled to 77 K to minimize the incoherent scattering of neutrons from phonons in the filter blocks.

Outside of the NG-6 shielding walls, neutrons travel more than a meter before reflecting from an 8.9 Å neutron monochromator (see Chapter 3). This monochromator reflects a narrow wavelength region near 8.9 Å through an angle of 61.2°. The placement of the monochromator was set so that the neutron trapping apparatus would not interfere with an existing structure (the “laser hut” in figure 4.2). Up beam from the 8.9 Å monochromator is a $\lambda/2$ filter (see section 3.8) – a pyrolytic graphite monochromator which reflects 4.45 Å neutrons out of the main beam and

dumps them into borated neutron absorbing material. Both monochromators are surrounded on all sides by lead shielding 10 cm thick. The monochromators are mounted on a welded aluminum table engineered to support the lead shielding. An exit hole in the lead shielding allows 8.9 Å neutrons diffracted by the monochromator to reach the UCN trapping apparatus.

A sliding table, which can be driven back and forth by a stepper motor, is mounted between the monochromator table and the trapping apparatus. This device, named the “slider,” has a stack of lead and a flight tube mounted on it. Either object can be placed along the 8.9 Å neutron flight path. When the flight tube is in place, 8.9 Å neutrons reflected from the monochromator may enter the apparatus. The flight tube is lined with borated shielding which absorbs, with minimal activation, any neutrons which scatter in the air. The stack of lead provides shielding for the trapping apparatus from gammas produced by neutron activation of the monochromator or its holder which exit the monochromator’s lead house through the hole left for the 8.9 Å neutrons. The slider, like the neutron beam blocker, is under computer control (see section 5.2.1).

4.1.1 Beam Alignment

The neutron trapping apparatus was aligned with the neutron beam using a combination of optical and neutron imaging techniques. First, images of the neutron beam were made using a two step imaging technique [119]. In this method, 20 cm square, 0.1 mm thick copper foils were placed in the monochromatic neutron beam (see chapter 3) and activated for 20 minutes (about 4 times the half-life of ^{66}Cu). Each activated foil was then used to expose an image plate for 10 minutes. The image plate is exposed by the electrons produced in copper (primarily ^{66}Cu) beta-




Figure 4.3: An image of the monochromatic neutron beam is shown. The image is generated in a two-step process [119]: activation of a copper foil in the neutron beam, followed by exposure of an image plate by betas from the decay of activated copper. The light wedge in the top right corner is a registration mark to preserve orientation of the foil. The image has been manipulated to heighten the contrast for display purposes. Hence, the image intensity is not necessarily linear in the neutron intensity in the image shown. The image shown maps an area 13.2 cm wide by 12.85 cm tall.

decay. The image plates were then scanned to produce two dimensional images of the neutron beam intensity (a typical image is shown in figure 4.3). Cadmium-bearing wires cross the foil and are imaged as an area of low neutron intensity. By shifting the foil holder so that the crossing point of the wire coincides with the centroid of the neutron beam, the beam center can be found. The beam center was determined to within about ± 5 mm at two foil locations 42 cm apart. The line of most intense flux is thus determined to $\pm 1.4^\circ$. A theodolite was used to transfer the beamline defined

by the two crosshairs to the neutron trapping cell.

A drawback of this technique is that it assumes that the maximum intensity of 8.9 Å neutrons coincides with the maximum intensity of the monochromatic beam. A more robust alignment technique, which should be used in the future, would involve measuring the time of flight spectrum as a function of position. By making a two dimensional map of 8.9 Å intensity with parallel foils at two distances from the monochromator, the beamline for 8.9 Å neutrons could be well defined.

4.2 Cryogenic Apparatus

4.2.1 Overview

The cryogenic apparatus used for magnetically trapping UCN is described in detail in [42] and is shown in figure 4.4. The principal components are a vertically mounted dilution refrigerator, capable of providing $> 400 \mu\text{W}$ of cooling power at 100 mK, and a magnetic trap. Both components are housed within an inverted-“T”-shaped cryostat. The cryostat has four thermal layers: a 300 K outer vacuum vessel, a 77 K radiation shield cooled by liquid nitrogen within a vertical jacket, a 4 K helium bath containing the superconducting magnets, and the superfluid-helium-filled cell which is coupled by a heatlink to the dilution refrigerator. Thermal isolation of these various stages is maintained by two vacuum spaces: the outer vacuum can (OVC) is a shared vacuum separating the 300 K, 77 K and 4 K regions, and the inner vacuum can (IVC) is a vacuum space separating the 4 K and < 250 mK regions. The contents of the original IVC and helium bath, including a 5 cm bore superconducting magnet, the neutron trapping cell, and a copper heat link, are shown in figure 4.5.

Hereafter, the changes made to the apparatus since 1999 will be described. The




Figure 4.4: The apparatus as it was in May of 1999 [42]. The components are described in the text.




Figure 4.5: The IVC as it was in May of 1999.

existing cryostat was modified to make room for a larger bore magnetic trap, and the accompanying detection and shielding systems. The IVC, as modified to accommodate the new magnet, cell, and heatlink, are shown in figure 4.6. The new magnet was designed to be the largest Ioffe-type magnet assembly that would fit inside of the existing cryostat. The IVC end cans, containing the cell ends and the coupling between the cell and the superfluid heatlink, were shortened to accommodate the increased




Figure 4.6: The helium bath and IVC after modification for the new magnet, heat link, and detection system.

length of the magnet. The horizontal section of IVC above the magnet was raised, and the “T” shaped mating section modified to allow insertion of a magnet assembly with such a large outer diameter. The basic techniques of component positioning (magnet support plates), stress relief (welded metal bellows coupling the different IVC components), and vacuum sealing (metal-metal indium seals) remained unchanged, though many of the experimental components implementing these techniques had to be remade.

Inside of the bore of the new, larger, magnet is inserted a larger neutron trapping and detection cell than in the previous apparatus. The cell is coupled to the dilution refrigerator using a superfluid heatlink (see section 4.2.3). The superfluid heatlink provides much higher thermal conductance and lower eddy current heating than the copper heatlink used previously. Changes to the helium fill lines through which the cell and heatlink are filled, and their thermal coupling to the dilution refrigerator are described in section 4.2.2.

4.2.2 Helium Fill Lines

When the apparatus was operated for the demonstration of neutron trapping (1998-1999) the lowest operating temperatures of the mixing chamber and cell were 190 mK and 250 mK, respectively [42]. The high mixing chamber temperature indicated the presence of a large heat load on either the cell or the mixing chamber. The cell to mixing chamber temperature difference was small compared to the temperature drop which such a large heat load would produce across the heat link (whose thermal conductance was measured in a separate experiment). This suggested that a large heat load was present on the mixing chamber itself. The ultrapure helium fill lines were suspected to be the source of that heat load. A series of test measurements (detailed below) verified that the large heat load on the mixing chamber was due to refluxing superfluid film flow in one of the ultrapure helium fill lines. Replacement of this fill line eliminated the excessive heat load and lowered the base temperature of the mixing chamber and buffer cell.

The source of the high heat load on the mixing chamber was determined through a series of tests. The dilution refrigerator, with the buffer cell and its fill lines extending to room temperature was decoupled from the horizontal section of the cell. The

refrigerator was then operated in a test dewar for quick turnaround time. With the original buffer cell and fill lines the base temperature of the buffer cell and mixing chamber were measured to be 50 and 60 mK respectively with no helium in the buffer cell and 250 and 425 mK with helium present. This measurement indicated that the presence of helium in the cell was necessary for the large helium heat load to be present.

The ultrapure helium which fills the cell is stored at room temperature as a gas in large stainless steel vessels (“dumps”). Four dumps, each with a volume of 450 L, store 3 atm of helium gas. The ultrapure helium gas handling system is all metal and had never been exposed to natural abundance helium. A sealed diaphragm compressor² is used to pump the ultrapure helium out of the cell and back into the dumps after a run. It can also be used to increase the condensation rate when the dumps are nearly empty. The ultrapure helium is passed through an all-metal liquid helium cold trap before entering the main cryostat, in order to remove non-helium impurities³.

The dilution refrigerator is equipped with three fill lines which connect the room temperature ultrapure helium gas handling system to the buffer cell. Two of these lines (called the “fill lines”) are 1.6 mm outer diameter stainless steel tubes which wrap around the still pumping line inside of the main helium bath, penetrate the 4K IVC top, pass through the 1K pot, are soldered to the outside of the still, and pass through the continuous heat exchangers before connecting to the buffer cell. These fill lines

²The diaphragm compressor was manufactured by KNF Neuberger, Trenton, NJ, (<http://www.knf.com>). In order to make the compressor helium leak-tight, it was necessary to remake the diaphragm seats from machined brass rather than cast aluminum. All NPT fittings were also replaced with VCR fittings.

³The “ultrapure” helium is extremely high purity in terms of its ^3He content, but contains a significant amount of nitrogen gas when it is first delivered. After passing through the helium coldtrap once, these impurities are entirely removed.




Figure 4.7: The design of the original (A) and modified (B) bypass lines are shown. The bypass line is a low impedance connection between the buffer cell and the room temperature ultrapure helium gas handling system.

are very well heat sunk but have high impedance. As such they are used for filling the cell but cannot be effectively used for pumping on it. The cell is pumped out through a “bypass” line (see figure 4.7). The bypass line has much lower impedance but is not heat sunk at as many points. The original bypass line (figure 4.7A) was only heat sunk at the 1K pot and the buffer cell. It was hypothesized that the cause of the heat load on the buffer cell was superfluid helium film flow through the ultrapure helium bypass line.

A large heat load on the buffer cell can be produced by superfluid film flow through the original bypass line. A well known property of superfluid ^4He is its tendency to produce a thin film along the wall of the vessel containing it. Hence superfluid film could flow up the walls of the original bypass line until the wall temperature was hot enough to boil the film off ($\sim 300\text{mK}$). The hot gas atoms then bounce around in the

line until they hit a colder point and stick (the refluxing helium tends to recondense at a point in the apparatus around 200-300 mK [120]). Since there was no part of the bypass line heatsunk at this recondensation temperature, the refluxing tended to heat the bottom end.

Heat loads on the mixing chamber due to the refluxing superfluid helium film were alleviated by heatsinking the bypass line to the 50 mK cold plate and to the still, and introducing a pinhole aperture. As shown in figure 4.7B, the fill line was modified by introducing a pinhole aperture heatsunk at about the temperature of the cold plate (between the step and continuous heat exchangers). This pinhole aperture, an 0.025" hole in 0.005" thick copper foil soldered to a standard 1/8" VCR gasket⁴, has a minimal effect on the impedance of the bypass line in the laminar flow regime (such as when pumping out the cell). Since the sticking probability for helium is close to 1 for the temperature of this gasket it serves as a collection point for all refluxing helium. The heat load of the refluxing is then taken up at the cold plate, where the cooling power is higher, rather than on the buffer cell and mixing chamber. When the mixing chamber and buffer cell were cooled with the modified bypass line they reached base temperatures of 50 and 60 mK respectively both with and without helium in the buffer cell.

4.2.3 Buffer Cell and Heat Link

The fill lines and heat link between the buffer cell and the horizontal cell (trapping region) were replaced when the apparatus was upgraded. In the old apparatus the heat link between the mixing chamber and the cell was provided by copper rods and

⁴VCR fittings are all-metal face-sealing vacuum fittings manufactured by the Swagelok Company, (<http://www.swagelok.com>).




Figure 4.8: A superfluid-helium filled cupronickel heat link connects the buffer cell with the cell.

braids, while the cell was filled with ultrapure helium through two 3.2 mm stainless steel lines[42]. This design was susceptible to eddy-current heating of the large copper mass when the magnet was ramped up and down. Furthermore, the heat link had poor thermal conductivity, limited by the braid. Since the new apparatus was designed for higher magnetic fields and with a smaller separation between the magnet coils and the heat link, estimates of the large size of the eddy current heating necessitated a new heatlink design.

Cell filling and thermal conductivity functions are now provided by a single large-diameter helium-filled cupronickel heat link (see figure 4.8). A continuous volume of helium extends from the buffer cell to the horizontal part of the cell. The heat link is made of large inner-diameter tube (varying from 1.2 to 2.2 cm). The limiting impedance of the system is the 1.3 cm diameter by 4.9 cm long tube penetrating

into the cell end caps. All fixed parts of the heat link and cell (including the cell end caps) below the buffer cell are made of 70-30 cupronickel⁵. The formed bellows providing stress relief in the heat link were made from thin-walled beryllium copper. Throughout the heat link, the high resistivity and small area perpendicular to the magnetic trap's field lines results in negligible eddy current heating. When the current was ramped at a rate of 2 A/s the cell warmed from 137 mK to 183 mK.

The use of superfluid helium rather than copper results in a much greater thermal conductance for this heat link compared with the previous one. The previous heat link had a measured conductance of $2.3 \times 10^{-4} \text{ W K}^{-1}$ [42]. By varying the power applied to a heater at one end of the cell while observing the mixing chamber and cell temperatures, the thermal conductance was determined. This measurement implies a thermal conductance of $0.81 \pm 0.04 T^3 \text{ [W K}^{-1}\text{]}$, where T is the temperature of the mixing chamber in Kelvin [65]. This agrees to within 20% with the conductance calculated from the known thermal conductivity of superfluid helium and the geometry of the cell. The intrinsic heat load on the cell (in the absence of windows) was also determined to be $360 \mu\text{W}$.

The buffer cell provides the thermal link between the superfluid helium heat link and the mixing chamber of the dilution refrigerator. As shown in figure 4.9, the buffer cell contains copper fins coated with silver sinter. This sintered coating has a large surface area in order to maintain high thermal conductivity between the mixing chamber and superfluid, even at temperatures $< 100 \text{ mK}$, where the Kapitza boundary resistance due to phonon mismatch is significant [55].

⁵70-30 cupronickel is a copper alloy with 70% copper and 30% nickel. It has a high electrical resistivity, thermal contraction and thermal conductivity similar to stainless steel. Unlike stainless steel, however, cupronickel does not have a martensitic phase, and does not become ferromagnetic, even in large magnetic fields and after many thermal cycles.




Figure 4.9: The new buffer cell is shown. The copper fins coated with silver sinter maintain thermal contact between the buffer cell and superfluid helium heat link.

The buffer cell also contains a helium level sensor. This sensor is a capacitor formed from two concentric stainless steel tubes. As the space between the two tubes fills with liquid helium, the capacitance increases due to the dielectric constant of helium. Condensation of ultrapure helium into the cell is controlled using the level meter so that when the cell is cold (the density of helium does not change significantly below 1.4 K) the level of helium within the buffer cell is approximately halfway up the fins. At temperatures below 4 K and above about 800 mK it was observed that good thermal contact is maintained between the cell ends and the mixing chamber as long as there is at least some liquid in the cell. This good thermal conductivity

is attributed to convection in the helium gas above the liquid due to the relatively high saturated vapor pressures at these temperatures. Below about 800 mK it is necessary to have the helium level within the buffer cell to maintain thermal contact between the cell and the mixing chamber. Due to this effect, even in the absence of the functioning level meter it is possible to determine when the cell is full from the temperatures of the cell ends and the mixing chamber.

The heat link requires rotatable indium seals and flexible bellows so that the orientation is not over-constrained during assembly and to minimize mechanical stresses during cooling (see figure 4.8). A single beryllium copper bellows assembly in the vertical section decouples the dilution refrigerator from the heatlink. The vertical section of the heatlink hangs straight down from the mixing chamber. The horizontal section of the heatlink is decoupled from the cell by two smaller beryllium copper bellows. The bellows are silver soldered to the cupronickel tubes and flanges of the heat link. The cell endcaps are soft-soldered to the cell tube (one of these solder joints must be made with the cell tube in place inside the magnet bore). Indium seals between the various flanges allow assembly of the heat link inside of the IVC.

The horizontal cell is capped off with different windows at the two ends. On the beam entrance side, the window is made from Teflon PFA or Teflon FEP (see figure 4.10). Teflon FEP and PFA are perfluorinated Teflons with low neutron activation and scatter. A superfluid-tight seal is formed by compressing the Teflon between the cell endcap and a compression ring [82]. The cell endcap and compression ring are made of 70-30 cupronickel, as is the rest of the cell. Compression is applied by 22 bolts on a 24 bolt pattern⁶. The bolts used for compression are 316 stainless

⁶One bolt must be left out to allow the tube connecting the heatlink to the cell to penetrate the endcap. The second bolt is left out to preserve symmetry and maintain even compression




Figure 4.10: A teflon window seals the beam entrance side of the cell.

steel, a nonmagnetic alloy chosen due to its similar thermal contraction to that of cupronickel. The other end of the cell has an acrylic window for optical access. The window is attached to a “snout” assembly as shown in figure 4.11 [121]. An acrylic window is glued to an acrylic tube using Stycast 1266 epoxy⁷. The acrylic tube is glued around the outside of a thin cupronickel fin with a wall thickness of 0.25 mm. The outer diameter of the cupronickel tube is machined to be 0.1 mm smaller than the measured inner diameter of the acrylic tube. The total length of the cupronickel fin is 12.5 mm, of which 6 mm overlaps with the acrylic. The snout assembly is sealed to the cell endcap using a standard indium seal.

After assembly, 6.35 mm thick cupronickel plates are inserted which rigidly connect

⁷Stycast 1266 is a clear epoxy, widely used at low temperatures, whose thermal contraction is well matched to that of acrylic. It is manufactured by Emerson & Cumming, (<http://www.emersoncumming.com>).




Figure 4.11: The “snout” assembly.

to the cell endcap and support the upper horizontal heatlink. The upper horizontal heatlink is held down against these support plates by kevlar strings, which loop once around the heat-link and also constrain the side-side motion of the cell (see figure 4.12). An additional kevlar string connects the horizontal heatlink support plate to the back plate of IVC end can. This prevents motion of the cell along the beam direction (z direction). The length of these z -positioning strings is set so that a gap of $2 - 4$ mm is maintained between the snout and the 4 K optical window (a similar gap results on the beam entrance side of the apparatus). Finally, one kevlar string at each end supports the cell from below. The strings that support the cell are the same design as was used in the previous version of the apparatus [42] and allow the vertical position of the cell to be set with respect to the magnet bore.




Figure 4.12: End on view of the IVC end can shows the kevlar strings for supporting and positioning the cell. The upper set of strings loops around the heatlink once. The lower string loops once around the cell. The z-positioning string makes a direct connection between the IVC end can and the horizontal heatlink support plate.

Because kevlar creeps up to 1% under load, and expands by a fractional amount of order 10^{-3} when cooled from 300 K to 4 K [122, 123], the cell was always positioned off-center at room temperature. There is a 3 mm difference between the outer radius of the cell and the inner radius of the magnet bore. The cell was positioned so that the top of the cupronickel cell tube was approximately 1 mm below the top of the magnet bore, when warm. Braided kevlar cords are used for support due to their high tensile strength and low thermal conductivity. The thermal conductivity of kevlar has been measured to be $2 \times 10^{-4} \text{ W cm}^{-1} \text{ K}^{-1}$ at 4 K [122]. Based upon this thermal

conductivity, the total heatload on the cell from all of the kevlar positioning strings is estimated to be $15\ \mu\text{W}$.

4.2.4 Blackbody Heat Loads

In order to match the increased magnet bore, larger neutron entrance windows, optical access windows, and lightguides were used, compared to our original apparatus. As a result of increasing the size of these objects, the blackbody heat loads on the low temperature portions of the apparatus increased dramatically. In this section, the origin, consequences, and elimination of these heatloads will be discussed.

The layout of the neutron entrance windows is shown in figure 2.1. Other than the increase in diameter, the design of the entrance windows is unchanged from that used previously [42]. Neutrons entering the apparatus pass through three 10 cm diameter vacuum windows, each of which is made of 0.25 mm thick Teflon PFA [82]. One of these windows, for the cell, is described above. The others are identical. Since PFA Teflon does not block 300 K or 77 K blackbody radiation, it is necessary to use an additional foil window at 77 K and 4 K for this purpose. As before, these windows are made from thin beryllium foils. The foils used are 0.05 mm thick, 10 cm diameter, light-tight beryllium windows⁸. The foils are thermally anchored at their edges to the cryostat's end flanges with aluminum compression rings. Using the known thermal conductivity of beryllium, the maximum temperature rise at the center of the foils, due to blackbody radiation, is calculated to be less than 1 K for each of the two foils. Therefore, no increased heatload is expected due to the neutron entrance side of the apparatus.

⁸Grade IF-1 Beryllium Metal X-Ray Windows, Brush Wellman Engineered Materials, Fremont, CA. (<http://www.beminc.com/web/bpdweb.nsf/>)

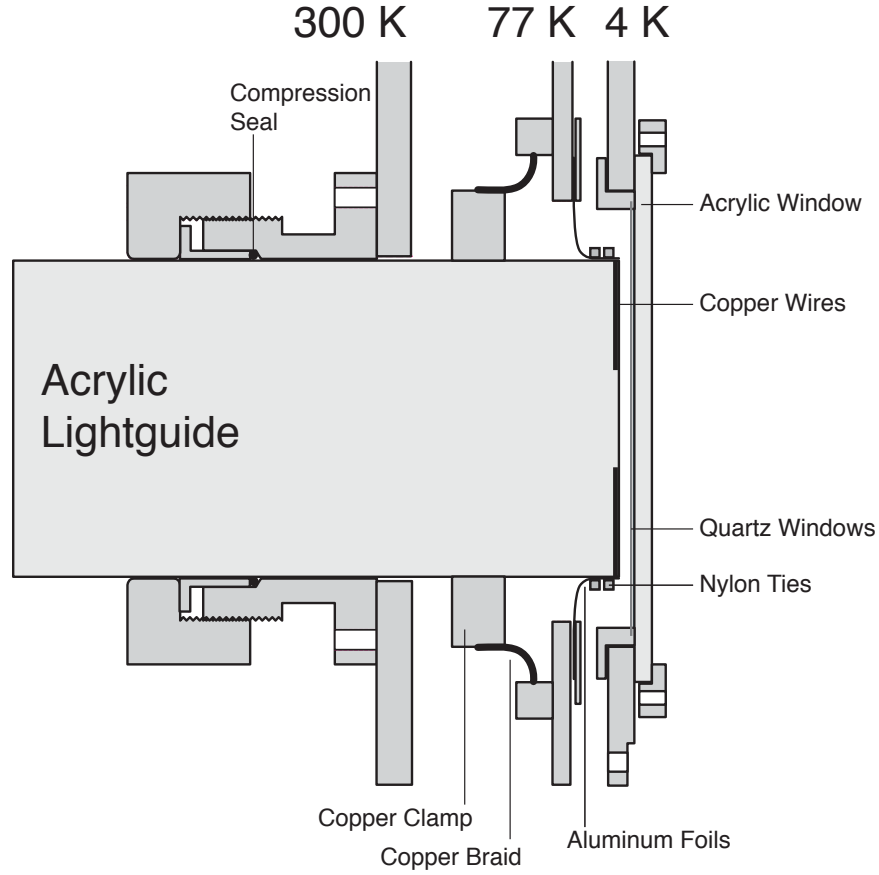


Figure 4.13: The thermal anchoring of the detection system is shown.

As is the case for the neutron entrance side, the size of the windows (and light-guide) used on the detection side were increased compared with our previous work. The final setup used on the detection side of the apparatus is shown in figure 4.13. The large blackbody problems are due to an unexpectedly high temperature of the “77 K” nitrogen shield. The primary heat input onto the nitrogen shield comes from the 11.4 cm diameter acrylic lightguide. Despite the relatively poor thermal conductivity of acrylic ($1.5 - 2 \times 10^{-3} \text{ W cm}^{-1} \text{ K}^{-1}$ for 77–300 K [124]), the short distance between the warm and cold ends of the lightguide results in a large heatload (3 W) on the nitrogen shield. The design of the nitrogen shield itself also contributes to its

high temperature. As shown in figure 4.4, the nitrogen jacket is only filled with liquid in the vertical section of the dewar. Below the end of the jacket, the nitrogen shield is made of 3.2 mm thick aluminum. In the “T” section of the cryostat, where the vertical and horizontal sections are mated, there are also two joints where the thermal contact is made simply by bolting the different parts of the aluminum shield together. This arrangement has a relatively high thermal resistance which, when combined with the high heat load from the lightguide, results in the “77 K” end flange on the detection side of the apparatus being much hotter than 77 K. The heat load on the colder parts of the apparatus is very sensitive to temperature changes of the 77 K flange since the power output in blackbody radiation goes as T^4 . High blackbody heat loads from 77 K on the 4 K components result in high helium boiloff and in the heating of the 4 K acrylic window to a temperature well above 4 K. The increased blackbody from this window can then put a large heat load on the cell that it cannot be cooled to a low enough temperature to trap neutrons.

Two clamps were used to thermally anchor the acrylic lightguide to the nitrogen end flange. The outer clamp is made from a split ring of OFE copper connected to the 77 K end flange by a copper braid. The ring is compressed around the acrylic using brass screws. Acrylic contracts more than brass or copper when it is cooled, so it is not certain that this clamp maintains good thermal contact at low temperature. The inner clamp is similar to one used on the previous apparatus [42]. Thermal contact is provided by a total of 336 strips of 5 cm wide and 30 μm thick aluminum foil⁹. The aluminum strips were secured to the acrylic lightguide using nylon cable ties. Since nylon contracts more than acrylic when cooled to 77 K, the force on the aluminum strips should increase as the lightguide is cooled and good thermal contact should be

⁹Reynolds Wrap Heavy Duty Aluminum Foil. Consumer Products Division, Reynolds Metals Company, Richmond, VA, (<http://www.reynoldskitchen.com>).




Figure 4.14: The scheme used for directly cooling the horizontal nitrogen shield is shown.

maintained.

The apparatus was cooled with both of these clamps installed and the nitrogen jacket full, but with both the “4 K” and 300 K portions of the apparatus at room temperature. In this configuration, the temperature of the “77 K” end flange near the clamping rings was measured to be 137 ± 2 K while the temperature at the center of the lightguide was 184 ± 4 K. At a temperature of 137 K, the blackbody radiation due to the end flange was ten times that which would be generated by the same flange at 77 K. In order to reduce this enormous heatload on the 4 K, it was necessary to extend the nitrogen reservoir so that liquid nitrogen was in much better thermal contact with the end flange. The setup used for directly cooling the horizontal 77 K shield is shown in figure 4.14. A 1.3 cm hole was drilled in the bottom of the existing nitrogen jacket. An indium seal is made at that location to a brass adapter fitting which is drilled through with an inner diameter of 1.0 cm and which mates externally to a preformed metal bellows piece. The metal bellows, which is necessary

for assembly of the apparatus, makes an indium seal to a fixed set of 2.5 cm copper refrigerator tubing which forms the horizontal nitrogen reservoir. The two halves of a split metal pipe are clamped around the bellows after assembly, to ensure that the bellows remain straight. From a “T” in the center of the horizontal section, where the joint to the vertical section is made, the copper tubing runs along the top of the horizontal nitrogen shield until it reaches the two end flanges. The tubing forms a horseshoe (Ω -shaped when viewed end-on) around each end of the apparatus about 1 cm behind the end flange itself. Copper fingers (17 on each side) 5 cm wide, 5 cm long and 1 mm thick are soldered on one side to the copper horseshoe and are clamped on the other between the end flange and the horizontal nitrogen shield. The horseshoe is attached to the horizontal nitrogen shield using an attachment strap (made from a split hose-clamp) which connects its two bottom ends.

In order to prevent vapor lock (calculations and observations both indicate that this occurred), it is necessary to put a second tube inside of the horizontal reservoir, through which vaporized gas can escape. The inner tube is 8 mm inner diameter, 1 mm wall thickness, FEP Teflon tubing. The tube extends from the top of the detector side horseshoe to the top of the T-section bellows. The top end of the tube is coupled to a 10 cm long stainless steel adapter with inner diameter 6.4 mm and wall thickness 0.9 mm. This adapter piece penetrates the brass adapter piece and the hole in the nitrogen shield. Stiff wires attached to the stainless steel adapter maintain a gap between it and the bottom of the nitrogen jacket. Inside of the nitrogen shield, the stainless steel inner-tube adapter is welded to a 1.0 cm inner diameter, 1.3 cm outer diameter tube which extends to the top of the nitrogen jacket. The exhaust pickup from the nitrogen shield is placed below the top of the inner tube exhaust line, so that liquid nitrogen will never cover the inner tube. Hence, liquid nitrogen fills the

lower reservoir through the gap between the inner and outer tubes and gas produced by the boiling of the liquid nitrogen flows out the inner tube with low impedance.

Following the implementation of this direct cooling system, the temperature of the nitrogen end flange was measured to be 82 K. The effect on the helium boiloff of cooling the end flange in this manner is discussed in appendix B. Because of the poor thermal conductivity of acrylic, the center of the lightguide may remain much hotter than the edge. A calculation using the finite element analysis package ANSYS¹⁰ indicates that if the portion of the lightguide clamped by the aluminum/nylon clamp is fixed at 90 K and the room temperature part of the lightguide is fixed at 300 K, then the center of the cold end of the lightguide will be at 195 ± 10 K. In order to reduce the temperature of the center of the lightguide further, 1 mm copper wires¹¹ were embedded in the cold face of the lightguide. Grooves of depth 1.25 mm were milled in the end of the lightguide with a pattern shown in figure 4.15. The pattern was chosen so that the distance between any point on the surface and a wire is less than 1 cm and covers 11% of the lightguide face. The copper wires are embedded in Stycast 1266 epoxy. The wires extend past the outer diameter of the lightguide and are bent around and compressed under the aluminum foils by the aluminum/nylon clamp. With the end flange at 82 K and the wires in place, the temperature at the center of the lightguide was measured to be 130 K.

Before implementing the direct cooling of the horizontal nitrogen shield, and embedding wires in the 77 K lightguide, it was impossible to cool the cell below ~ 800 mK. Again, the poor thermal conductivity of acrylic is the culprit. The radiative heatload on the 4 K acrylic window (which has a 15.2 cm diameter open area)

¹⁰ANSYS 5.7, ANSYS, Inc., Canonsburg, PA, (<http://www.ansys.com>).

¹¹Puratronic grade copper. Sigma-Aldrich, (<http://www.sigmaaldrich.com>).




Figure 4.15: The pattern of wires embedded in the cold face of the 77 K lightguide is shown. The lightguide has overall diameter 11.4 cm. The grooves are 1.25 mm deep.

is sufficient to heat the center of the acrylic window to substantially above 4 K. The blackbody radiation from this hot window then puts enough of a heatload on the cell that it cannot be cooled to its base temperature. Even after the lightguide is cooled to 130 K, the heat load on the 6.4 mm thick acrylic 4 K window from this radiation is 0.17 W. Given the acrylic's low thermal conductivity of $5 \times 10^{-4} \text{ W cm}^{-1} \text{ K}^{-1}$ [55], the center of the window may reach a temperature of $> 100 \text{ K}$ if the edge is pinned¹²

¹²In order to ensure that this pinning occurred, it was necessary to supplement the thermal carrying capacity of the 1.0 cm thick stainless steel 4 K end flange with a 1 mm thick copper plate. The two plates are bolted together at the edges and held together over their entire area with vacuum grease. The temperature at the outer




Figure 4.16: The 4K window assembly is shown, viewed from the inside of the IVC, looking out. The outer flange, with bolt circle, compresses the indium seal between the acrylic window and the 4 K flange. Part of the quartz window compression flange can be viewed inside of the opening in the acrylic window compression flange. The cloverleaf shape of the four crystalline quartz window pieces is shown.

at 4 K.

A search was made for a material which would absorb the 130 K blackbody, and which had higher thermal conductivity than acrylic. An excellent conductor at low temperature, which can absorb about 95% of the 130 K blackbody power spectrum is crystalline quartz. However, even 5% of 0.17 W is much more power than the dilution refrigerator can remove at a temperature of 100 mK. Hence, crystalline quartz cannot be used alone, but instead is used together with acrylic to make a composite window. The geometry of the composite window is shown in figure 4.16. On the outside of the acrylic window, four pieces of crystalline quartz are glued¹³. Each piece is cut from

edge of the window assembly was then measured to be 5.5 K.

¹³Crystalline quartz wafers were provided as complimentary samples by Sawyer

a 10 cm diameter wafer with thickness 0.33 mm. In order to thermally anchor the quartz to the acrylic, and to prevent loss of light in the detection system due to surface reflections at the interface, the quartz windows are glued to the acrylic using Stycast 1266 epoxy. In thermal dunk tests of individual pieces, the quartz windows did not delaminate, nor did they crack. Upon cooldown of the entire assembly, however, the quartz windows did crack and partially delaminate. They provided sufficient cooling, however, that the refrigerator was able to cool the cell below 250 mK for extended periods of time. The detection efficiency likely suffered due to these cracks. In the future, we plan to replace the thin quartz crystals with a single 1.5 mm thick, 15.3 cm diameter crystalline quartz window. This window will be thermally anchored at the edges only.

4.3 Magnetic Trap

4.3.1 Magnet Design

The magnetic trap used in this work is in the Ioffe configuration (see section 1.5.1). The implementation of this conceptual design with real magnets requires certain compromises. Since infinite current bars do not exist, they are approximated by racetrack-shaped coils (see figure 4.17). Four of these racetrack coils are arranged to produce the radially-confining magnetic quadrupole. Axial confinement is provided by two solenoid assemblies with identical current senses. Within each solenoid assembly a large solenoid (“pinch coil”) provides axial confinement and two smaller solenoids with opposite current sense (“bucking coils”) cause the magnetic field to fall off faster. The bucking coils both increase the length of the magnetic trap, and reduce the magnetic




Figure 4.17: The configuration of superconducting coils in our Ioffe trap is shown. Four racetrack-shaped coils form a radial quadrupole. Two sets of solenoids with the same current sense pinch off the magnetic trap along the axis.

field in the turnaround region of the racetrack coils.

The trap depth is defined as the difference between the maximum and minimum values of the magnetic field within the confinement region. In our case the maximum magnetic field value is found at the radial edge of the trap, which is defined by the location of the physical wall (TPB-coated GoreTex) at $r = 4.2$ cm as shown in figure 4.18. The actual trap depth is set at the radial edge of the trap not for $z = 0$, as shown in the figure, but near the end of the trap, at the boundary between the bucking and pinch solenoids. Here the fringing fields of the solenoids counter the quadrupole field and reduce the trap depth. The minimum (but non-zero) magnetic field is located along the center axis of the trap. The trap depth is 1.1 T for an operating current of 170 A.

All of the coils in our trap are wound with superconducting wire ¹⁴. The process of designing the magnetic trap depends on trading off size and depth of the magnetic trap by altering the shape and placement of the various coils, and is described in detail in [65]. For the geometry of coils used in our magnetic trap (see figure 4.17)

¹⁴The superconducting wire used is composed of NbTi superconductor embedded within a copper matrix. The copper:superconductor ratio is 2:1. The copper is necessary to carry current in the case of a quench. The wire was purchased from Supercon, Shrewsbury, MA, (<http://www.anomet-supercon.thomasregister.com/olc/anomet-supercon/home.htm>).

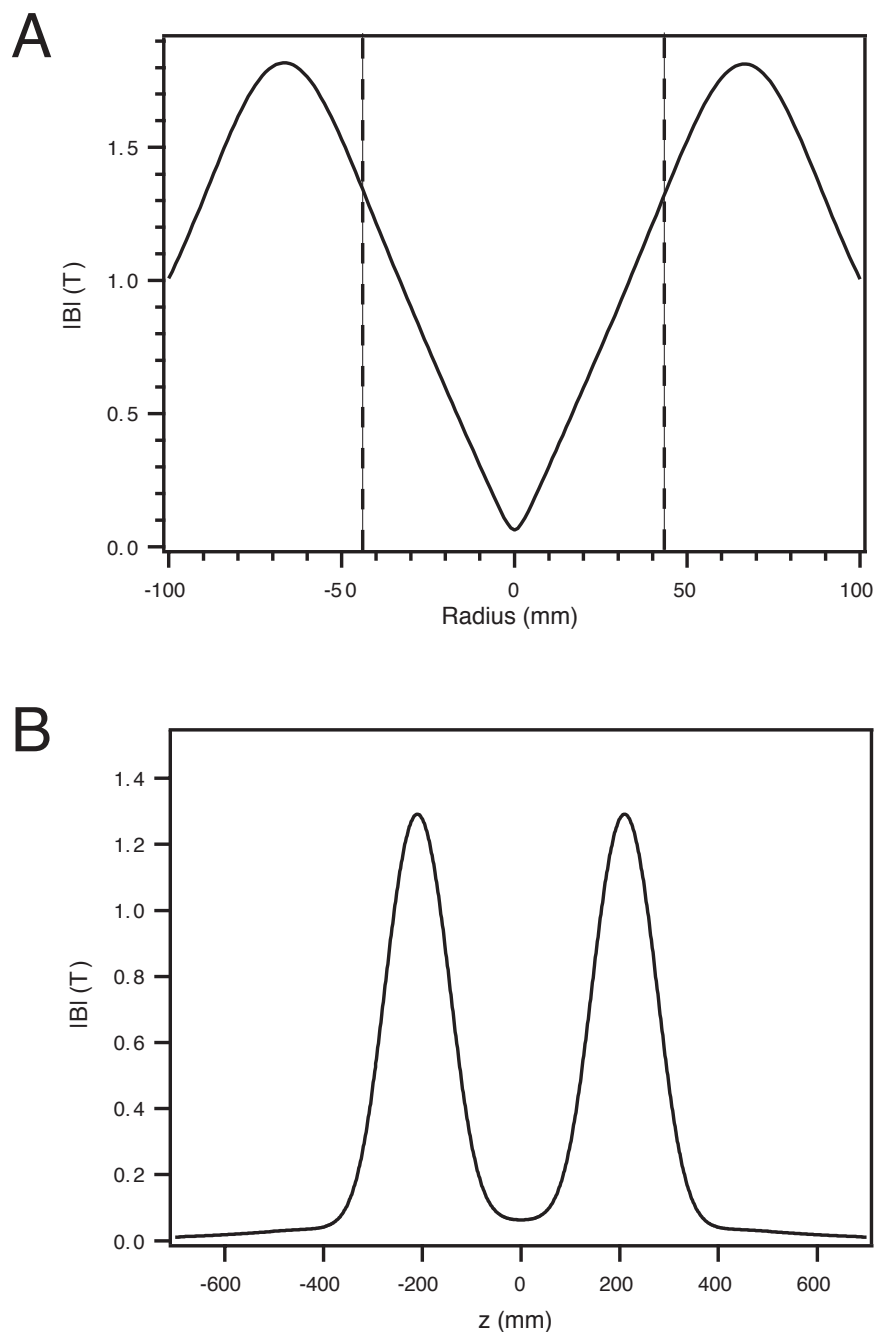


Figure 4.18: The magnetic field is shown as a function of position within our trap. A. The field strength as a function of radius at $z = 0$ (the center of the trapping region). B. The field strength is shown as a function of z for $r=0$ (along the axis of the trap). The field profile shown is for an operating current of 170 A. The physical wall, defining the trap depth, is located at $r = 4.2$ cm.

and the measured load line of the wire used for winding them, the maximum possible operating current is 320 A.

In an ideal magnet, it would be possible to raise the current of the magnet to its design value with no quenches. Usually at least a few quenches are required, and the process of repeatedly quenching the magnet at higher currents is called “training” the magnet. A magnet which undergoes a large number of quenches without reaching the design current is said to suffer from “excessive training.” As will be described below (see section 4.3.3), our magnetic trap never approached its design performance. The magnet assembly (and some of its components, when tested separately) quenched numerous times at lower currents.

A magnet quench occurs when some event heats part of the superconducting wire above the critical temperature T_C and it becomes non-superconducting. Dissipation of the magnet’s stored energy through resistive heating of the normal portion of the magnet produces additional heat which is conducted along the magnet wire, driving more of the magnet normal and propagating the quench. Within several milliseconds it is possible for the entire stored energy of the magnet to be dissipated¹⁵. If the stored energy is high and the quench propagation time is slow compared with the energy dissipation time, a small length of wire may heat up enough to vaporize, destroying the continuity of the magnet. The probability of such a catastrophic event depends on the location within the magnet where the quench occurs and details of the quench dynamics (quench propagation speed and thermal conductivities of the materials involved). Hence, excessive training of the magnet is undesirable for several

¹⁵In normal operation, our magnet is supplied with a constant current by a power supply. When a quench occurs, the resistance of the magnet becomes too great for the power supply to provide the programmed current (magnet power supplies typically are low voltage, high current devices), and the power supply automatically trips.

reasons: risk of damage to the magnet, consumption of cryogenic liquids¹⁶, and time wasted.

In magnet construction, one aims to minimize the frequency of magnet quenches. Design and construction of our magnet included several specific elements to this end. The superconducting wires in all of our coils are embedded in epoxy. Epoxy prevents the wires from moving in response to the forces they experience as the magnet current is ramped up. Small wire movements can generate sufficient heat to initiate a quench. Similarly, epoxy-rich regions have a tendency to crack under stress, releasing energy and inducing quenches. The coils were wound to minimize epoxy-rich regions by close packing of the wires and insertion of fiberglass between wire layers. Furthermore, the quadrupole assembly was radially prestressed with sufficient force to prevent coil movement during energizing of the magnet. The prestress was applied by aluminum collars surrounding the titanium magnet form assembly (see figure 4.19). The collars were machined to have only a $25\text{ }\mu\text{m}$ clearance from the outer diameter of the assembled quadrupole. A series of seven collars fit over the entire length of the quadrupole assembly. Differential thermal contraction of aluminum and titanium produced a prestress along the length of the magnet [65].

4.3.2 Quench Protection Circuit Design

Despite careful design and construction of a superconducting magnet, there always remains a finite probability of a quench. Since large quenches can be catastrophic, a prudent measure is implementation of a system for protecting the magnet from any

¹⁶In any quench, the stored energy in the magnet is eventually deposited in the helium bath surrounding the magnet, resulting in the rapid evaporation of liquid helium. The latent heat of evaporation of liquid helium is 2.56 kJ l^{-1} [55]. Quenches of our magnet deposited up to 50 kJ of energy into the helium bath.




Figure 4.19: The magnetic quadrupole assembly was prestressed by the differential contraction between a series of tight-fitting aluminum collars and the titanium magnet form.

quenches which do occur. Such a quench protection circuit typically relies on the rapid detection of quenches and the dissipation of part of the magnet's stored energy in an external dump resistor.

Quench protection then consists of two problems: quench detection and energy dissipation. A schematic of the quench detection circuit used for our magnet is shown in figure 4.20. The magnet assembly consists of 10 coils connected in series. The connections between each set of three coils in the two solenoid assemblies are made without the use of lugs or voltage taps. For the purposes of quench detection we thus consider each solenoid assembly to be one coil. The quench is detected by monitoring the voltages across each of the six coils. This voltage consists of both inductive and resistive components. When the magnet is being ramped, a non-zero inductive voltage results on each coil due to its self inductance and the mutual inductance between the various coils. A resistive voltage is only present when the magnet goes normal in




Figure 4.20: The quench detection circuit used for the magnetic trap assembly.

a quench. In order to distinguish between resistive and inductive voltages we rely on the symmetry of the magnet assembly¹⁷. The magnetic fluxes penetrating any of the quadrupole coils should be almost identical to that penetrating either of its neighbors. Similarly, the fluxes penetrating the two solenoid assemblies should be almost the same.

A voltage tap is attached to each of the lugs where the electrical connections between serial coils are made. The voltage across each coil is measured using a unity gain difference amplifier¹⁸. All of the voltages are summed, with weights α_i to produce

¹⁷The same technique for distinguishing between the resistive and inductive voltages outlined herein should work even in the absence of symmetry between the coils. By appropriately choosing the values of α_i in equation (4.1), the inductive voltages can be cancelled, though the amplification of a resistive voltage will be different for each coil.

¹⁸Burr-Brown INA117. These difference amplifiers were selected because they are rated to withstand ± 500 V on the inputs, with respect to ground. Since the dump resistor has a resistance of approximately 2Ω these amplifiers can be used with the magnet operating at currents up to 250 A.

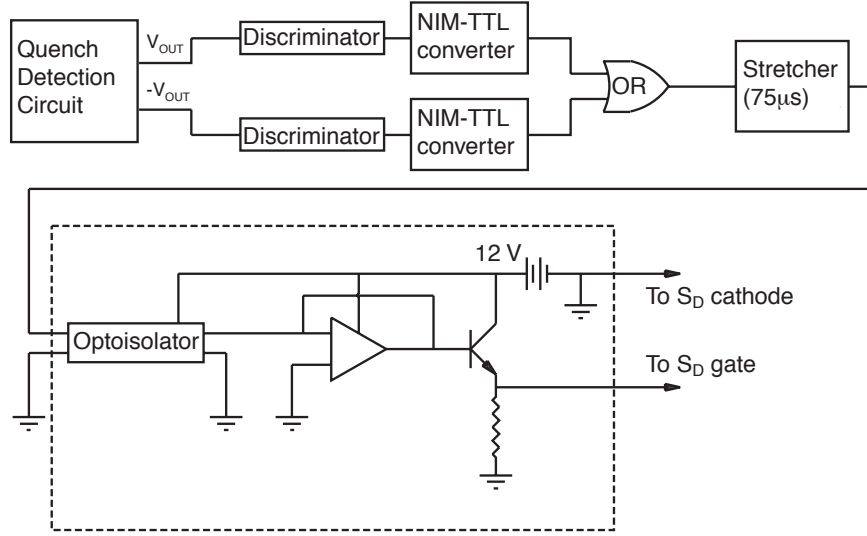


Figure 4.21: Trigger circuit for the quench protection system. This circuit monitors the resistive voltage from the magnet. If the resistive voltage spikes, a pulse is output to the dump SCR which causes it to conduct. The region inside the dashed box is shown in detail in figure 4.22.

the total output voltage V_{out} , where

$$V_{out} = \alpha_1 V_1 - \alpha_2 V_2 + \alpha_3 V_3 - \alpha_4 V_4 + \alpha_5 V_5 - \alpha_6 V_6. \quad (4.1)$$

In the case where the weights are appropriately adjusted (by varying the potentiometers shown in figure 4.20) the inductive voltages present during magnet ramping can be completely cancelled and the output voltage V_{out} will be only due to any net resistive voltage. The sign of V_{out} in the presence of a true resistive voltage can be positive or negative, depending on which coil the quench initiates in.

Since the discriminators used to detect a voltage spike on V_{out} (see figure 4.21) are NIM devices which only look for negative pulses, it is necessary to monitor both V_{out} and $-V_{out}$. The outputs of the two discriminators are converted from NIM logic signals to TTL logic signals and OR'ed. This OR'ed signal is true when a resistive voltage of more than the discriminator threshold (1.00 V) is present on any of the

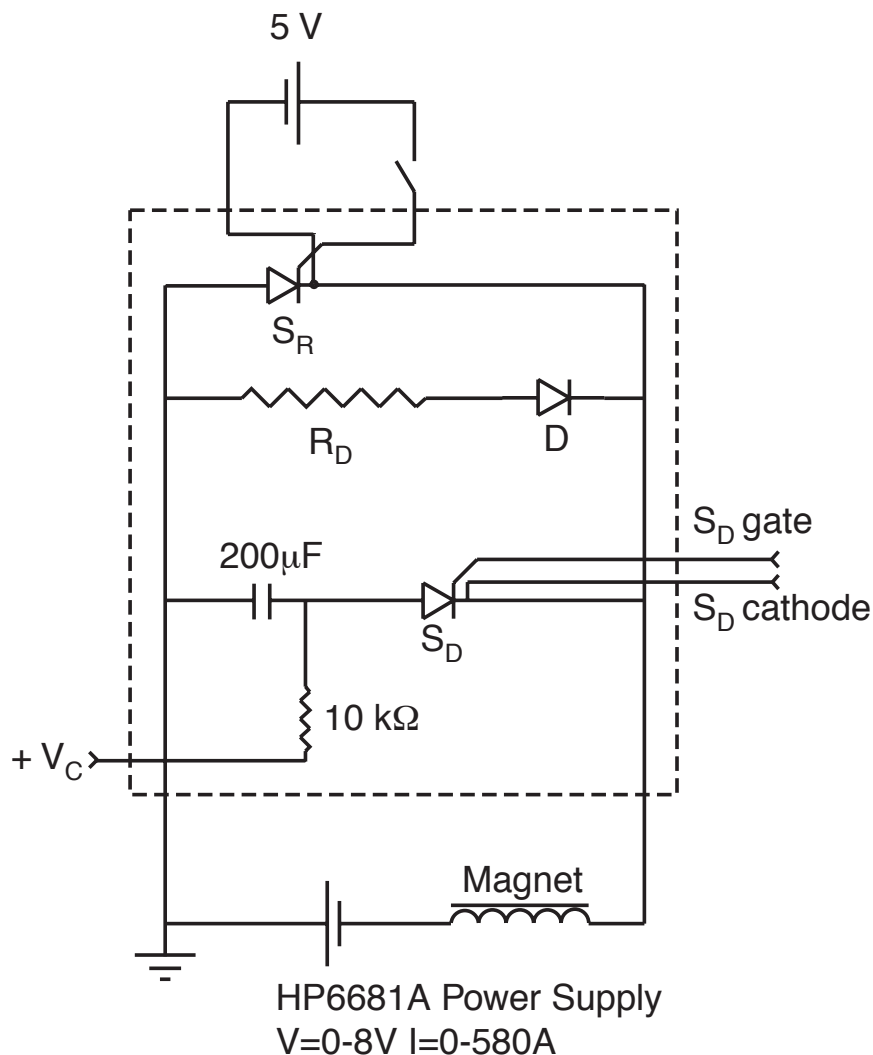


Figure 4.22: The quench protection circuit. Normally the run SCR, S_R , is conducting and the dump SCR, S_D , is non-conducting so current flows from the power supply through the magnet. When the quench detection circuit detects a quench, then the run SCR is disabled and current flows through the dump resistor, R_D .

magnet coils. If this condition is true, then a quench has occurred and it is necessary to trigger the quench protection circuit.

The quench protection circuit is shown in figure 4.22. The design is a modified version of that found in [125]. The magnet power supply is a Hewlett-Packard 6681A, GPIB-programmable DC power supply. This power supply can output up to 580 A at

0–8 V. One terminal of the power supply is connected to the main electrical ground of the local power panel with a braided copper grounding strap. When the circuit is in normal operation, almost all of the current flows through the run SCR¹⁹, S_R , and through the magnet. In order to reach this operating mode, a 5 V floating power supply is briefly switched on (using a manual switch) to supply the gate-cathode bias necessary to put the run SCR into conduction.

A 200 μF capacitor is biased with a positive voltage V_C (ranging from 50 – 100 V at various times²⁰). When a quench is detected, the discharge SCR, S_D , is put into conduction by the quench detection circuit. A trigger from either discriminator (see figure 4.21) feeds a pulse stretcher which outputs a 75 μs pulse. This long pulse is fed through an optoisolator and buffering amplifier to produce a 12 V (floating) bias on the dump SCR's gate-cathode junction. The optoisolator is necessary since the dump SCR's cathode may be floating at a high voltage during the quench due to large resistive and inductive voltages across the magnet. The pulse stretcher is used so that the dump SCR is kept conducting long enough to discharge the capacitor fully. The discharging capacitor counteracts the magnet current which would normally be

¹⁹A Silicon Controlled Rectifier (SCR) is a three terminal (gate, anode, and cathode) solid state device similar to a transistor. The connection between anode and cathode is normally non-conducting, but can be made conducting by applying a positive bias voltage between the gate and the cathode. The anode-cathode voltage drop of the SCRs used is approximately 1.2 V. Once in the conductive state, the SCR will remain conducting until the gate-cathode bias voltage is removed and the anode-cathode current stops flowing for longer than t_q (15 μs in our case). The run SCR is water cooled since it dissipates up to $1.2 \text{ V} \times 250 \text{ A} = 300 \text{ W}$.

²⁰The primary effect of varying the voltage on the capacitor is to change the discharge time of the capacitor. At some point during the magnet test, the quench protection circuit failed to fire and it was mistakenly believed that the cause could be that the V_C was too low. At that point V_C was raised from 50 to 100 V. It was later determined that the quench protection circuit failed to fire due to incorrect configuration of the discriminators, but V_C was left at 100 V.

flowing through the run SCR for a time $\tau = CV_C/I_{max} \approx 40 \mu\text{s}$ for the worst case of $V_C = 50 \text{ V}$ and $I_{max} = 250 \text{ A}$. As a result, the run SCR goes into non-conducting mode since no current flows through the run SCR for a time longer than t_q and there is no gate-cathode voltage (it was removed right after initially putting the run SCR into conduction). Thus, after the capacitor is fully discharged current flows only through the dump resistor R_D ²¹. The diode D is present to prevent the capacitor from discharging through the dump resistor.

The dump resistor is made from a 44 m length of 2 mm diameter soft steel wire. The total resistance was 1.86Ω . The resistor had mass 1070 g so that even if all of the energy from a 225 A quench were dissipated in the resistor it would be heated by $\sim 50 \text{ K}$ (neglecting air cooling). The L/R time constant of the magnet-dump resistor circuit is approximately 0.5 s.

All of the wires connecting components of the magnet circuit were made with 3/0 welding cable. At the top of the dewar, the 3/0 welding cable connects to vapor cooled current leads²². The vapor cooled leads pass through 4 helium exhaust baffles in the upper section of the vertical dewar and end just below the lowest baffle. A copper/superconducting bus bar is attached to the bottom of each vapor cooled lead and extends to several centimeters below the 1K Pot intake line. A twisted pair

²¹The capacitor is not removed from the circuit until the pulse on its gate ends and the current flow through it goes to zero. Since the magnet and capacitor form an LC circuit, current will continue to flow from the capacitor to the magnet until the capacitor reaches a net negative voltage of $-V_C$. At that point no current will be flowing through S_D . Since the gate pulse ($75 \mu\text{s}$) is chosen to have a length shorter than the LC time constant ($\approx 400 \mu\text{s}$), then the first time the current through S_D reaches 0 conduction will stop. After that point the capacitor will be recharged to its original voltage V_C . The $10 \text{ k}\Omega$ resistor between the capacitor and its charging power supply ensures a charging time constant much longer than the discharge time.

²²American Magnetics Model L-250 vapor cooled current leads are rated to 250 A. American Magnetics, Oak Ridge, TN, (<http://www.americanmagnetics.com>).




Figure 4.23: The quadrupole and octupole configurations for the racetrack coils are shown.

of copper/superconducting wires makes a long (approximately 20 cm) solder joint with the bus bar and extends down to the T-section of the dewar where it joins the magnet's leads. The helium exhaust from the dewar must be routed through the vapor-cooled leads to provide cooling. The magnet may be operated as long as the helium level is above the bottom of the bus bar.

4.3.3 Magnet Training and Quench Protection Performance

Magnet training was performed in several stages. Before attaching the four racetrack coils to the magnet form, the coils were trained one at a time, each reaching a current of at least 275 A [65]. The entire “quadrupole assembly” of four racetrack coils was assembled and tested in January 2001. It was later determined that the electrical connections between the coils had been made in the wrong order and instead of being connected as a quadrupole, the assembly actually formed an octupole (see figure 4.23). Although the forces on the assembly are certainly different in the octupole and quadrupole configurations, the testing of the octupole trap reveals im-




Figure 4.24: The training curve of the octupole assembly (without solenoids). The coil in which the quench initiated is indicated, if it was determined.

portant information about the performance of the quench protection system, the risks due to high energy quenches, and the training behavior of the magnet. Unfortunately the octupole geometry is not useful for trapping neutrons due to its much lower trap depth. The octupole trap has a depth of 0.37 T at 170 A, as opposed to a depth of 1.1 T for the quadrupole trap at the same current²³.

Initially the octupole assembly was trained without solenoids. The training was performed in a 10" inner diameter test dewar. The training curve is shown in figure 4.24. Other than the anomalously low quench #7 (in which the helium exhaust

²³For UCN production, the lower trap depth of the octupole geometry is partially compensated for by its flatter radial field profile. The resulting larger phase space for UCN production in the octupole is an advantage over the quadrupole at equivalent trap depths. At equivalent operating currents, however, the much higher trap depth of the quadrupole dominates production considerations.

was not directed through the vapor cooled current leads), the maximum current increases from quench to quench. The very slow (~ 0.7 A/quench) increase from quenches 9 to 45 is excessive training. We do not have a specific model for the linear nature of the training curve and its small slope. The linearity suggests the elastic deformation of some part of the system, perhaps the compression rings. The resistance of the assembly was $272\ \Omega$ at 300 K and $18\ \Omega$ at 77 K. The inductance of this assembly was approximately 1 H.

In some of the quenches from the octupole training sequence it was possible to determine in which coil the quench initiated. The trigger output of the quench detection circuit was used to trigger a four channel digital oscilloscope. Each of the four digitized channels measured the output voltage of one of the difference amplifiers. A typical trace of voltage on the four coils as a function of time is shown in figure 4.25. In this case, the quench initiates in coil 4 as can be seen from the non-zero voltage on coil 4, which begin about $140\ \mu\text{s}$ before the quench protection circuit triggers. Once the voltage on coil 4 has begun to oscillate, oscillating voltages are induced on the other coils as well, until the resistive voltage grows large enough to trip the quench detection circuit²⁴. A disproportionate number of quenches occurred in coil 4. This may be an indication of some flaw in the winding of coil 4, or some asymmetry in the compression scheme.

²⁴The long delay between the initiation of the quench and the triggering of the protection circuit is due to the high (1 V) threshold on the quench detection discriminators. It was found that if this threshold was lowered, the quench detection circuit would falsely trigger. While ramping the current up, there are always small voltage oscillations due to imperfect subtraction of the inductive voltage. There are also short time scale voltage spikes, likely due to small cracking of the epoxy below the level necessary to initiate a quench of the entire magnet. Reliable operation with a more sensitive threshold may be possible if V_{out} is passed through a low-pass filter before discrimination.




Figure 4.25: The voltage on each of the four coils in the octupole assembly is shown as a function of time during a quench (#45). The voltage on coil 4 is the first to oscillate, beginning at a time $140 \mu\text{s}$ before the quench protection circuit triggers.

Following the training of the octupole assembly, the entire octupole magnet trap was trained in the 10" inner diameter test dewar. The outer diameter of the trap assembly was approximately 9.75". Each of the solenoid assemblies has resistance 150Ω at 300 K for a total resistance of 570Ω for the trap. The inductance of the entire octupole trap was measured to be 2.6 H. The training curve is shown in figure 4.26. The two groups of low quenches are believed to be false triggering of the quench protection circuit on noise, rather than real quenches. When the quench protection circuit trips, however, the sudden large voltage placed across the magnet by the capacitor can induce a quench. Therefore it was not possible to determine which quenches were initiated inside the magnet and which were induced by the quench protection circuit triggering on voltage noise. For most of the quenches from




Figure 4.26: The training curve of the first assembly of the octupole trap (with solenoids). The low quenches are believed to be false trips of the quench detection circuit.

28–36 the quench protection circuit did not fire (as determined by monitoring the temperature of the dump resistor). At the time it was believed that this was due to insufficient voltage on the capacitor²⁵ so V_C was raised to 200 V. On the next quench (#37), the quench protection circuit again did not fire and the magnet was damaged. The cause of damage was determined to be a spark between the two solenoid leads

²⁵It was later determined that the failure of the quench protection circuit to fire was due to the use of $50\ \Omega$ pass-through attenuators in an attempt to increase the discriminator threshold beyond 1 V. In reality the non-inverted output, V_{out} , of the quench protection box could not drive $50\ \Omega$ due to the use of a $10\ \text{k}\Omega$ output resistor R_2 (see figure 4.20). If the quench occurred on one of the two coils which produced a negative pulse on V_{out} the quench detection circuit would not fire. If the quench occurred on either of the other two coils then the detection circuit continued to work (due to the additional inverting amplifier's low output impedance). The correction of this flaw is described below.

where they crossed over each other and were pinched between the aluminum solenoid form and the inner wall of the test dewar²⁶. The spark vaporized portions of the two leads in addition to several layers of wire in one of the bucking coils.

After re-winding the ruined bucking coil, the two solenoid assemblies in series (without the octupole) were trained in the bucket dewar²⁷. The training curve for the solenoids is shown in figure 4.27. The first quench of the solenoids occurs near the current of the last quench of the octupole trap and the training curves for both of those assemblies follow approximately the same slope. This suggests that the training of the solenoids may have limited the training of the entire trap. Quench #3 occurred while the magnet was sitting at a stable current of 170 A (not being ramped). This spontaneous quenching was unexpected, and was seen on several other occasions (see below). The training curve levels off between 220 A and 230 A for quenches 16 to 25. This behavior was taken as an indication that the maximum performance of the solenoid assembly had been reached. In all further tests with which involved the solenoids and octupole or quadrupole assembly, the maximum solenoid training

²⁶At the time the only insulation on the leads was mylar tape (Scotch Tape #56, 3M Corporation). Following the catastrophic loss of the solenoid, all leads were insulated with Teflon spaghetti tubing rated to > 1000 V.

²⁷For these tests the solenoid assemblies were mounted with less of an axial separation than they would be when part of the magnetic trap. This was done intentionally to raise their inductance as high as possible in order to test the performance of the quench detection circuit under situations as similar as possible to those in which it had failed to trigger. In this configuration the assembly had a measured inductance of 1.9 H. The quench protection circuit did not fire after the first quench. The flaw in the quench detection circuit described above (too high of an output impedance on the non-inverting AMP) was then discovered and the output impedance was lowered by changing R_2 to $500\ \Omega$ (see figure 4.20). This had the effect of reducing the gain of the summing amplifier from unity to about $1/20$. The discriminator thresholds were hereafter set to 0.15 V (or about 3 V disbalance signal). The maximum possible output of the disbalance signal was measured to be 14.2 V. There were no further problems with the quench protection circuit failing to fire.




Figure 4.27: The training curve of the two solenoid assemblies in series (without the octupole).

current of 230 A was never exceeded. This offers additional support for the theory that the magnet performance was limited by the training of the solenoids.

Following the training of the solenoid assembly, the entire octupole trap was re-assembled and trained. The training curves from this testing are shown in figure 4.28. This assembly of the octupole trap was cooled down twice in the test dewar for training. Following quench #12 of the first cooldown, the vapor cooled current leads plugged with ice when the magnet was left overnight. It was then necessary to warm up to unplug the leads. The apparent continuity between the training curves for the first and second cooldowns suggests that the magnet “remembers” its training from cooldown to cooldown. This is consistent with a model where the function of the training is to crack epoxy rich regions in the coils.

As part of the testing of the octupole trap assembly, the performance of the quench




Figure 4.28: The training curves of the second assembly of the octupole trap (with solenoids). A second cooldown was required to continue training after the current leads plugged following quench #12 of the first cooldown.

protection circuit was also tested. The voltage of a short section of the dump resistor, $V_{section}(t)$, (with resistance, $R_{section} = 0.072 \Omega$) was measured by a digital oscilloscope. The scope was triggered by the quench detection trigger. Under the simplest model, after the quench protection circuit has fired, the system can be considered to be a closed circuit with a resistor (the dump resistor and the magnet's resistance in series) and an inductor (the magnet) in series. In this case the total voltage drop around the closed loop must be zero and:

$$-L \frac{dI}{dt} + IR = 0 \quad (4.2)$$

where $L = 2.8$ H is the inductance of the magnet, $R = R_D + R_{magnet}$ is the total resistance of the dump resistor and magnet, and I is the current. Both I and dI/dt are determined from the measurement of $V_{section}(t)$. Since the temperature of the

dump resistor rises from 20 °C to a maximum of about 50 °C during the largest quenches, it is assumed that $R_{section}$ and R_D do not vary with time. Then the current is simply

$$I = \frac{V_{section}}{R_{section}} = \frac{V_{section}}{0.072 \Omega} \quad (4.3)$$

The derivative is calculated from the time sequence of the voltage. The total resistance of the system can then be calculated from

$$R = \frac{L}{I} \frac{dI}{dt} \quad (4.4)$$

With I and R calculated in this way, we can compute the total energy dissipated in the system from

$$E_{total} = \int I^2 R dt \quad (4.5)$$

and compare this with the total energy dissipated in the dump resistor

$$E_D = \int I^2 R_D dt \quad (4.6)$$

where $R_D = 1.86 \Omega$ was measured separately by running a low current through the resistor while measuring the voltage drop across it. If this methodology is correct, then E_{total} should also equal the total stored energy in the magnet $LI^2/2$.

The results of the analysis of the quench protection circuit during the training of the octupole trap are shown in table 4.1. The measured and predicted energy stored in the magnet are very consistent. The fraction of the magnet's stored energy which is dissipated in the dump resistor decreases linearly with current, as shown in figure 4.29.

At high currents the quench protection circuit offers very little protection. One possible explanation is that for large quenches the dump resistor becomes small compared to the resistance of the magnet (see figure 4.30) and more power is dissipated in the larger resistor. In this case it should be possible to increase the protection

Table 4.1: The performance of the quench protection circuit for protecting the octupole trap. The dump resistor $R_D = 1.86 \, \Omega$. The inductance of the magnetic trap is 2.6 H. The various parameters are calculated using the model in the text. The efficiency of the quench protection circuit $\eta_D = E_D/E_{total}$ is also shown.

Quench #	I_{quench} [A]	$LI^2/2$ [J]	E_{total} [J]	E_D [J]	η_D
1	125	20323	20682	17354	0.839
2	166	35823	35761	19021	0.532
3	171.2	38102	38224	17992	0.471
4	165	35393	35545	19250	0.542
5	180.7	42448	43050	22885	0.532
6	183	43535	44158	20194	0.457
7	174	39359	40202	22692	0.564
8	182.2	43155	43907	22793	0.519
9	185.4	44686	45281	19397	0.428
10	185.7	44830	45961	22724	0.497
11	193	48424	49335	20007	0.406
12	195.8	49839	50777	22139	0.436
13	187	45460	46099	19640	0.426
14	195.9	49890	50690	19139	0.378
15	197.4	50657	51761	19759	0.382
16	205.5	54899	56139	18181	0.324
17	213.9	59479	60912	17557	0.288
18	213.7	59368	60623	17623	0.291
19	225.6	66164	67336	16802	0.250
20	227.4	67224	68109	16611	0.244
21	230.2	68890	69572	16162	0.232
22	221.3	63666	63921	15778	0.247
23	232	69971	70715	16407	0.232
24	233.3	70758	71528	16402	0.229

efficiency by raising the resistance of the dump resistor (see below for a limited test of this hypothesis). As the dump resistor becomes larger, then the voltage $I_{max}R_D$ across the magnet becomes very large and the risk of damaging the magnet due a spark (as in the case of the solenoids, above) becomes significant²⁸. It is also worth

²⁸The leads of the magnet are insulated with teflon spaghetti tubing which is rated to > 1000 V. Since the magnet is surrounded by liquid helium, which has a large




Figure 4.29: Fraction of the magnet's stored energy dissipated in the external dump resistor.

noting that the resistance rises faster in the larger quenches.

The ultimate solution to the quench protection problem lies in the direction of building magnets with lower inductance. Since the time constant for an R-L circuit is R/L , in a lower inductance magnet the current will drop faster. This will lead to more power dissipated in the external resistor since the resistance of the magnet will be lower for any given current. The inductance can be decreased by using larger diameter wire to wind the magnet coils. This requires an increase in the operating current in order to maintain the same trap depth. The increased operating current

dielectric constant, the effective insulation of the magnet leads to any external surface is probably even higher. At some point, however, the formvar insulation of the magnet wires themselves will break down and layer-layer sparking within the coils may occur. In any event, we were limited to the much lower maximum voltage of 500 V, which was the maximum rating of the difference amplifiers used in the quench detection circuit.




Figure 4.30: The resistance of the entire system (dump resistor + magnet) is shown as a function of time following the triggering of the quench detection circuit. Several representative quench currents are shown. At higher currents the maximum resistance of the magnet is larger and the quench protection circuit is less effective.

means that the dump resistor must be decreased (so as to not exceed the maximum voltage). In the end, going to higher current/lower inductance is an improvement, since the total stored energy is proportional to $LI^2/2$ so that for a fixed stored energy, which is equivalent to a fixed trap depth, the inductance drops more quickly than the current increases. Although operating at higher currents introduces additional complexity (multiple power supplies, large copper bus bars rather than welding cable for interconnects, high boiloff from large vapor cooled leads), established technical solutions (such as high- T_C current leads) exist. There is a large body of experience in manufacturing low-inductance, high-current magnets for accelerator magnets (an example of a high current quadrupole magnet appropriate for magnetic trapping of UCN is discussed in section 7.7).

The efficiency of the quench protection circuit is of great concern because we desire to lower the risk of a quench damaging the magnet. A catastrophic quench did occur in the octupole trap while it was installed in the main apparatus. Following the training shown in figure 4.28 the octupole trap was installed in the main apparatus and re-trained. The magnet was quenched six consecutive times. The first quench was at 159 A and the quench current increased rapidly until quenches 4 to 6 were all at 210 A. The magnet remained cold but un-powered for three days, after which the current was raised to 238 A without the magnet quenching. On the following cooldown, a similar behavior was observed, where the magnet quenched twice at 205 A and 200 A, was operated at 200 A for 3 days, was turned off for 3 days and then ramped to 225 A without quenching. A possible explanation for this behavior is slow “training” of the magnet while it is running at high current, as the constant stress on the magnet cracks epoxy-rich regions without producing a quench of the entire magnet. Over time, enough epoxy-rich regions may be broken down that a higher current can be achieved without quenching. On two further cooldowns, the magnet was operated semi-continuously (taking a series of positive and negative runs as described in chapter 5) for three and a half days at 225 A without incident.

A spontaneous quench in the octupole assembly resulted in catastrophic failure of the magnet. After several hours of semi-continuous operation at 225 A, the magnet quenched while sitting at a steady current of 225 A. The magnet was not being ramped when it quenched, nor could any other external event be identified as the cause of the quench. The quench apparently started in part of the lead connecting two coils, since a section of lead was vaporized, as were several layers of wire in one of the quadrupole coils where the lead entered the coil. No model exists for the cause of this spontaneous quench.




Figure 4.31: The training profile of the quadrupole trap (with solenoids) is shown.

When the magnet was disassembled to replace the destroyed coil, it was discovered that the racetrack coils had been connected incorrectly. This had resulted in the “quadrupole” assembly actually being an octupole assembly, as shown in figure 4.23. The magnet assembly was then reassembled, using a spare racetrack coil, with the quadrupole assembly properly wired. The entire quadrupole trap was then trained in the test dewar. The training profile of the quadrupole trap is shown in figure 4.31. In order to reduce the probability of another catastrophic failure of the magnet, we determined that we would train the magnet up to a current at which the energy dumped in the magnet during a quench was no more than one half that dumped into the octupole trap during a 225 A quench. From table 4.1 we can see that the maximum dumped energy under this rule should be no more than 25 kJ. Using the same technique described above for the octupole trap, the performance of the quench

Table 4.2: The performance of the quench protection circuit for protecting the quadrupole trap. Quenches in which the magnet was manually tripped are denoted with a quench number preceded by T. The dump resistor had resistance $R_D = 1.86 \, \Omega$ for all quenches except quenches 11, T4 and T5, in which $R_D = 2.54 \, \Omega$. The inductance of the magnetic trap is 3.5 H. The various parameters are calculated using the model in the text. The efficiency of the quench protection circuit $\eta_D = E_D/E_{total}$ is also shown.

Quench #	I_{quench} [A]	$LI^2/2$ [J]	E_{total} [J]	E_D [J]	η_D
T1	50.0	4412	3593	3446	0.959
T2	100.0	17650	13969	14455	1.031
T3	125.0	27578	21714	23661	1.087
1	146	37623	36836	28587	0.776
2	148	38661	36923	28860	0.782
3	154.3	42022	41004	29948	0.730
4	160.7	45580	45212	31050	0.687
5	160.6	45524	44972	31031	0.690
6	164.0	47471	47795	31428	0.658
7	166.9	49165	49754	31794	0.639
9	170.2	51129	51485	32419	0.630
10	160.0	45184	44743	30987	0.693
11	172.3	52398	53463	32506	0.608
T4	100.0	17650	17233	17545	1.023
T5	150.0	39713	39547	28412	0.718
12	171.2	51731	52210	28186	0.540
13	171.8	52094	55533	34388	0.619
14	175.0	54053	57952	34718	0.599
15	175.7	54486	58808	34863	0.593
16	181.4	58079	63084	35290	0.559

protection circuit was studied while training the quadrupole trap, and is shown in table 4.2. When quenched at 181.4 A, the energy dumped into the magnet is about 10% above our target maximum. Hence, it was determined that the magnet should be operated at a current no higher than 180 A.

For trips T4 and T5 and quench #12, the dump resistance R_D was increased from $1.86 \, \Omega$ to $2.54 \, \Omega$. Contrary to the expectation that the quench protection efficiency would increase for larger R_D , a lower quench protection efficiency was observed

(compare for example quench 12 and quench 13). We do not have a model for this observation. The behavior of the quench protection circuit as R_D is varied deserves further study.

Given the fairly long delay in detecting quenches (see figure 4.25), it was thought that the decreased efficiency of the quench protection circuit at high currents might simply be due to late detection. In order to test this hypothesis, the quench protection circuit was tripped with the magnetic trap at 150 A. Tripping the quench protection circuit puts a large voltage across the magnet which then leads to a quench. In a trip, the delay between quench initiation and quench detection is zero. Since the same quench protection efficiency measured for a 150 A trip is consistent with that for a similar size quench (see table 4.2), the theory of low protection efficiency due to late detection can be ruled out.

When the quadrupole trap was cooled down in the main apparatus, following the above training, additional training was required in order to achieve an operating current of > 170 A. The training curve of the quadrupole trap in the main apparatus is shown in figure 4.32. Since the current seemed not to increase much beyond 170 A after numerous quenches, a running current of 170 A was selected for UCN magnetic trapping measurements.

4.4 Detection System

4.4.1 Overview

The development of the detection system, its optimization, and the details of its performance are described in depth in [65]. As described in section 1.5.3, the detection system converts EUV photons from helium scintillation to blue photons, pipes them




Figure 4.32: The training profile of the quadrupole trap (with solenoids) in the main apparatus is shown. Based upon this training curve an operating current of 170 A was selected for the trapping experiments.

out of the apparatus, and detects them with room temperature photomultipliers. Each of these functions will be described herein.

The low temperature portion of the detection system is shown in figure 4.33. EUV scintillation light is produced by the recoil of neutron decay electrons through the superfluid helium filling the trapping region. The trapping region is surrounded by a 1 mm thick cylinder of GoreTex, with inner diameter 8.4 cm. On the GoreTex is evaporated a layer of $200 - 400 \mu\text{g cm}^{-2}$ of the organic fluor 1,1,4,4,-tetraphenyl buta-1,3-diene (TPB). TPB absorbs photons over a broad wavelength region from the soft UV to X-rays and emits blue light with a spectrum peaked at 440 nm and a width of approximately 50 nm [63]. A variety of fluors were studied and TPB was selected as the one with the highest fluorescence efficiency (FE, the ratio of blue photons




Figure 4.33: The low temperature part of the detection system is shown.

emitted to EUV photons absorbed) was chosen [62]. The FE of an evaporated TPB film was measured to be 3.7 ± 0.1 relative to Sodium Salicylate [62], which has been independently measured to have an absolute FE of 0.37 [64]. Combining these two measurements, the absolute FE of evaporated TPB is 1.4.

Since neutron decays can occur anywhere in the trapping region, it is important to extract blue scintillation light produced all along the trap's outer diameter. For this purpose, the TPB fluor is evaporated onto GoreTex²⁹, a material with high

²⁹GoreTex GR gasket material was purchased in 1 mm thick sheets. W. L. Gore and Associates, Inc., (<http://www.gore.com>). All GoreTex gasket material is sold with the GoreTex trade name prominently displayed on one side of the material (the side not coated with TPB, in our case). Gasket material may be purchased with the logo painted on with a red paint or simply embossed in the GoreTex itself. Rabbit tube tests of both types of GoreTex sheet indicated that the sheet with the painted logo contained about 3 times as much iron as the sheet with the embossed logo. Hence, the gasketing material used has the GoreTex logo embossed on the back side. While quantitative analyses of absolute impurity levels in the GoreTex were not made, the rabbit tube tests did clearly indicate the presence of manganese and aluminum, in addition to iron.

diffuse reflectivity in the blue³⁰, which maintains its high reflectivity even after large radiation doses [127]. Since the trap length (22 cm) is only a few times its diameter (8.4 cm) a large fraction of the blue light should reach an end of the trapping region after only a few bounces.

At the detector end of the trapping region is an optically clear, neutron-absorbing, beamstop. This beamstop is made from high purity boron oxide³¹, (B_2O_3). The blue transmission of the beamstop was measured using an LED and a photodiode. Since boron oxide is a hygroscopic material whose surface becomes cloudy when exposed to air, the transmission of the beamstop was measured as a function of time, while the beamstop was exposed to atmosphere. While the initial transmission of the boron oxide was 87.9%, it dropped to 82.2% after fifteen minutes of exposure to room air, and 73.5% after 30 minutes. It is thus necessary to polish the beamstop after any exposure to air, store it in a desiccator, and minimize exposure to atmosphere during insertion into the cell. During assembly, the lightguide and beamstop are the last items inserted into the cell and the snout indium seal are made immediately after their insertion. Prior to inserting the beamstop, the cell is flushed for several minutes with dry nitrogen gas. Once the beamstop is in place and the indium seal made, the cell is alternately pumped and flushed with ultrapure helium to remove any remaining water vapor. The effective exposure time of the boron oxide is thus reduced to about five minutes.

The beamstop is mounted on the end of a 8.7 cm diameter, 43.2 cm long, ultra-

³⁰GoreTex is manufactured from expanded polytetrafluoroethylene (PTFE). The diffuse reflectance of pressed PTFE powders has been measured to be $> 99\%$ for blue light [126]. GoreTex and other white teflons are widely used in high energy physics for wrapping detectors to increase light collection efficiency.

³¹The manufacturer specifies the Boron Oxide as $> 99.99999\%$ purity. All-Chemie Ltd., Mt. Pleasant, SC, (<http://www.allchemie.thomasregister.com>).

violet transmitting (UVT) lightguide. The beamstop is held to the lightguide by FEP teflon tape³² around its outer diameter. The lightguide transmits the blue light pulse to the end of the cell, where it passes out through a window (see figure 4.11). The lightguide is made from UVT acrylic since UVT acrylic has less blue light absorption than standard, UVA (ultra-violet absorbing), acrylic [65].

Since the acrylic lightguide contracts more than the cupronickel cell tube, care was taken to fix the position of the lightguide within the cell, relative to the 4 K window. The position of the lightguide is set by three stainless steel pegs driven radially into the lightguide, which are pressed up against a step in the cell endcap during assembly. The pegs are 3 mm diameter and 1 cm long and are inserted into 5 mm deep milled holes evenly spaced around the circumference of the lightguide. The distance between the end of the lightguide and the pegs is chosen to set a gap of approximately 2 mm between the end of the lightguide and the inside of the snout window. Without such pegs, the gap could be much larger, reducing the detection efficiency. On one occasion the gap was set too small and the acrylic lightguide popped the window off of the snout between 77 K and 4 K.

The windows and optics at 4 K and warmer are shown in figure 4.34. The 4 K window and 77 K lightguides are described in detail in section 4.2.4. The 4 K window is a composite structure of 6 mm thick acrylic onto which 0.33 mm thick crystalline quartz wafers are glued using Stycast 1266 epoxy. During the cooldowns in which data was taken, the quartz wafers cracked and partially delaminated, resulting in lowered detection efficiency. The 77 K lightguide is an 11.4 cm diameter UVT acrylic lightguide. The cold face of the lightguide is embedded with 40 wires of diameter 1 mm (see figure 4.15). This wire pattern, needed for cooling the face of the lightguide,

³²McMaster-Carr, New Brunswick, NJ, (<http://www.mcmaster.com>).




Figure 4.34: The detection system optics are shown.

covers 11% of the lightguide’s face, but reduces the light collected by 30%³³. At room temperature, light is split by an aluminized splitter before being coupled to two photomultiplier tubes. The results of a measurement of the overall detection efficiency of the detection system are reported in section 6.6.

4.4.2 Photomultipliers

The photomultipliers used in this work are Burle 8854, 12.7 cm, head-on bialkali photomultipliers³⁴. The 8854 is a 12.7 cm diameter version of the 5.1 cm diameter Burle 8850, which was used in our earlier work [42] and which we characterized extensively. The 885x family of photomultipliers is distinct for its cesiated gallium-

³³The reduction in light collection efficiency was measured in an independent calibration. Light collection is reduced by an amount greater than the coverage fraction due to the finite thickness of the wires and the angular spread of light incident on the face of the lightguide.

³⁴Burle Industries, Lancaster, PA, (<http://www.burle.com>).




Figure 4.35: The “spillover” of single photon light pulses to larger photoelectron pulses is shown. A Burle 8850 PMT was exposed to a single photon LED light source. The difference between the PMT’s pulse height spectra with the LED on and off was measured. The curve shown is the fraction of single photon LED pulses which result in an observed PMT pulse of \geq threshold photoelectrons.

phosphide first dynode. This dynode has very high gain compared to the beryllium copper dynodes typically used in PMTs. The high gain of the first dynode results in excellent discrimination between noise and single photoelectron pulses. Moreover, it results in low gain dispersion, so that one, two, and three photoelectron pulses can be separately distinguished in the PMT’s pulse height spectrum (see figure 5.6).

Several characteristics of the 885x family photomultipliers are of concern for this work. Although the PMTs have excellent gain dispersion, it is not perfect. It is possible for single photon events to result in large pulses that fall outside of the “single photoelectron peak” in a standard PMT pulse height spectrum. A measurement of this “spillover” effect is shown in figure 4.35. An LED was operated at a very low volt-

age, so that it produced infrequent, time-uncorrelated, single photons. Pulse height spectra with the LED on and off were measured and subtracted. The tube studied had been used in the laboratory environment for at least a year and demonstrated the signs of helium afterpulsing (see below). Hence, this measurement of the spillover includes the effect of large afterpulses to some extent.

A second effect, which is particularly important is afterpulsing. It has been well-established that even single photon events can result in large photoelectron afterpulses [128, 129, 130]. These afterpulses are caused by ionization of residual gases (primarily helium) inside of the photomultiplier. Positive ions are accelerated in the opposite direction from electrons, and may be driven back into the photocathode or dynodes, releasing a delayed pulse of many electrons. Pulses of up to 20 photoelectrons have been seen after single photoelectron initial pulses [130]. These afterpulses have a characteristic delay from the initial pulse which depends on the mass of the ionized gas, the dynode voltages, and the dynode separation. In our earlier work, we observed afterpulsing at about 500 ns, with 8850 PMTs operated at 1700 V. Afterpulsing tends to increase over time, when PMTs are used in a laboratory environment with excess helium concentrations (such as helium boiloff from cryogenic experiments) [129]. There may also be a longer-time component to afterpulsing, on the tens of μ s time scale, though some authors have reported this effect[131] and others found no evidence for it[132]. The combination of a large rate of single photon luminescence backgrounds with the possibility of multi-photoelectron afterpulses and/or spillover, introduces the possibility of observing single-photon luminescence signals even with multi-photoelectron thresholds.

Since the gain of the photomultipliers is very sensitive to magnetic fields, great care was taken to shield the PMTs from the large fringe magnetic field produced by




Figure 4.36: The mu-metal and active compensation coils used for shielding the PMTs from the magnetic trap's field are shown.

our trap. The primary field felt by the PMTs is due to the solenoids. The magnetic shielding setup is shown in figure 4.36. The shielding includes both passive (high magnetic susceptibility material) and active (two compensating coils which partially cancel the trap's field at the PMTs) components. The setup and tuning of these components is described in greater detail in [65].

4.5 Shielding

Different types of shielding are used to minimize the various backgrounds to which the experiment is susceptible. Time varying backgrounds can be caused by neutron activation or luminescence. Neutron activation is minimized through careful selection of materials (as described in chapter 2) and through the use of neutron shielding. Some of the neutron shielding materials luminesce (see section 4.5.1). To the extent possible, non-activating, non-luminescing light-absorbing materials are placed between the neutron shielding and the detectors for the purpose of reducing luminescence (see section 4.5.2). Since the liquid helium is an efficient scintillator for all forms of ionizing radiation, lead shielding is used to reduce the ambient gamma background (see section 4.5.3) and muon veto paddles are used to detect and veto cosmic ray muons (see section 4.5.4).

4.5.1 Neutron Shielding

As described in chapter 2, neutron activation is an extremely problematic background to the measurement of the neutron lifetime. In order to minimize activation, as many neutrons as possible are absorbed in neutron shielding. The portions of the apparatus exposed to neutrons are shown in figure 2.1. The neutron shielding consists of three main sections: interlocking neutron shielding outside of the cell, a tube of neutron shielding along the length of the cell, and a beamstop at the end of the cell.

The neutron shielding outside of the cell on the beam entrance side of the apparatus is shown in figure 4.37. The shielding outside of the cell is made from interlocking pieces of grade AX05 boron nitride. This grade of BN is a high purity, easily machinable, material made by hot pressing hexagonal boron nitride powder. The absorption length of 8.9 Å neutrons in AX05 boron nitride is 60 μm . All of the shielding out-

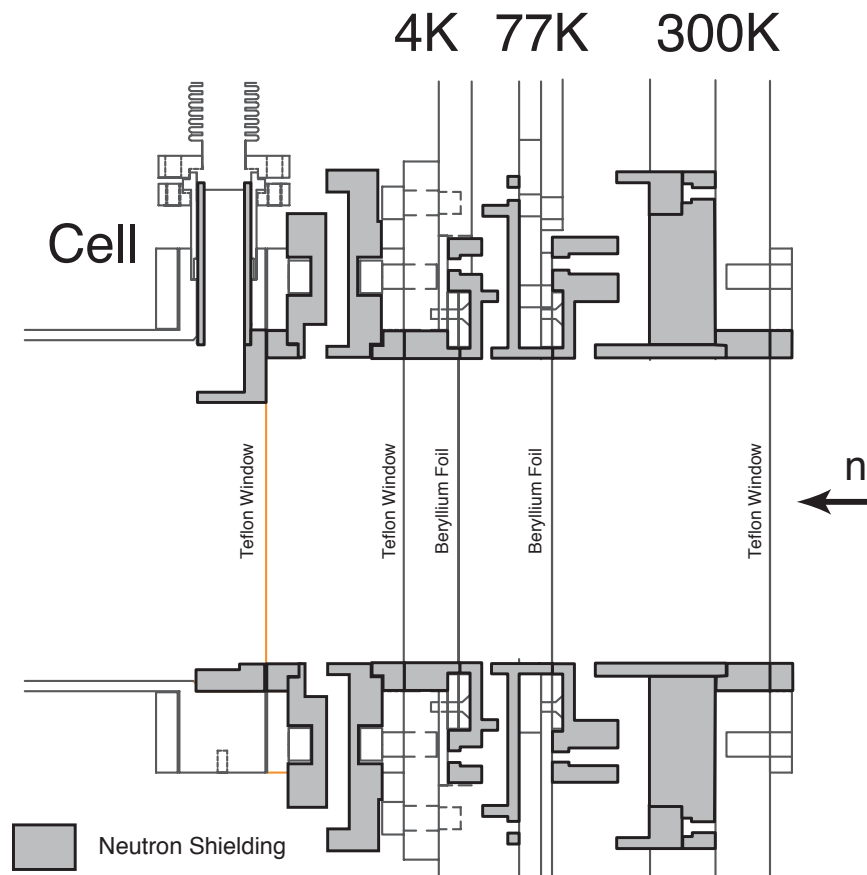


Figure 4.37: Neutron shielding on the beam entrance side of the apparatus is shown. All neutron shielding shown is grade AX05 hexagonal boron nitride.

side of the cell is at least 3 mm thick, so the neutron shielding is completely black. The only concern then is neutrons which scatter multiple times and do not hit the shielding. As can be seen from figure 4.37, at least 3 scattering events, two from the BN itself are necessary in order for neutrons to escape the shielded region. Neutron activation tests of boron nitride have shown no measurable levels of activation, using gamma spectroscopy and beta spectroscopy. Certificates of analysis provided by the manufacturer for the lots of boron nitride used for shielding show < 1 ppm of all tested elements except for manganese (1 ppm), tungsten (5 to 19 ppm, depending on the lot), zinc (3 – 5 ppm) and calcium (71 ppm in one lot, < 1 ppm in all others)

[133]. Chemical analysis was performed by optical emission spectroscopy. None of these elements activates at a level of concern.

Inside of the cell, boron nitride is also used for neutron shielding. As shown in figure 2.1, a tube of neutron shielding lines the cell for the entire length of neutron exposure. This tube is made from four interlocking pieces of grade AX05 boron nitride. Each piece is roughly 30 cm long, has an 8.88 cm inner diameter and a wall thickness of 3.4 mm. The tubes are joined by overlapping adjacent tubes for a length of approximately 1 cm, over which the outer and inner tubes are each reduced to one half of their nominal wall thickness. A single layer of FEP teflon tape around the outside of the joints holds the boron nitride tubes together while they are inserted into the cupronickel cell tube. The BN tube and graphite light shield (see below) are assembled as a single unit, with the TPB coated GoreTex inside, and that entire unit is inserted into the cell in one piece.

Neutrons are absorbed at the end of the trapping region by a boron oxide beamstop. The optical transmission and luminescence properties of the beamstop are described above. The beamstop was 9.6 mm thick Boron Oxide of specified purity $> 99.99999\%$. The absorption length for 8.9 Å neutrons in boron oxide is 60 μm , so the beamstop is black to neutrons. Activation tests of the boron oxide beamstop, using beta spectroscopy, are described elsewhere [65]. The boron oxide beamstop was used because it had much lower luminescence and much less primary activation than lithium fluoride (fluorine activates), with only slightly worse transmission. A boron/lithium glass was also investigated as a possible beamstop but was excluded due to a very strong secondary activation caused by the reaction $^{16}\text{O}(\text{t},\text{n})^{18}\text{F}$, where the tritons are produced by neutron capture on ^6Li . As described in section 2.4.4, the boron oxide is also susceptible to secondary activation by the reaction $^{10}\text{B}(\alpha,\text{n})^{13}\text{N}$.

4.5.2 Luminescence Reduction

Since boron nitride luminesces strongly (see section 2.4.4), a light absorbing tube is placed between the boron nitride and the detection region. The light absorber is made from a 1 mm wall thickness tube of grade AXF-Q1 Poco Graphite, which slides inside of the boron nitride. As with the boron nitride, the graphite tube is made from a series of interlocking tubes, each about 15 cm long. The efficacy of graphite as a luminescence reducer for boron nitride luminescence is demonstrated in section 2.4.4, as is the lack of luminescence from graphite itself.

4.5.3 Gamma Shielding

Backgrounds due to gamma rays coming from outside of the apparatus originate in naturally radioactive isotopes in the building materials of the facility, or may be produced by guides and neutron shielding of other experiments in the guide hall. In order to reduce the background due to those sources, a lead shield was constructed around the horizontal part of the cryostat by stacking lead bricks ($5 \times 10 \times 20$ cm). A floor of lead 5 cm thick was supported just below the bottom of the cryostat (leaving just enough room for the muon paddles). Walls of lead with thickness 10 cm were built around all sides of the cryostat, leaving holes for the very ends of the PMT cans to penetrate, and for the neutrons to enter the apparatus. A top with thickness 10 cm was also built over the horizontal section of the cryostat. The shielding has various gaps due to imperfect stacking of the lead bricks, and for the various holes necessitated by the geometry. The largest hole in the gamma shielding is for the vertical section of the cryostat. The effects of this lead shielding on the measured constant background are described in section 6.2.1.

4.5.4 Muon Veto

In order to veto cosmic ray muons which pass through the apparatus at a rate of about $1 \text{ cm}^{-2} \text{ min}^{-1}$ [5], muon veto paddles are used. These paddles of plastic scintillator coupled to PMTs detect the Cerenkov radiation produced by muons passing through them. As described below (see section 5.1.1), scintillation events observed by the main detector in coincidence with events on the muon paddles are rejected. The paddles and their positions in this work are just as described in [42].

Chapter 5

Data Acquisition and Analysis

5.1 Data Acquisition Hardware

An overview schematic of the data acquisition system (DAQ) is shown in figure 5.1. The data acquisition system is controlled by a Power Macintosh G3/266 computer¹. This computer has three slots which can hold cards communicating via the PCI bus. The total number of slots is increased by using the third internal PCI slot to hold a PCI bridge extender card. This card connects to an external PCI expansion chassis with 7 additional PCI slots². Of the nine total PCI slots, four are used in this generation of the DAQ (2 for waveform digitizers, 1 for the multifunction digital I/O and counter/timer card, and 1 for a GPIB interface card). In the future, two additional slots will be used for counter cards which will allow monitoring of event rates from several stages of the NIM logic circuit (this is analogous to the “binner” data acquisition system used in our earlier work [42]).

¹Apple Computer. Cupertino, California. <http://www.apple.com>.

²Aurora Technologies. Brockton, MA. <http://www.auroratech.com>. The expansion chassis used was Model XP-7.

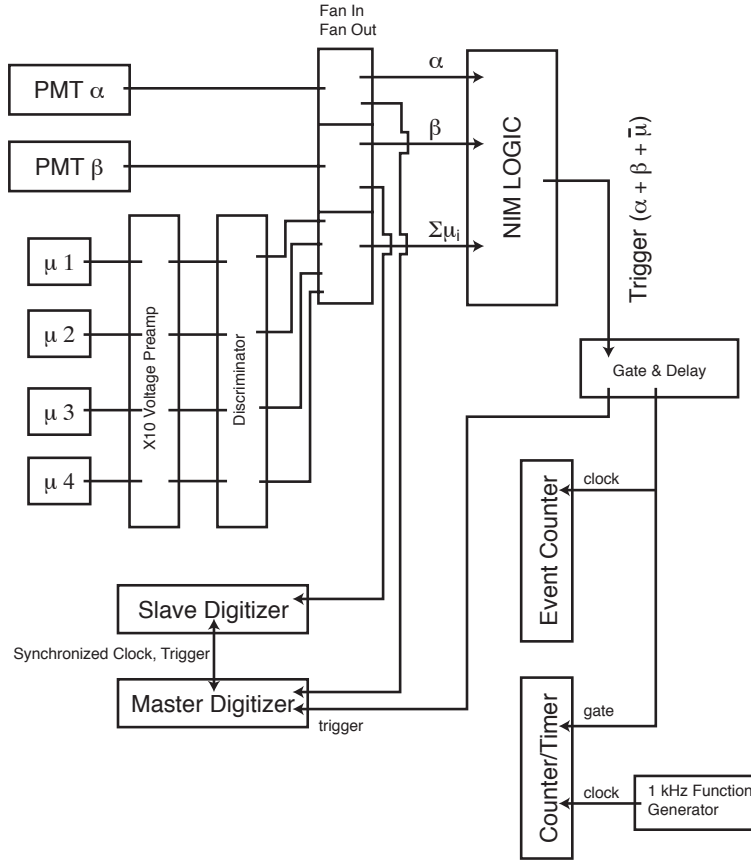


Figure 5.1: A block diagram of the data acquisition system.

The goal of the DAQ is to record multi-photon scintillation events originating within the cell, and reject events produced by cosmic ray muons. Hence, the data acquisition system triggers when coincident pulses are observed on the two main detection photomultipliers which are not coincident with an event on any of the muon paddles. For each trigger, the data acquisition records the time of occurrence and the digitized waveforms from the two main PMTs.

As described in section 4.4.2, two 5 inch diameter photomultipliers (Burle 8854) observe light produced within the detection system. These two PMTs, denoted α and β , are operated in pulse counting mode. To partially compensate for inherent gain

differences in the two photomultipliers used, they were biased at different voltages, $V_\alpha = +2250$ V and $V_\beta = +2200$ V. The anode signal of each photomultiplier was connected to a separate input channel of a linear fan in/fan out³. One of the fan out's outputs is used as the input signal for waveform digitization. Another output is used in the NIM logic circuit which produces the overall system trigger.

Another important input to the NIM logic circuit is the veto signal from the muon paddles (see section 4.5.4). The three muon paddles ($\mu_1 - \mu_3$) under the cell are biased at -2400 V. The fourth muon paddle (μ_4) uses a different model PMT and was biased at -1850 V. The anode output of each muon-detecting PMT is amplified with a $\times 10$ voltage amplifier⁴ before being input to a discriminator⁵. The discriminator produces a digital output pulse whenever the input voltage crosses below a threshold value. Amplification of the PMT pulses is necessary before discrimination in order to increase the size of the PMT pulses beyond the minimum value of the discriminator threshold (-10 mV for the model used). The thresholds of each discriminator were varied while observing the coincidence rate ($\alpha + \beta + \sum \mu_i$). The thresholds were set to maximize this coincidence rate⁶. The chosen muon thresholds lead to rates from each

³Phillips Scientific. Ramsey, NJ. <http://www.phillipsscientific.com>. Model 744 Quad Linear Gate Fan In/Fan Out. The fan out contains unity gain, 250 MHz bandwidth, buffering amplifiers which produce multiple independent outputs for each input signal. This particular fan out also has the property (important for the waveform digitizers, see section 5.1.2) that it limits the output voltage to the range ± 2.5 V.

⁴LeCroy Research Systems. Chestnut Ridge, NY. <http://www.lecroy.com/lrs/>. Model 612 PM Amplifier.

⁵Phillips Scientific. Ramsey, NJ. <http://www.phillipsscientific.com>. Model 708 Octal 300 MHz Discriminator.

⁶Although it would seem that the maximum coincidence rate would occur for the minimum threshold, this was not observed to be the case. As the threshold was lowered into the noise the output rate of pulses from the discriminator was observed to drop. This is likely due to hysteresis in the comparator within the discriminator introduced by the manufacturer to prevent small fluctuations in pulses crossing through

paddle which are much higher than the expected muon detection rate. Nevertheless, the rates are low enough that the occurrence of false coincidence between the muon paddles and the detection PMTs is low. The outputs of the four muon discriminators are fed into one of the channels of the same PS 744 quad fan in described above, which linearly adds the four signals. The output signal from this device ($\Sigma\mu_i$) contains a true logic pulse whenever an above-threshold event occurs on any muon paddle. The measured rate of $\Sigma\mu_i$ is 6.0 kHz, while the coincidence rates $\alpha + \beta$ (while observing the constant background only) and $\alpha + \beta + \Sigma\mu_i$ are 36.3 and 17.1 Hz respectively. The rate of false vetoing is 0.1 Hz⁷.

5.1.1 Trigger Logic

The analog signals from the two main detector PMTs and the muon veto signal are combined using NIM logic modules to generate the master trigger for the DAQ. A schematic of the hardware used is shown in figure 5.2. A trigger is produced whenever an above-threshold coincidence occurs between PMT α and PMT β which does not coincide with an event on any muon paddle.

The analog signal from each PMT, after being passed through the linear fan out (see above), is discriminated using the same module that discriminates the signals from the muon paddles. The entire discrimination module can be inhibited by a digital input/output line on the NI-DAQ multifunction timing card (see section 5.1.3). This allows the DAQ software to start and stop acquisition of events with a sharp time

the threshold from producing multiple output pulses.

⁷If the rate $\Sigma\mu_i$ is constant in time, then false vetoing is proportional to the countrate $\alpha + \beta$. The overall effect on the data is thus a reduction in the detection efficiency. Any time dependence in the muon veto rate will result in a time varying detection efficiency and a systematic bias.

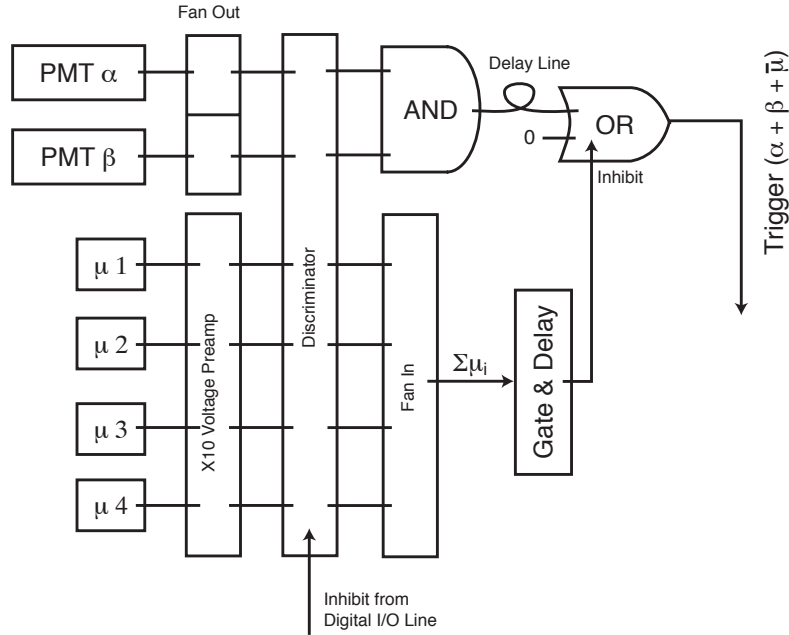


Figure 5.2: A block diagram of NIM logic modules used to produce the DAQ trigger is shown.

edge. The discriminator threshold for each of the PMT signals (α , β) is set just below the dip between the first photoelectron peak and the noise, so that the discriminator produces a pulse for all single photoelectron pulses and some small number of noise pulses. The output pulses from the PMT discriminators are each 20 ns wide and are delayed approximately 15 ns from the time the PMT voltage crosses the threshold. These outputs are input to an “AND” coincidence module⁸, which puts out a 194 ns wide pulse whenever a coincidence occurs within a window with an effective width of 40 ns. The coincidence output pulse passes through a 48 ns coax delay cable before being input to an “OR” coincidence module (one channel of an identical but separate LRS 622). The delay cable is used to match the delay of the gate and delay for the muon signal.

⁸LeCroy Research Systems. Chestnut Ridge, NY. <http://www.lecroy.com/lrs/>. Model 622 Quad Coincidence Module.

The function of the “OR” coincidence module is simply to allow vetoing of the $\alpha + \beta$ signal whenever a muon event occurs. The “OR” module’s inhibit control is connected to the lengthened, slightly delayed muon signal. As described above, the output of the muon discriminator channels are fanned in to produce one logical signal which is effectively the “OR” of all muon channels. In order to make the muon vetoing window wide enough to veto any coincident helium afterpulses which might be produced by a highly energetic muon passing through the cell, it is necessary to pass the muon logic signal through a gate and delay module⁹. The gate and delay module produces a pulse 308 ns wide and delayed by 52 ns. This delayed signal is used to inhibit the “OR” coincidence module as described above.

The final trigger signal is thus the output of the “OR” coincidence module. This module produces a NIM logic “TRUE” pulse whenever above-threshold PMT pulses from PMTs α and β have occurred within a 40 ns window, and at least 10 ns before or 100 ns after any event on a muon paddle. The output trigger signal is fanned out using a gate and delay. Simultaneous logic signals from the gate and delay are sent to the digitizer card and timing card as triggers. A delayed output is fed to a multichannel scaler for real-time visual monitoring and display of the countrate.

5.1.2 Waveform Digitization

Waveform digitization is triggered by the NIM logic circuit and is performed by two Signatec PDA500 waveform digitizers¹⁰. These 8-bit digitizers have an input bandwidth of 500 MHz and a digitization rate of up to 500 MS/s. In our earlier work with waveform digitization it was determined that identical pulse height spectra are

⁹Module produced in-house at NIST

¹⁰Signatec, Inc. Corona, CA. <http://www.signatec.com>




Figure 5.3: A typical PMT pulse is shown as digitized by the Signatec PDA500 master card. Each point is 4 ns.

obtained at digitization rates of 250 MS/s and 500 MS/s[42]. In order to minimize the amount of data acquired, the digitization cards are operated at 250 MS/s. Each card contains 4 megabytes of RAM which can be segmented so that many events can be recorded, each within its own small section of the card's memory. For each event, 64 samples are recorded, thus spanning a time window of 256 ns. Typical PMT pulses from neutron decay events are ~ 10 ns wide and large scintillation events from high-energy cosmic rays passing through the helium may be as much as 100 ns wide. A typical digitized pulse is shown in figure 5.3.

The signals from the two detection PMTs, α and β , are each recorded by a separate digitizer card, called the master and slave cards, respectively. The master card provides clock and trigger signals for the slave card so that the time bases of the two cards are completely synchronized. Although the two cards were supposedly

tested together by the manufacturer in this master/slave configuration, the slave card displayed several hardware problems which had to be compensated for in the data analysis software (see section 5.3).

Since the trigger produced by the NIM logic is delayed with respect to the analog signal input to the digitizers, the cards were operated in a pre-trigger mode. Each card is in fact digitizing the input signal constantly. Each sample produces one byte which is loaded into a circular buffer of fixed depth (in our case 80 samples, or 320 ns). When the card receives a trigger, it transfers the first 64 samples from the circular buffer to the next segment of the card's main memory. This has the effect of recording the digitized waveform for the time window from 320 to 64 ns before the trigger. Given the delays in the trigger logic, this places the main event pulse around 25% through the 64 sample window.

The analog to digital converter (ADC) within each waveform digitizer has a fixed input range of ± 270 mV. The input analog signal is attenuated and shifted before being fed to the ADC in order to keep it within the appropriate range. An attenuation of 16 dB was determined to allow clear discrimination of the single photoelectron peak from each PMT while also allowing large events (up to about 50 p.e.) to be detected. A large attenuation was also necessary to maintain reliable operation of the digitizing cards. When pulses with voltages (after attenuation and shift) outside the input range were input, the card would behave in an unexpected manner: the card would trigger constantly and record a digitized signal showing rail to rail noise at approximately the sampling frequency. We guaranteed that the input signal could not exceed the voltage which produced this effect by using a large input attenuation and a fanout with a maximum output voltage range¹¹ of ± 2.5 V. As shown in figure 5.3 the offset

¹¹With a 16 dB attenuation, a -2.5 V signal would be attenuated to -396 mV. Since the signal is offset almost to the positive rail, the maximum pulse amplitude

of the analog signal was chosen so that the background level is close to the maximum allowed value (255), therefore allowing the maximum range of negative-going pulse sizes.

5.1.3 Time Tagging

The trigger signal is also input to the time tagging circuit. Time tagging is performed by one of the counters on a National Instruments PCI-6025E card¹². The 24-bit counter operates as a counter/timer with a 1 kHz clock input provided by an SRS DS345 function generator¹³ with a specified frequency accuracy of 5 ppm. The NIM logic trigger is provided on the gate input of the counter/timer. When a pulse is input to the gate, the counter/timer latches its current value into an on-board register, generates an interrupt, and continues to count. The data acquisition software (see section 5.2) reads out the clock register upon receiving the interrupt. If a second event occurs before the register can be read and cleared, then an error bit is set by counter and recorded when the counter is read out. In the case where this error bit is set, one or more time tags have been missed. In order to know exactly how many events have occurred, the trigger signal is also used as the clock signal for the second counter on the same PCI-6025E card. This counter then acts as an event counter, counting the total number of triggers.

which would remain above the negative rail is almost 540 mV. Therefore the fanout cannot generate a negative pulse large enough to exceed the ADC's input range.

¹²National Instruments. Austin, TX. <http://www.ni.com>

¹³Stanford Research Systems. Sunnyvale, CA. <http://www.srsys.com>

5.1.4 External Control

In addition to recording input signals, the data acquisition system also controls several external devices. Through digital input/output lines in the National Instruments PCI-6025E card, digital control signals are provided to: the power supply for the detection PMTs (so that it can be turned on or off), the beam blocker (blocking or unblocking the beam), and the slider (covering the entrance to the apparatus with a pile of lead or moving the lead out of the way in favor of a flight tube). One of the digital I/O lines can also be used to inhibit the discriminator module in the NIM logic circuit, hence providing control over the enabling or disabling of data acquisition. The magnet power supply is controlled over GPIB by a National Instruments NI-GPIB PCI to GPIB interface card, allowing the DAQ to control and monitor the output current and voltage.

5.2 Data Acquisition Software

The data acquisition software controls the entire data acquisition process. The DAQ software is written in C++ and C and runs under the MacOS 9 operating system.

5.2.1 Experimental Sequence

In each run, the data acquisition software performs a sequence of actions (see figure 5.4). Two types of runs are taken: “positive” and “negative” runs. Positive runs are runs where the magnetic trapping field is maintained at its maximum value throughout the run. Hence ultracold neutrons which are produced by the superthermal process, in the appropriate spin state, while the beam is on are trapped. The




Figure 5.4: The state of various devices as a function of time is shown during standard positive and negative runs.

number of trapped neutrons builds up until the beam is turned off and the observation of decays begins. In negative runs, the magnet current is kept low while the beam is on, so that ultracold neutrons which are produced are not accumulated in the trap. By observing the signal in a negative run, the backgrounds in the experiment are measured. Subtraction of negative from positive runs then should yield the signal due to trapped neutrons.

The use of a background subtraction technique is necessitated by the large backgrounds in the system. Comparisons have been made between the alternative running strategies of background correction and fitting “positive” data to a known signal plus background model [134]. Even in the case where the time dependent background is known, *a priori*, to be a single exponential decay with a lifetime of 213.56 s (the decay lifetime of neutron-activated aluminum), it is found that background correction is preferable for any non-zero amplitude aluminum signal. Those comparisons were

motivated by our previous results, where the observed time-dependent background fit well to a single exponential decay (aluminum) [42]. In the current work, however, no model for the time dependence of backgrounds in our experiment is known *a priori*. Hence it is impossible to fit the “positive” runs to a complicated multi-parameter function of signal plus background. Background subtraction is the only viable strategy. The optimal times for filling the trap and observing decays has been studied for both operational strategies: background correction [134], and fitting the signal plus background to a single model [135, 136]. To optimize the sensitivity to the neutron lifetime signal, we chose to load the trap for 2100 s and observe decay events for 2700 s. These values are near the calculated optimum values when using a background correction strategy with 3400 trapped neutron expected per load, and our known fluorine decay signal and constant backgrounds (see section 6.2).

At the start of the run, the photomultiplier tubes are turned on and the magnet is set to its initial current (170 A for a positive run, 5 A for a negative run)¹⁴. Data is taken for 100 s before the beam is turned on. This allows observation of the background rate before the run, which serves as an important diagnostic. The compensation coils for the photomultipliers (see section 4.4.2) are always kept at their operating current of 8 A. When the magnet current is low, the “compensation” of these coils has the opposite effect and decreases the efficiency of the PMTs to zero.

¹⁴The magnet current is set to 5 A rather than 0 A for a negative run because of the quench protection circuit (see section 4.3.2). If the output voltage of the power supply drops below the voltage drop of the “Run SCR” (1.2 V), all current is shunted to the dump resistor. A power supply providing 5 V across the “Run SCR” must then be manually switched in to reactivate the quench protection and allow normal flow of current through the magnet. In order to avoid the necessity of flipping this switch at the appropriate time in each negative run, the current is always maintained just high enough to keep the “Run SCR” in normal conduction. Only 0.5% of the UCN trapped with the magnet at 170 A are trapped with the magnet at 5 A.




Figure 5.5: The countrate is shown over the first 100 s of the run for both positive and negative runs. The data shown are the average of 37 positive and 38 negative runs with $T_{cell} < 250$ mK. The data is cut to require at least 2 photoelectrons per PMT. The countrate drops to zero in the negative runs due to the magnetic field around the PMTs.

In addition, the magnet current is high at the end of each run. It takes about 200 s for the magnet to ramp down. Thus the countrate over the first 100 s of a negative run drops with a characteristic time profile (see figure 5.5). The countrate at the beginning of a positive run is constant.

At $t = 100$ s, the detection PMTs are turned off and data acquisition is inhibited. The computer moves the slider so that the flight tube replaces the lead in the neutron path. The computer signals by a series of beeps for the operator to open the main beam shutter, which must be operated manually for safety reasons. Opening of the beam shutter is always complete by $t = 110$ s, at which time the computer moves the beam blocker out of the neutron beam. This defines the time at which neutrons are

allowed into the apparatus ensures that it occurs at the same relative time in each run.

While the beam remains on, the DAQ simply waits. At $t = 2100$ s the magnet current is set to 170 A. In a positive run this has no effect since the current was already 170 A, but in a negative run the magnet starts to ramp up. It takes the magnet 95 s to ramp to full current. This occurs just before the beam is turned off.

As when the beam is turned on, the beam is turned off in such a way that the true beam off time is under the control of the computer, not the operator. At $t = 2195$ s, the DAQ signals the beam blocker to block the beam. At $t = 2200$ s the slider moves the lead back into the beam path and a series of beeps instructs the operator to close the beam shutter. By the time the operator has a chance to close the shutter, the beam blocker has completely closed, so the actual beam off time is consistent from run to run. At $t = 2200$ s the PMTs are turned back on. With the beam now off, the countrate on the PMTs is not high enough to damage them, but the initial countrate is still high. These high background countrates render the initial seconds of data less useful (due to the statistical and systematic uncertainties of subtracting two large numbers \sim kHz). In order to minimize the acquisition of less useful data, we then wait until $t = 2225$ s to re-enable data acquisition. Data is then acquired until the end of the run at $t = 4900$ s.

5.2.2 Data Acquisition Loop

While data acquisition is enabled, the DAQ software asynchronously performs several tasks. Nothing in the software requires “real-time” operation. The software was written to run on the MacOS 9 operating system which implements “cooperative multi-tasking” (wherein applications must call the system’s task manager when they

are willing to allow another operation to take CPU cycles) but not true preemptive multi-tasking. The MacOS Finder (file manager), which normally is always running, was forced to quit and the data acquisition software was the only program running during data acquisition. This method of operation allowed the program to have routine control of the CPU so that all of its tasks were performed within the necessary time windows.

The DAQ software routinely checks the event counter to determine how many triggers occurred. Since the Signatec cards have a memory of 4 Mb and use 64 bytes to store each event they can store up to 65536 events before they must be read out. The DAQ also maintains a buffer for storing timestamps which holds a similar number of events. When the computer notices that the event counter has exceeded 62000 events, it pauses data acquisition to clear all buffers¹⁵. The NIM logic circuit is inhibited so that no additional triggers can be received. The memory of the two Signatec cards is read into system RAM over the computer's PCI bus. The PCI standard specifies a maximum transfer rate of 100 megabytes/s. In practice it takes about 500 ms to read out both cards. The time stamp buffer is also copied to a separate memory location. Once these buffers are cleared and the signatec cards are reset, the NIM logic circuit is de-inhibited and data acquisition resumes. In practice a typical 4900 s run requires clearing the buffers twice (once during the run and once upon its completion).

Whenever data acquisition is enabled, an interrupt handler is prepared to deal with incoming timestamps. When a timestamp is latched by a gate signal on the

¹⁵We pause to clear the buffers at 62000 events rather than 65536 to allow ourselves leeway. If the computer doesn't immediately notice that 62000 has been exceeded, or if it is interrupted before inhibiting acquisition then there is a significant margin of error before the buffers in fact overflow. In practice, the buffers almost always contain < 62010 events when read out.

counter/timer an interrupt is generated. The interrupt handler services the card's interrupt by transferring the 24-bit timestamp and the 1 error byte into a timestamp buffer in system RAM. The counter/timer continues to count clock pulses throughout. If another gate pulse is detected while the first interrupt is still being serviced then an error flag is set that will be read out with the next event. Approximately 3% of the timestamps contain this error flag. The data analysis software corrects for these "missing" timestamps.

5.3 Data Analysis Software

At the end of each run, the data acquisition software writes the data stream to disk. Each data file is then analyzed using the data analysis software. The data analysis software reads in a stream of blocks containing tagged information. Blocks may contain information describing the settings of the DAQ software, timestamps, or digitized waveform data. Information blocks are used to set some of the software's parameters. Timestamps are simply read into memory. Signatec data is read into memory and "parsed" (see below) to produce a list of event areas (the integrated area of the PMT voltage waveform). The timestamps and areas (from master and slave cards) are then matched up and compared to a set of criteria by the "cutter". The time and areas of events which pass all of the data cuts are then output for further analysis.

5.3.1 The Parser

The parser is the core of the data analysis software and is responsible for converting digitized waveforms into pulse areas. The goal of this conversion is to replicate the energy resolution of a charge-coupled preamplifier.




Figure 5.6: A typical pulse height spectrum of the Burle 8854 is shown. The spectrum was produced using a charge-coupled preamplifier (Canberra Model 2005 Scintillation Preamplifier), shaping amplifier ($1\ \mu\text{s}$ shaping time), and MCA. The PMT was observing backgrounds in the cell.

In applications where pulse energy resolution is more important than timing, photomultiplier output is typically passed through a charge-coupled preamplifier, a shaping amplifier, and multi-channel analyzer (MCA). This set of electronics converts the PMT's anode charge to a channel in the MCA and yields optimal pulse height spectra (see figure 5.6 for a typical pulse height spectrum from a Burle model 8854 PMT produced in this manner). This method of producing pulse height spectra, when combined with a PMT with low gain dispersion, allows clear discrimination of the one, two and sometimes higher photoelectron peaks.

The purpose of the parser is to integrate the pulses digitized by the signatec card. Integrating this voltage waveform should yield a value proportional to the anode charge. If the parser functions properly, a high resolution pulse height spectrum

should be obtained which allows good discrimination between 1 and 2 photoelectron events, and thus allow separation of single-photoelectron backgrounds (such as luminescence) from multi-photoelectron signals.

A variety of approaches could be used to perform this integration. The simplest algorithm would simply be to sum the values of all 64 points within each event. In practice this yields a pulse height spectrum where the photoelectron peaks are smeared together. As will be discussed below, part of the reason for this is that the background voltage (when a pulse is not occurring) changes throughout the data taking. Even if this is corrected for (in the manner described below) and an average background is subtracted from each point, the pulse height spectra produced using this algorithm are unsatisfactory (see figure 5.7).

The parser produces much cleaner pulse height spectra if it integrates each pulse within an event independently. Since an event is 64 samples long and PMT pulses may be as narrow as 10 ns, there may be several pulses in any given event. It has been determined empirically that the best pulse height spectra are generated when each event is assigned only the area of its “main” pulse (that is the pulse which causes the NIM logic to trigger).

Before describing the details of how each pulse is identified and integrated it is important to understand the general algorithm and the method of selecting several essential input parameters. The general function of the parser is exposed in figure 5.8. The parser locates the points `startcross` and `finishcross` which are the first and last points, respectively, in the pulse which go below the threshold value `thresh`. The area A is computed as

$$A = \sum_{i=\text{startcross}}^{\text{finishcross}} (V_i - \text{backgroundave}) \quad (5.1)$$

That is the digitized waveform is summed between `startcross` and `finishcross` and




Figure 5.7: An unsatisfactory pulse height spectrum of the Burle 8854. The area is computed by summing all 64 points in each event. An average background (computed by averaging the last 16 points of the last 500 events) is subtracted from each point before summing. This pulse height spectrum was produced by histogramming the area of all events digitized by the master card throughout the 37 positive runs in the Cold1 data set (see section 6.1).

the average background value `backgroundave` is subtracted from each point. This process is repeated for all of the pulses in the event. The pulse which contains the point `trigpos` between `startcross` and `finishcross` is considered to be the main event and it is the area of this pulse that is used in cutting the data and generating pulse height spectra.

The correct operation of the parser depends crucially upon the proper selection of the parameters `trigpos`, `backgroundave`, and `thresh`. Initial attempts were made to specify these parameters manually, but it was determined that at least in the case of `trigpos` and `backgroundave`, they can be determined automatically by the parser from looking at the data itself. One technique which aids in determining




Figure 5.8: A typical digitized pulse is shown. The parser calculates the area of the pulse by summing the value of the points in the waveform between **startcross** and **finishcross**. These are the points at which the waveform crosses the threshold value **thresh**. The average background value **backgroundave** is subtracted from each point.

these parameters is calculating an “average event” by averaging all events in a given data block (1 card memory) point by point. A typical average event is shown in figure 5.9. In the average event, one sees a clear pulse shape corresponding to a multi-photoelectron event. From the background at the beginning of the event, the pulse falls monotonically to a (negative) peak before rising back to the background. The average value of **trigpos** can easily be ascertained simply by finding the position within the event at which the minimum value occurs.

Calculation of the average event can also allow detection and correction of two intermittent hardware problems with the slave digitizer card. The value of **trigpos** calculated from the average event is 16 ± 2 every time the master digitizer is read out




Figure 5.9: A typical “average event” is shown. This event is calculated by averaging, point by point, the 62001 events in one data block from the master digitizing card. Data shown is from the first master readout of file 080901008.

and the first time in each run when the slave card is read out. Occasionally when the slave card is read out it will skip the first 4 bytes in the card memory. Since all analysis of the data is performed relative to the first byte returned by the card, the 4 bytes missing from the first event has the effect of shifting 4 bytes from the beginning of each event into the end of the previous event. This effect can also be seen as a shift in the location of `trigpos` as calculated from the minimum position in the average event. Hence data blocks where `trigpos` is calculated to be ≤ 14 are padded with 4 bytes at the beginning of the data block to compensate for this hardware error.

Byte Ordering Problem

The other intermittent hardware problem with the slave digitizer card can be seen clearly in the average event shown in figure 5.10. The average event shown with the




Figure 5.10: An average event from the slave card is shown (solid line). This event has been distorted during readout of the slave card. The corrected event is also shown (dashed line). Data shown is from the second slave readout of file 080901008. Two rotations were applied to correct the data.

solid line is calculated from the raw data returned during card readout. The shape of the pulse is very jagged and not at all like what we see in the master card average events (see figure 5.9). This is due to a misordering of the bytes by the slave card. A flaw in the multiplexer/de-multiplexer unit within the slave card, which converts 4 8-bit samples into a single 32-bit word and back again, is likely responsible for the misordering. The misordering consists of one or more rotations or one or more shifted rotations. These effects are shown in figure 5.11. The basic unit of data correction is a rotation of byte order, where a sequence (1 2 3 4) is reordered to (4 1 2 3). More than one application of this rotation may be necessary to correct the data. Although in most instances the regular byte order is restored after one or more rotations starting at the first byte in the data file, occasionally it is necessary to perform one or more




Figure 5.11: Operations for correcting mis-ordered slave card bytes are shown. All operations are performed on the data in 4 byte sections. The data may be corrected by one or more rotations. The rotations may begin with the first byte in the file, or the second (“shifted rotation”).

“shifted rotations” where the same rotation function is applied but starting on the second byte in the data file rather than the first. The parser routine tries all 7 possible orderings of the data (0–3 rotations and 0–3 shifted rotations). The parser determines that the byte order is correct when the average pulse monotonically decreases from the first point 10% of the way between the background and the minimum value until reaching the minimum value. Whatever function was necessary to correct the “average event” data is then applied to the entire data block before further analysis.

Determining the Background Level

The algorithm used for determining the parameter `backgroundave` has a strong effect on the results of the data analysis. From the average event shown in figure 5.9, we can see that the last 16 points in the average event are at the background level. The first few points in the data file also are background, but since the position of the

main peak varies from event to event it is harder to determine how many points at the beginning of each event may be considered background. Averaging points which are part of a peak into the background will lower its value and lead to spurious area calculations. A simple algorithm for determining **backgroundave** would then be to simply average the last 16 points in the average event. This would yield a value of **backgroundave** which could be used for every event in the data block.

Unfortunately this algorithm is unsatisfactory, as it results in a non-physical jump up in the calculated countrate at the boundary between the first and second data blocks. This effect is seen in data which has been cut with a non-zero requirement for minimum pulse area. Because **backgroundave** is miscalculated for the first data block, as will be shown below, the area assigned to events near the end of the block is too small. Thus fewer events from the first block pass the cut and the apparent countrate is lower in that block than it should be.

In order to analyze the data properly, **backgroundave** must be computed independently for each event rather than for entire data blocks. This is because the background level (baseline PMT anode voltage) is changing with time (see figure 5.12). Here **backgroundave** is computed by averaging the last 16 points in each event. These individual values of **backgroundave** are smoothed in a boxcar average over 500 events. The value of **backgroundave** is initially much lower after the beam has been on than it is before but then settles down to a steady level. It was the effect of averaging in this time period of lower **backgroundave** in the first data block that lead to the discontinuity in countrates between the first and second data block. With **backgroundave** computed for each event by averaging the last 16 points of the 500 previous events, this discontinuity is eliminated.

The fact that **backgroundave** changes with time after the beam has been on is




Figure 5.12: The time dependence of the **backgroundave** parameter is shown over an entire run. Two readouts of the master digitizer card are shown. Here **backgroundave** is computed by averaging the last 16 points in each event. These individual values of **backgroundave** are smoothed in a boxcar average over 500 events. The time gap where the beam is on is indicated. For the run shown, **backgroundave** reaches a steady level around event number 15400, which occurs at $t = 2320$ s. Data shown are from the first and second readouts of the master digitizer in file 080901008.

evidence that the PMTs must “settle down” when they are first turned back on. This change in **backgroundave** may be caused either by the exposure of the PMT (while turned off) to very large light levels while the beam is on or by the voltage divider heating back up after having 2100 s to cool down while the PMTs were turned off. Although this shift in baseline is accounted for by the method of computing **backgroundave**, more troubling is the possibility that the gain of the PMTs changes with time. An important improvement in the data acquisition system in the future would be the incorporation of gain compensation. This could be accomplished fairly easily by regularly pulsing an LED with the same pulse shape. By comparing the

shape of the digitized test pulses, shifts in the gain could be detected and corrected for in data analysis.

Parsing Threshold

The final parameter that governs the parser's action is the parsing threshold, **thresh**. As shown in figure 5.8 this parameter describes the threshold for what is considered a pulse. Unlike **backgroundave**, which is an average of many points and hence can have non-integer values, **thresh** can only take integer values from 0–255. In practice the optimal threshold was determined by manually varying the threshold. If **thresh** is set too low then a large number of events are parsed in which no pulse is found at **trigpos**. Similarly if **thresh** is set too high, then a large number of events find a pulse before the main pulse. For all of the data analyzed below, **thresh** was set to the highest value for which the number of events with a pulse before the main pulse is small.

Once all of the parameters are set, the integration takes place from **startcross** to **finishcross**. The details of setting **startcross** and **finishcross** and performing the integration are described fully in Appendix E. A listing of the C code for the parsing routine is also found therein. The parser assigns a main pulse area to each event recorded by the master or slave card.

5.3.2 Cutting the Data

Once the parser has integrated all of the events, it is necessary to correlate the timestamp with the area data and perform any desired cuts. The number of events read out from the master and slave cards is equal to the total number of events counted by the event counter, N_{events} . This is in all cases greater than the total number of

timestamps which are recorded, $N_{timestamps}$. In order to match times with areas it is necessary to account for the missing timestamps. Some number of timestamps, N_{error} (typically 2 – 3% of N_{events}), are marked with an error flag which indicates that another timestamp was received while that timestamp was being read out. Since $N_{error} < N_{events} - N_{timestamps}$ some of these error flags correspond to more than one missed timestamps. It is assumed that in no more than 2 timestamps are missed while the counter/timer is being read. The fraction of the error flags which correspond to two missing timestamps is

$$F_{2missing} = \frac{N_{events} - N_{timestamps} - N_{error}}{N_{error}} \sim 0.1 \quad (5.2)$$

The total number of timestamps in the data file is then increased to N_{events} by adding one timestamp after each event marked with the missed timestamp error flag. A second timestamp is randomly added with probability $F_{2missing}$ after these events. Since timestamps are only missed when multiple timestamps come in rapid succession, and since the readout time for the counter/timer is on the order of a few milliseconds, it is accurate enough for our purpose to assign to these added timestamps the same time as the timestamp with the error flag.

Once we have a time assigned to each event, we can perform cuts on the data. We typically want to restrict the data to events with areas within a given band. It is convenient to think of the desired area thresholds in terms of numbers of photoelectrons. Since we particularly desire to separate the largely single photoelectron luminescence background from the signal, the determination of photoelectron thresholds is important. A typical pulseheight spectrum generated by histogramming the areas output by the parser is shown in figure 5.13. Thresholds (in area units) for 1, 2, and 3 photoelectrons are determined visually by inspecting the pulse height spectrum. For a higher n photoelectron thresholds (Th_n) we can calculate the threshold




Figure 5.13: A typical pulse height spectrum is shown for data analyzed by the parser. The 1, 2, and 3 photoelectron thresholds are shown. This pulse height spectrum was produced by histogramming the main pulse area of all events digitized by the master card throughout the 37 positive runs in the Cold1 data set (see section 6.1). Events where no main pulse is detected are assigned an area of 0.

from the measured thresholds Th_2 and Th_3 as:

$$Th_n = Th_3 + (n - 3) \times (Th_3 - Th_2) \quad (5.3)$$

Higher photoelectron thresholds are calculated from Th_2 and Th_3 without regard for Th_1 since the optimal value of Th_1 may be affected both by the gain of the PMT and the parsing threshold, while the other thresholds should only depend on the PMT.

5.4 Data Analysis Procedure

The data taken in this experiment are analyzed in sets (see section 6.1). Each set contains n_+ positive and n_- negative runs. Using the data analysis software described

above, all of the runs in a given set are parsed and cut using identical parameters and thresholds. The result of this procedure is a data file for each run which contains a monotonically increasing sequence of event times which have passed the data cuts. The data file also contains the time and location of any dead time due to card readout. This data are then read into Igor Pro¹⁶ for further analysis.

The events in each run are binned to construct a time histogram. Let us consider bins of arbitrary width t_{bin} where the number of counts in the i th bin of the j th run of type $p = +$ or $-$ is $x_{i,j,p}$. We also construct for the j th run a weighting vector $w_{i,j,p}$ which allows us to perform deadtime correction. We set $w_{i,j,p} = 1$ unless a time bin contains deadtime, in which case $w_{i,j,p}$ is simply set to the fractional live time of the bin.

Data analysis is performed on summed, binned data. All of the runs are summed to calculate:

$$x_{i,p} = \sum_{j=1}^{n_p} x_{i,j,p}. \quad (5.4)$$

The weighting vectors are also summed over j to form $w_{i,p}$ and can be used to calculate the deadtime-corrected countrate $r_{i,p}$ [Hz] for each bin:

$$r_{i,p} = \frac{x_{i,p}}{w_{i,p} t_{bin}} \quad (5.5)$$

The statistical error on the measurement of $r_{i,p}$ is simply

$$\Delta r_{i,p} = \frac{\sqrt{x_{i,p}}}{w_{i,p} t_{bin}} \quad (5.6)$$

While valuable information about the nature of backgrounds can be obtained from analysis of the summed, binned positive and negative data alone (see section 6.2.2),

¹⁶Wavemetrics, Inc. Lake Oswego, OR. <http://www.wavemetrics.com>. Igor Pro 4.03 is a commercial data analysis package for Macintosh and PC. It contains a standard suite of graphing, curve fitting, and data manipulation routines. Data analysis can be automated using a proprietary, C-like macro language.

information about trapped neutrons is obtained by subtracting positive and negative runs.

The difference signal for the i th bin, d_i , is computed from the difference of the summed, binned, weighted positive and negative runs as:

$$d_i = r_{i,+} - r_{i,-} \quad (5.7)$$

The statistical error on the estimated difference for the i th bin is:

$$\Delta d_i = \sqrt{(\Delta r_{i,+})^2 + (\Delta r_{i,-})^2} \quad (5.8)$$

Weighted, non-linear least squares curve fits are then performed on the data vector d_i with the weighting $1/\sigma = 1/\Delta d_i$. Curve fits are performed to the function $y = y_0 + Ae^{-kt}$. The lifetime and its error are estimated as

$$\tau_{-\Delta-}^{+\Delta+} = \frac{1 + (\frac{1}{k-\Delta k} - \frac{1}{k})}{k - (\frac{1}{k} - \frac{1}{k+\Delta k})} \quad (5.9)$$

where Δk is the estimated standard deviation in the parameter k .

Chapter 6

Results and Discussion

The neutron trapping apparatus (as described in Chapter 4) was operated on several occasions in 2001 for the purpose of demonstrating magnetic trapping of UCN. A monochromatic beam of neutrons was used to measure the neutron lifetime to as high an accuracy as possible. Upon each cooldown except the final one, the excessive heat loads described in section 4.2.4 prevented the cell from cooling to a temperature low enough that the UCN upscattering lifetime is long enough to permit observation of magnetic trapping.

This chapter contains the results from the last of the above-mentioned cooldowns of the apparatus, consisting of 18 days of data taking. The different experimental conditions under which data were taken are described in section 6.1. The apparatus was run both with ultrapure ^4He (“trapping runs”) and with natural abundance helium (^3He runs”). Trapping runs were performed for a range of cell temperatures. In the ^3He runs, ^3He is present at sufficient concentration to absorb any trapped UCN with a time constant of order 1 s, however less than 1% of the cold neutron beam is absorbed by the ^3He .

An analysis of the observed backgrounds, both constant and time-varying, is found

in section 6.2. We can conclude from this analysis that for all detection thresholds, the majority of the events producing the constant background originate in the helium within the cell, rather than the acrylic lightguides. A large time-dependent background is also observed at all thresholds. At single photoelectron thresholds the largest component is likely due to neutron-induced luminescence (see section 2.4).

The presence of time-dependent backgrounds necessitates taking positive and negative runs (as described in section 5.2.1). A difference in the countrate versus time between positive (trapped UCN + backgrounds) and negative (backgrounds only) runs may be evidence either for magnetic trapping of UCN in the positive runs, or for imperfect subtraction of the backgrounds. If the backgrounds are not identical in the positive and negative runs then subtracting them will leave a residual difference which may mimic a trapping signal. The measurements made with natural abundance helium are used to determine if a putative trapping signal is caused by imperfect background subtraction. A true trapping signal seen in the ultrapure ^4He should disappear from otherwise similar runs taken with significant amounts of ^3He present. Any signal due to imperfect background subtraction should be observed in both ultrapure and ^3He runs.

A measurement of the detection efficiency of the apparatus was performed in a separate cooldown of the apparatus (see section 6.6). A ^{113}Sn beta line source placed in the center of the detection region was used to calibrate the detector's efficiency. Since the observed detection efficiency is low, one would expect that the number of observed neutron decays will depend strongly on the threshold.

An analysis of the data requiring detection of at least one photoelectron per PMT is found in section 6.3. A difference signal is observed between positive and negative trapping runs. Comparison with ^3He runs shows that this signal is due to imperfect

background subtraction. Subtraction of the ^3He countrate from the trapping countrate should subtract the signal due to imperfect background subtraction, leaving only the trapping signal. The resulting signal does appear consistent with trapped neutrons, however the statistical uncertainties inherent in the multiple subtractions make this result unconvincing.

Demonstration of magnetic trapping can be clearly established by analyzing the data using higher thresholds. When two photoelectrons are required per PMT, a trapping signal is seen in the trapping runs but not in the ^3He runs (see section 6.3). The trapping signal varies in amplitude and lifetime with threshold (see section 6.4) and temperature (see section 6.5). The observed lifetime for trapping runs in cold, ultrapure helium is shorter than expected. Further measurements will be necessary to determine the cause of this apparent discrepancy.

From the observed neutron decay rates and the independently measured detection efficiency, one can extract the effective UCN production rate within the apparatus (see section 6.7). The measured UCN production rate is 0.4 to 1.0 of the production rate predicted by the Monte Carlo used for design of the monochromator and upgraded apparatus, depending upon the value assumed for the trap lifetime.

6.1 Data Sets

Sets of positive and negative runs were collected under a variety of experimental conditions. Trapping runs, with ultrapure ^4He in the cell, were taken in several different temperature ranges: < 250 mK, 300–350 mK, 400 mK, 500 mK, 750 mK, and > 1 K. A set of ^3He runs were also taken, with the cell at temperatures < 250 mK

Table 6.1: For each dataset, the running conditions (type of helium used and temperature of the cell, T_{cell}), and the number of positive and negative runs are shown. Note that some additional runs were taken but excluded from the analysis because of problems in the data files.

Name	Helium	T_{cell} [mK]	# Positive	# Negative
Cold1	ultrapure	< 250	37	38
325 mK	ultrapure	300–350	13	13
400 mK	ultrapure	400–470	16	15
500 mK	ultrapure	500–625	28	26
750 mK	ultrapure	750 ± 1	18	16
Warm	ultrapure	> 1000	22	20
^3He	natural	< 250	15	15
Cold2	ultrapure	< 625	94	92

but full of natural isotopic abundance helium ($^3\text{He}/^4\text{He}$ approximately 10^{-7})¹ rather than ultrapure ^4He . The names, characteristics, and numbers of runs in each of these data sets are shown in table 6.1.

Within all data sets, positive and negative runs were alternated². Certain systematics, such as activation of an isotope with a decay lifetime on the order of several times the length of an experimental run, can be minimized through this procedure. While this pattern was generally adhered to, it was occasionally changed for various practical reasons. For example, helium transfers were performed while the beam was on and the magnet off during negative runs, in order to maximize running time.

¹The isotopic ratio of $^3\text{He}/^4\text{He}$ for helium which originates from natural gas sources, as does all commercially available bottled helium gas, is $1 - 2 \times 10^{-7}$ [55].

²Runs were taken following a pattern such as “+ – – + – + + –” and so on.

Hence, the need to transfer helium sometimes lead to altering the running pattern. Furthermore, as will be seen below, runs from different datasets were often interspersed. So, while a general alternation of positive and negative runs was followed, no particular pattern was strictly adhered to for every dataset.

Data in the “Cold1” data set were taken with the cell cooled to the minimum achievable temperature. This temperature varied somewhat from run to run but was in all cases less than 250 mK. A total of 40 positive and 42 negative runs were taken in this temperature range. The data were taken in three contiguous blocks of time for a total of 135 hours of running time. After the first block of 39 hours, in which 12 positive and 11 negative runs were taken, the dilution refrigerator was intentionally warmed to approximately 0.5 K for 30 hours to take data which are included in the 400 mK and 500 mK data sets. The refrigerator was then cooled back down and a second block of cold data (21 positives, 23 negatives) was taken for 72 hours. Within this block of data are 5 positive and 6 negative runs where the signal was observed for 5400 s after the beam was turned off, rather than the usual 2700 s.³ After the second block of data, the dilution refrigerator could no longer provide sufficient cooling to keep the cell temperature below 250 mK (due to a plug in one of the refrigerator’s condensing lines)⁴. Thus warmer data was taken for approximately 7 days. After opening the second condensing line on the dilution refrigerator, it was again possible

³This subset of the data was taken to see if at long times the background was constant. Poor statistics on this data subset made it impossible to draw meaningful conclusions from the long time parts. The data up to the usual 2700 s after the beam, is included in the main Cold1 dataset.

⁴Our dilution refrigerator is equipped with two parallel condensing lines. Each line has its own liquid helium cold trap and condensing impedance. Normally the refrigerator is operated with only one condensing line open at a time. We have found that the first condensing line often plugs after several days of running. This plug probably forms due to the high flow rates experienced during the initial cooling of the fridge and the condensation of the helium into the cell.

to cool the cell to less than 250 mK and a third block (24 hours running time) of Cold1 data (7 positives, 8 negatives) was taken. In the analysis of data from the Cold1 data set, 3 positive runs and 4 negative runs are excluded from consideration⁵.

The data in the “325 mK” dataset were taken at cell temperatures between 300 and 350 mK. The data were taken in two blocks. In the first block (20 hours running time) the refrigerator was slowly warming up as its cooling power decreased (the first condensing line was plugging). The 7 positives and 7 negatives in this first block range in cell temperature from 310–355 mK. The second block of data (17 hours running time), taken 8 days later, was taken while the refrigerator was capable of cooling to less than 250 mK (the second condensing line was in use). In this case the cell temperature was stabilized using a temperature controller⁶. The resistance of the thermometer on the beam entrance side of the cell was stabilized to $3400 \pm 2 \Omega$, corresponding to a temperature of 330 mK. While the cell temperature was stabilized at 330 mK the mixing chamber temperature remained stable at 296 mK. Under these running conditions, 6 positive and 6 negative runs were taken. All 13 positive and all 13 negative runs in the data set were considered in the analysis.

⁵Two of positive runs and one of the negative runs are excluded because the byte ordering problem (see section 5.3.1) in one of the waveform digitizers could not be corrected for those runs. One positive and one negative are excluded because the nature of the run as recorded in the lab notebook was inconsistent with the observed behavior of the countrate during the 100 s of observation before the beam is turned on (see figure 5.5). Finally, two negative runs were excluded due to large amounts of noise in the data (large non-statistical spikes in the counting rate for short periods of time). Smaller amounts of noise in other data runs were corrected by eliminating narrow time windows of the data file and marking those sections of the data dead for purposes of deadtime correction.

⁶Temperature control was maintained using a TDS530 temperature controller [Company Info?]. The temperature controller uses a PID feedback loop to vary the power output to a resistor on the detector side of the cell in order to maintain tight control on the measured resistance of a thermometer on the beam entrance side of the cell.

Runs included in the “400 mK” dataset were taken at cell temperatures from 400–470 mK. Of these runs, 3 positives and 2 negatives were taken at various temperatures in this range as the dilution refrigerator temperature slowly drifted up. The remainder of the runs were taken as a contiguous block using the temperature controller to stabilize the cell temperature at 400 mK. The set contains a total of 17 positive and 16 negative runs collected in 48 hours of running time. One positive and one negative run had to be excluded from the data set during analysis⁷.

All runs with cell temperatures in the range 500–625 mK are included in the “500 mK” data set. The data set was taken in several time blocks. Three short blocks of running (6.5 hours, 8 hours, and 10.5 hours) were separated from each other by interludes of a few hours⁸. The majority of the running time in this set occurred during one contiguous 60 hour block. A total of 28 positive and 29 negative runs were completed. Three of the negative runs were excluded from the data analysis⁹.

The data in the 750 mK dataset were taken with the cell temperature stabilized using the temperature controller. The mixing chamber temperature also remained stable at 733 ± 2 mK throughout this data set. A total of 18 positive and 18 negative runs were taken in one contiguous block of 51 hours of running time. One of the negative data files was lost and another was excluded¹⁰.

The “warm” data set comprises 22 positive and 22 negative runs. The data

⁷The positive run and the negative run were excluded due to inability to correct the byte ordering problem, and large noise, respectively.

⁸At this point temperature control was maintained by manually adjusting the power output to a heater on the mixing chamber. Hence relatively poor control was obtained and the cell temperature slowly drifted. Hence runs in these three blocks of the “500 mK” set were interspersed with runs in the “400 mK” set.

⁹In two cases runs were excluded when the magnet was ramped up late after a helium transfer. The third negative run is excluded due to an abnormal time profile.

¹⁰This run was excluded because the magnet ramped up late.

represent 61 hours of running time and were taken in one contiguous block of time. Two negative runs are excluded from consideration¹¹. The warm dataset was taken after the first condensing line of the dilution refrigerator plugged sufficiently that the mixing chamber temperature could not be kept below 600 mK (for taking data in the 500 mK data set). The warm data was then taken by operating the refrigerator effectively as a pumped ^4He system¹². For all of the runs in the warm dataset the cell temperature was in the range 1.0–1.2 K.

After data were collected at all of the above described temperatures, the ultrapure ^4He was removed from the cell and replaced with natural abundance helium. The cell and mixing chamber were heated while a diaphragm pump removed the helium gas from the buffer cell bypass line, placing it into the ultrapure helium dumps (see section 4.2.2). Once 99% of the ultrapure helium was pumped out of the cell, the condensing procedure was repeated using natural abundance helium from a high pressure gas cylinder was condensed into the cell. With the cell temperature below 225 mK, a total of 15 pairs of positive and negative runs for the ^3He dataset were taken in a single contiguous block of running time (42 hours). All of the runs were considered valid for use in the data analysis.

As will be described below, it may be appropriate to combine data from different operating temperatures. In this manner the “Cold2” dataset is the combination of the Cold1, 325 mK, 400 mK, and 500 mK datasets.

¹¹In one case the magnet ramped up late and in another the byte ordering problem could not be corrected.

¹²No power was applied to the still, mixing chamber, or cell heaters. The still was pumped on with a rotary pump only, not a turbo pump.

6.2 Backgrounds

In order to understand what accuracy is achievable in a measurement of the neutron lifetime, one must carefully consider the backgrounds. Calculations presented below (see section 7.8) of the accuracy achievable for a fixed running time, depend strongly on the expected backgrounds. In considering what changes might be made to improve the apparatus, it is important to know as much as possible about the current backgrounds, what causes them, how they vary with threshold, and how they might be suppressed. We consider the background as consisting of two components, the “constant” or “flat” background, which does not vary with time (or may vary slowly with time) and the time-varying background. By their very nature, time-varying backgrounds must be caused by the neutron beam being on¹³, since that is the effect that breaks the time symmetry of the system. Constant backgrounds may come from a variety of sources.

6.2.1 Constant Background

Constant backgrounds are caused by scintillation of high energy charged particles in the helium or within the acrylic. Since muons should be vetoed with high efficiency by the muon paddles (see section 4.5.4), the dominant source of constant backgrounds is expected to be compton scattered electrons. The gamma radiation which leads to compton scattering may originate either outside or inside the apparatus. Since the lead shielding (see section 4.5.3) is imperfect, some gammas originate within the

¹³Time-varying backgrounds may also be due to drifts in the PMT gains or backgrounds. These effects should not be correlated from run to run. There may, however, be shifts in the gain of a PMT as it settles down, following the neutron beam exposure period, which could be correlated from run to run and might be indistinguishable from any other neutron beam caused background.

guide hall, primarily as prompt gammas from capture of neutrons on the cadmium shielding which is almost universally used on neutron scattering instruments. Decay of naturally occurring long-lived radioactive isotopes (such as uranium and thorium and their daughters) inside of the lead shielding also may contribute to the constant background.

The constant background is significantly reduced by the shielding techniques employed (see section 4.5). The background coincidence countrate, requiring at least one photoelectron per PMT, decreased from 86 Hz to 19 Hz when the lead shielding was constructed and the muon veto was enabled. Of this reduction the vast majority comes from constructing the lead gamma shield, since the muon paddles were measured to veto 17 Hz at this threshold.

The portion of constant background due to gammas in the guidehall has been characterized by observing the change in background rate within the apparatus when the reactor has shut down. The constant background coincidence rate, requiring at least one photoelectron per PMT, was observed to drop to $54 \pm 4\%$ of its initial value immediately after the reactor shut down. This tells us that at least 45% of the constant background is originating outside of the apparatus and could be eliminated through better shielding. The constant background remaining when the reactor is off could be due to radioactive sources either outside or inside of the lead shielding. Natural radioactivity of materials inside of the cryostat may well account for the reactor off constant background. All structural materials contain Uranium, Thorium and Potassium to some extent. Aluminum Alloy 6061, from which most of the cryostat is made, has been measured to have 0.42 decays/minute/gram due to small amounts of ^{232}Th included in the alloy [70]. Given the roughly 10 kg of aluminum present inside of the lead shielding, approximately 70 decays per second are occurring, each




Figure 6.1: The pulse height spectra of PMT α are shown with the cell empty (dashed line) and full (solid line) of superfluid helium.

producing a 583 keV gamma.

Another important question regarding the constant background is what fraction of it is generated from scintillations in the helium and what fraction is due to Cerenkov radiation in the acrylic lightguides. Two simple measurements were performed to determine the relative contribution as function of threshold of these two sources. The apparatus was set up and the cell was cooled down in the normal fashion to 4 K. With the muon paddles in place and the lead house completely constructed, but with no helium in the cell, data was taken for 1000 s. The pulse height spectra of one of the two PMTs are shown in figure 6.1. The other PMT has substantially the same pulse height spectrum. After completely filling the cell with ultrapure helium and cooling it to 3 K, but otherwise leaving the experimental conditions unchanged, a second 1000 s background measurement was made. The pulse height spectrum from

Table 6.2: Countrates in coincidence with the cell empty or full of ultrapure helium. The rates are tabulated as a function of the threshold number of photoelectrons required on each PMT. Thus, for a 2 photoelectron threshold, the total countrate of events where at least 2 photoelectrons are detected in coincidence on each PMT is reported. Also shown is the fraction of the background (± 0.01) at each threshold which is due to the acrylic.

Threshold (p.e.)	Countrate [Hz]		Fraction due to acrylic
	Empty	Full	
1	19.2	30.5	0.63
2	9.3	17.4	0.53
3	5.4	12.5	0.43
4	3.7	9.8	0.38
5	2.7	8.1	0.34
6	2.2	6.9	0.31
7	1.7	6.0	0.29
8	1.4	5.3	0.27
9	1.2	4.7	0.25
10	1.0	4.2	0.24
15	0.5	2.8	0.19
20	0.3	2.1	0.16

this run is also shown in figure 6.1.

Since the mass of acrylic in the detection system is ten times the mass of helium in the detection region, it may be natural to expect that 90% of the constant background results from detection of gammas in acrylic. Contrary to that expectation, figure 6.1 shows that roughly equal amounts of constant background light originate in the helium and acrylic. Table 6.2 shows the countrate in coincidence as a function of the threshold on each PMT. These results demonstrate that even at the lowest reasonable operating threshold (2 photoelectrons/PMT) only half of the constant background originates within the acrylic. At higher thresholds this fraction is reduced even further. Future modifications to the detection scheme which might allow removal of all of the acrylic light collection system (see section 7.6) would thus result in at most a factor of

2 reduction in the constant background. These schemes may nonetheless be very attractive due to their other main feature: significant increases in overall detection efficiency.

6.2.2 Time Dependent Backgrounds

Developing a clear picture of backgrounds is significantly more complicated from the time-varying than constant backgrounds. Time varying backgrounds may be caused by neutron activation of materials or neutron induced luminescence. In practice it may be impossible to determine which of these sources is responsible for a time-varying background simply by observing the countrate as a function of time for trapping runs.

The coincidence countrate for the average of all of the negative runs in the Cold1 dataset is shown in figure 6.2. When data taking begins, 25 s after the beam is turned off the coincidence countrate is hundreds per second, but falls rapidly. The countrate then falls more slowly and has not reached a constant level by the end of the run at 2700 s after beam off. A falling function with several time constants, such as this one, may be composed either of multiple exponentials (such as from activation of several different isotopes) or of a power law (t^{-power}). Fits to both of these possible functions are about equally good. For example, when binned in 5 s bins, the data can be fit to three exponentials:

$$y_{coinc} = (25.1 \pm 0.1) + (1988 \pm 42) e^{-t/(16.0 \pm 0.1)} \\ + (21.9 \pm 0.7) e^{-t/(148+9-7)} + (9.3 \pm 0.4) e^{-t/(955+76-66)} \quad (6.1)$$

with $\chi^2 = 1.11$ or to a power law plus a single exponential:

$$y_{coinc} = (22.2 \pm 0.2) + (2422 \pm 112) e^{-t/(14.6 \pm 0.2)} + (363 \pm 20) t^{-(0.59 \pm 0.01)} \quad (6.2)$$




Figure 6.2: The coincidence count rate for the negative runs in the Cold1 dataset is shown. The data have been cut to require at least 1 photoelectron in each PMT. The curve fit shown is to the three exponentials (as described in the text). The two curve fits are virtually indistinguishable to the eye on this scale.

with $\chi^2 = 1.10$. The short exponential (14 – 16 s lifetime) observed is likely due to activation of fluorine, which has a lifetime of 15.9 s [67] and is a major component of the GoreTex reflector.

We cannot easily determine by curve fitting alone whether the observed background is due to activation (suggested by the three exponentials) or luminescence (suggested by the power law fit). An activation rate of some few tens per second of two different isotopes is not unbelievable in and of itself. Since we have measured (see section 2.4) that at least 99.9% of the luminescence events result in uncorrelated single photons, the effect of luminescence can be determined by computing the accidental coincidence rate. Although the rates on each PMT are not normally recorded by the




Figure 6.3: The singles (1+ photoelectron) count rate for PMT β is shown at certain sampled points during a run. The curve fit shown is that described in equation 6.3.

DAQ¹⁴, the single PMT rates were occasionally measured using a NIM counter. A measurement of the variation in the singles rate of PMT β with time at certain points within one run is shown in figure 6.3. A best fit of this data to a power function plus a positive offset yields

$$y_{singles} = (0 \pm 525) + (74.4 \pm 1.5 \times 10^3) t^{-0.29 \pm 0.01} \quad (6.3)$$

The singles rate of PMT α was not measured independently for the same run. If

¹⁴This is a significant flaw in the DAQ which is in the process of being eliminated. We have simple counter cards which will be read out once per second. These cards are model PCI-DIO-48H/CTR16 counters from Measurement Computing, Middleboro, MA, <http://www.measurementcomputing.com>. With 32 bits per counter, achieved by ganging two 16 bit counters, a total of 16 channels will be available using two cards. This will allow monitoring of individual PMT count rates as well as various hardware coincidence rates ($\alpha + \beta$, $\alpha + \beta + \mu_i$, etc.).

we assume that the singles rates of the two PMTs are roughly equal¹⁵, then we can estimate the accidental coincidence rate¹⁶ due to luminescence as $y^2 \times 20 \times 10^{-9}$ where y is the luminescence countrate measured in equation (6.3). This yields an estimated accidental coincidence rate of

$$y_{\text{accidental}} = 111 t^{-0.58 \pm 0.02} \quad (6.4)$$

The estimated accidental coincidence rate is thus only enough to explain about 1/3 of the total observed time-dependent background. While there is some uncertainty in the accidental coincidence rate, it seems likely that the observed time dependent background with single photoelectron thresholds is due to a combination of luminescence and activation.

Even if it is present in the single photoelectron coincidence data, it should be dramatically suppressed in coincidence data with higher thresholds. It is expected that the luminescence background consists entirely of uncorrelated single photon events (see section 2.4). With single photoelectron thresholds, this means that luminescence should only be detectable due to accidental coincidence. The excellent gain dispersion of our photomultipliers allows strong discrimination between one and two photoelectron events. If the probability of a single photon event “spilling over” to produce an n photoelectron event is $P_{\text{spill}}(n)$, then the accidental coincidence rate observed with single photoelectron thresholds will be reduced by a factor of $P_{\text{spill}}(n)^2$ for n pho-

¹⁵The two PMTs used were provided with very similar quantum efficiencies and were operated at nearly identical gains. Barring some unforeseen asymmetry in the setup the rates of the two PMTs should be relatively close.

¹⁶The accidental coincidence rate is the product of the countrate of the one PMT times the probability that a pulse on the other PMT occurs at the same time. For PMTs with countrates R_1 and R_2 [s^{-1}] and a total coincidence window (the sum of the output pulse widths of the two PMT discriminators) T [s] wide the accidental rate is simply $R_1 R_2 T$ [s^{-1}].




Figure 6.4: The coincidence count rate for the negative runs in the Cold1 dataset is shown. The data have been cut to require at least 2 photoelectrons in each PMT.

toelectron coincidence data. We have measured $P_{spill}(2) < 0.05$ (see section 4.4.2). Thus, by raising the thresholds to two photoelectrons per PMT, more than a factor of 1000 reduction in any luminescence background is expected. Applying such cuts to the Cold1 dataset does not reduce the overall time-dependent backgrounds nearly that much, as shown in figure 6.4.

Binning the data in 5 s bins and fitting to the same two models as were used for the single photoelectron coincidence data yields a curve fit to three exponentials:

$$\begin{aligned}
 y_{coinc} = & (12.5 \pm 0.1) + (1823 \pm 80) e^{-t/(15.0 \pm 0.2)} \\
 & + (3.4 \pm 0.6) e^{-t/(132+45-27)} + (1.2 \pm 0.3) e^{-t/(871+433-218)} \quad (6.5)
 \end{aligned}$$

with $\chi^2 = 1.11$ or to a power law plus a single exponential:

$$y_{coinc} = (12.2 \pm 0.1) + (1863 \pm 74) e^{-t/(14.8 \pm 0.2)} + (65 \pm 18) t^{-(0.66 \pm 0.06)} \quad (6.6)$$

Table 6.3: The amplitudes and lifetimes of the two longer exponentials in a three exponential fit are shown for negative runs from several of the data sets. The data have in all cases been binned in 5 s bins and cut to require at least one photoelectron per PMT.

Data Set	A_2	τ_2	A_3	τ_3	χ^2
Cold1	21.9 ± 0.7	$148 + 9 - 7$	9.3 ± 0.4	$955 + 76 - 66$	1.11
500 mK	18.9 ± 0.8	$162 + 14 - 13$	8.7 ± 0.6	$986 + 122 - 97$	1.00
Warm	23.1 ± 1.3	$127 + 11 - 9$	9.3 ± 0.5	$900 + 84 - 71$	1.03
^3He	23.5 ± 1.6	$134 + 14 - 12$	8.9 ± 0.7	$816 + 104 - 82$	1.27
Cold2	20.7 ± 0.5	$156 + 6 - 6$	9.1 ± 0.3	$961 + 54 - 46$	1.07

with $\chi^2 = 1.11$. As with the single photoelectron data, equally good fits are obtained either to three exponentials or an exponential plus the power law. The amplitude of the power function is decreased by only a factor of 5 in the two photoelectron data, relative to the single photoelectron data. This is not consistent with our model of luminescence being due to uncorrelated single photon events unless $P_{spill}(2) = 0.45$, which is much larger than our measured value. Hence it seems more likely that the time varying background seen with two photoelectron thresholds is due to activation of some material. Given the large cold neutron flux, and the small observed backgrounds (only a few per second decay rates), and our high detection efficiency either for betas or gammas, this seems like a very reasonable model.

Due to the consistency of our extracted fit parameters across data sets and thresholds, additional weight is added to the model that some of the single photoelectron, and all of the multi-photoelectron, time-dependent background is due to activation. When we fit the single photoelectron data from different temperature datasets to the triple exponential model we get surprisingly consistent results, as shown in table 6.3. If the signal were due to luminescence, this would not be expected, since the behavior of the luminescence signal (at least from Boron Nitride) is known to vary with

Table 6.4: The amplitudes and lifetimes of the two longer exponentials in a three exponential fit are shown for various thresholds (in photoelectrons/PMT). In all cases the data are from negative runs from the Cold1 data set and it have been binned in 5 s bins before fitting.

p.e./PMT	A_2	τ_2	A_3	τ_3	χ^2
1+	21.9 ± 0.7	$148 + 9 - 7$	9.3 ± 0.4	$955 + 76 - 66$	1.11
2+	3.4 ± 0.6	$132 + 45 - 27$	1.2 ± 0.3	$871 + 433 - 218$	1.11
3+	2.5 ± 0.5	$130 + 53 - 29$	0.9 ± 0.2	$922 + 596 - 260$	1.06
4+	1.9 ± 0.5	$123 + 73 - 34$	0.8 ± 0.2	$812 + 474 - 219$	1.06
5+	2.4 ± 1.3	$90 + 51 - 24$	0.7 ± 0.1	$734 + 279 - 158$	1.04
8+	1.0 ± 0.4	$129 + 142 - 44$	0.5 ± 0.2	$774 + 574 - 231$	1.00
10+	0.8 ± 0.2	$161 + 146 - 52$	0.4 ± 0.1	$1018 + 2418 - 421$	1.05
15+	0.5 ± 0.1	$236 + 281 - 83$	0.2 ± 0.1	$1514 + \infty - 1086$	1.15

temperature between 250 mK and 1 K [42, 66]. Similarly, if we examine the variation in the fit parameters with threshold for the Cold1 data set, we see that the lifetimes remain consistent even at high thresholds (see table 6.4). Even at high thresholds where the signal is small and the error bars large, the lifetimes extracted by fitting remain consistent¹⁷. It is almost inconceivable that any luminescence is observable when requiring 15 photoelectrons per PMT (since $P_{spill}(15) \sim 10^{-4}$), unless there is some multi-photon luminescence mechanism. In considering all of the above analysis together, it seems most likely that all of the time varying background at thresholds of at least two photoelectrons per PMT is due to activation. It seems unlikely that any information about which material is actually activated can be extracted from the above curve fits. There are likely more than 2 activated isotopes (plus fluorine),

¹⁷Given the large number of parameters used in these nonlinear fitting models, Depending on the initial parameter guesses, other choices for the lifetimes may fit the data equally well. While the exact lifetimes extracted may not provide any insight into the activation processes at work, the fact that the same fits are consistent with the data over many thresholds and temperatures is indicative of the fact that the background has a consistent time-dependence, even if that time dependence is not uniquely determined.

and so extracting individual lifetimes may well be impossible. The best approach to reducing such backgrounds is to reduce the number of neutrons hitting materials in the cell and to choose the cell materials carefully (see chapter 2).

6.3 Evidence of Neutron Trapping

In order to search for evidence of trapped neutrons it is necessary to consider the difference between positive and negative runs. Ideally, the difference countrate (computed as described in section 5.4) would be due only to trapped neutrons if the backgrounds in positive and negative runs were identical. There are certain mechanisms which may lead to imperfect background subtraction. It is known, for example, that the luminescence rate from boron nitride changes with magnetic field [66]. Hence luminescence backgrounds may be different in positive and negative runs. Furthermore the presence of the magnetic field focuses neutrons in one spin state and defocuses those in the other. Hence the distribution of neutron captures (on shielding materials and impurities) may be different in positive and negative runs. In order to be sure that any observed difference signal is due to trapped neutrons and not imperfect background subtraction (due to one of the above effects or something similar), runs must be compared from trapping experiments made with ultrapure and natural abundance helium. Any difference signal observed in the ^3He dataset (runs taken with natural abundance helium) must be due to imperfect background subtraction. Since the contribution to the time-dependent background from activation and luminescence varies strongly with threshold, any effect of imperfect background subtraction may also be different for different thresholds.

With single photoelectron thresholds, imperfect background subtraction is a significant problem and prevents direct analysis of the difference data. The single photo-




Figure 6.5: The difference count rate for the (a) Cold1 and (b) ^3He datasets are shown. Both data sets have been cut to require at least 1 photoelectron in each PMT.

electron difference count rates for the Cold1 and ^3He datasets are shown in figure 6.5. The presence of a time-dependent signal in the ^3He data indicates that the single photoelectron backgrounds are different between positive and negative runs. It is thus necessary in general to analyze the data by requiring more photoelectrons per PMT.

One possible method for extracting useful information from the cold single photoelectron data is to subtract the difference count rates of the Cold1 and ^3He datasets. One should expect imperfect background subtraction to be equally problematic in both datasets. Therefore the difference count rate in the ^3He dataset is, in effect, a measurement of the imperfect subtraction. By the same logic that is used to subtract negatives from positives, the ^3He difference count rate can be subtracted from the Cold1 difference count rate to extract a signal which must be due to trapped neutrons. The result of such a subtraction is shown in figure 6.6. This data fits well




Figure 6.6: The difference between the countrates in the “Cold 1” and the ^3He dataset is shown. The data have been cut to require at least 1 photoelectron in each PMT. A curve fit to a single exponential plus a constant offset is shown.

($\chi^2 = 1.05$ to a single exponential plus a constant offset¹⁸. While the resultant lifetime $\tau = 856^{+1385}_{-327}$ s is consistent with the neutron lifetime, the error bars (dominated by the poor statistics of the ^3He dataset) limit the utility of this result.

By analyzing the data with higher thresholds, one may hope to get away from problems caused by imperfect background subtraction. Requiring at least two photoelectrons per PMT results in a flat difference signal for the ^3He dataset, as shown in figure 6.7. With these thresholds, a non-zero difference signal is still observed in

¹⁸Unless otherwise noted, all of the data that follows has been fit to a single exponential plus a constant offset. The data are binned into 50 s bins. The fitting procedure only considers data from $t = 100$ s to $t = 2650$ s. The high countrates at early times (due to Fluorine activation) increase the probability of imperfect background subtraction. The data analysis procedure occasionally misses events at the end of the run, so the last bin is excluded.




Figure 6.7: The difference countrate for the ^3He dataset is shown. The data have been cut to require at least 2 photoelectrons in each PMT. In the fit shown, the lifetime is held equal to the lifetime extracted from the Cold1 data (see figure 6.8).

the Cold1 data, as shown in figure 6.8. Fitting the Cold1 difference data to a single exponential plus a constant offset yields:

$$y = (0.24 \pm 0.02) + (0.86 \pm 0.20)e^{-t/(312+116-66)} \quad (6.7)$$

with a χ^2 per degree of freedom of 1.11. Fitting the ^3He difference signal to an exponential fixed to the lifetime extracted from the Cold1 dataset results in an amplitude of -0.22 ± 0.19 . The difference between the amplitudes for the Cold1 and ^3He datasets is thus 3.9σ . Hence we can state with $> 99.9\%$ certainty that a difference exists between the Cold1 and ^3He datasets, which is evidence for the magnetic trapping of UCN.

The lifetime extracted from the Cold1 dataset is reasonably consistent with the lifetime of 660^{+290}_{-170} s obtained in our original demonstration of neutron trapping [42].




Figure 6.8: The difference countrate for the “Cold 1” dataset is shown. The data have been cut to require at least 2 photoelectrons in each PMT.

Comparing the two results as decay rates $k = 1/\tau$, we have $(3.22 \pm 0.87) \times 10^{-3} \text{ s}^{-1}$ and $(1.52 \pm 0.46) \times 10^{-3} \text{ s}^{-1}$ respectively. Adding the errors in quadrature, the difference between the results is 1.7σ . In decay rate space, the lifetime extracted from the Cold1 dataset is 2.4σ different from the world average value of $885.7 \pm 0.8 \text{ s}$ for the neutron lifetime[5]. Therefore we have evidence for trapping of UCN, but the observed trap lifetime appears to be shorter than the accepted value of the neutron lifetime.

6.4 Higher Thresholds

An important consistency check can be performed by analyzing the data at various thresholds and comparing the extracted trap lifetimes. For the Cold1 dataset, the extracted fit parameters are shown in table 6.5 for thresholds of 1 to 15 photoelectrons

Table 6.5: The fit parameters extracted from fitting the difference data (positive - negative) are shown as a function of the cutting threshold (in photoelectrons/PMT) for the Cold1 dataset. The data are fit to the function $y = y_0 + Ae^{-t/\tau}$. The χ^2 per degree of freedom is shown as a measure of the goodness of the fits.

p.e./PMT	y_0	A	τ	χ^2
1 ^a	0.25 ± 0.18	1.08 ± 0.18	$886 + 1385 - 327$	1.06
2	0.24 ± 0.02	0.86 ± 0.20	$312 + 116 - 66$	1.11
3	0.16 ± 0.02	0.70 ± 0.14	$366 + 128 - 75$	1.10
4	0.14 ± 0.02	0.70 ± 0.14	$323 + 103 - 63$	1.21
5	0.12 ± 0.02	0.67 ± 0.13	$312 + 96 - 59$	1.23
8	0.08 ± 0.01	0.46 ± 0.10	$332 + 119 - 69$	1.05
10	0.07 ± 0.01	0.36 ± 0.10	$309 + 135 - 72$	1.02
15	0.01 ± 0.02	0.13 ± 0.03	$830 + 1462 - 323$	1.05

^aFor the single photoelectron threshold only, the data considered was (Cold1 difference countrate) - (³He difference countrate).

per PMT. With the exception of the single photoelectron (different extraction procedure) and 15 photoelectron (extremely small amplitude) thresholds, the extracted lifetimes are all within one standard deviation of each other. The lifetimes extracted from one and two photoelectron thresholds are discrepant by 1.9σ . This difference may suggest that the short lifetime seen in all of the higher threshold data is due to imperfect background subtraction. Poor statistics for the ³He dataset make the double background subtraction technique unfeasible for the higher threshold data, since the resulting errors would be exceedingly large. The evidence of persistent problems due to imperfect background subtraction is extremely troubling, and suggests that in the future even greater care should be taken to eliminate time-dependent backgrounds. One experimental approach which may allow this is described in section 7.3.

6.5 Temperature Variation of the Trap Lifetime

The variation in the extracted trap lifetime with temperature is also an important consistency check. At higher temperatures, two phonon upscattering of UCN should reduce the trap lifetime. The theoretically calculated two phonon upscattering rate is proportional to T^{-7} [54] and this rate has been measured to become comparable to the neutron lifetime at about 750 mK [52]. At temperatures of 500 mK and below, upscattering should make a negligible change in the observed lifetime at the level of accuracy of these lifetime measurements. Therefore, if the data are consistent with a simple model of a single short trap lifetime we should see a consistent amplitude and lifetime extracted from all datasets colder than 500 mK and a shorter lifetime and smaller amplitude signal in the 750 mK and warm datasets.

The parameters extracted from fitting the various datasets to a model containing a single exponential decay plus constant offset are shown in table 6.6. In all cases, the data have been cut to require two photoelectrons per PMT. As expected, the amplitudes of the 750 mK and Warm datasets are significantly reduced from those of the colder datasets. The Warm dataset is consistent with a flat background, as is expected from the short (tens of seconds) ^4He upscattering lifetime at this temperature (> 1 K) [52]. A graphical representation of the extracted decay rates ($k = 1/\tau$) for the different data sets is shown in figure 6.9. As can be seen from the figure, the extracted decay rate varies across the different temperature sets with $T < 600$ mK. The maximum discrepancy, between the 400 mK and Cold1 datasets, is 2.4σ . The Cold2 dataset, which includes the data from all four datasets with runs colder than 600 mK, is consistent with all of the datasets other than the 400 mK data. On balance, it seems likely that all of the data below 0.5 K are consistent with a single exponential lifetime. The smallest statistical errors come from the Cold2 data, which

Table 6.6: The fit parameters extracted from fitting the difference data (positive - negative) are shown for the various temperature datasets. All fits are performed on data cut to require at least two photoelectrons per PMT. The data are fit to the function $y = y_0 + Ae^{-t/\tau}$. The χ^2 per degree of freedom is shown as a measure of the goodness of the fits.

Data Set	y_0		A	τ	χ^2
Cold1	0.24 ± 0.02		0.86 ± 0.20	$312 + 116 - 66$	1.11
$^3\text{He}^a$	0.004 ± 0.031		-0.22 ± 0.19	$312 + 0 - 0$	1.03
325 mK	0.10 ± 0.03		0.83 ± 0.47	$232 + 314 - 85$	1.08
400 mK	-0.09 ± 0.14		0.96 ± 0.11	$1083 + 802 - 323$	1.04
500 mK	-0.04 ± 0.05		0.83 ± 0.11	$674 + 283 - 154$	1.03
750 mK	-0.04 ± 0.09		0.43 ± 0.12	$802 + 2898 - 353$	0.91
Warm	-0.04 ± 0.18		0.27 ± 0.13	$1371 + \infty - 848$	1.04
Cold2	0.10 ± 0.02		0.76 ± 0.07	$525 + 104 - 75$	0.90

^aFor the ^3He dataset, the lifetime is held constant and equal to the value extracted from the Cold 1 dataset.




Figure 6.9: The extracted decay rates are shown for the different data sets. The temperature axis used is for representation only.




Figure 6.10: The difference count rate for the Cold2 dataset is shown. The data have been cut to require at least 2 photoelectrons in each PMT.

is a combination of all of the colder data. The Cold2 data, cut to require two photoelectrons per PMT, are shown in figure 6.10. In that case, the best value for the trap lifetime is 525^{+104}_{-75} s. This result is within the error bars of our previous trapping result [42]. This result is 2.5σ discrepant from the accepted neutron lifetime value of 885.7 ± 0.8 s [5].

6.6 Calibration of the Detection Efficiency

In order to convert between the observed rate of neutron decays and the number of UCN produced and trapped in the apparatus, it is necessary to know the detection efficiency. The detection efficiency, or probability of detecting a trapped neutron which decays, depends on the amount of the decay energy which is taken away by

the electron. Depending on the detector setup it may also depend on the neutron's position at the time of decay. We have performed several measurements, independent of the trapping runs, in order to determine the detection efficiency.

In the design phase of the detection system, measurements of different setups were made using a room temperature cell testing apparatus [65]. This apparatus contained a detector insert of TPB evaporated onto goretex. High purity Argon gas, which acted as the scintillation medium in place of liquid helium, was continuously flowed through the system. An alpha source produced EUV scintillations in the Argon gas. In addition to many other tests performed with this system, the relative detection efficiency was measured for three positions. The strongest signal was measured with the alpha source in the center of the cell, 17.8 cm from the end of the lightguide. With the alpha source 14 cm from the lightguide the relative detection efficiency was 92%. With the alpha source 30.5 cm from the end of the lightguide, the relative detection efficiency was 32%. The detection efficiency decreases as one moves away from the center of the cell in either direction. As one moves towards the lightguide, the solid angle for EUV to hit the TPB decreases (since the beamstop and lightguide are not coated with TPB). As one moves away from the lightguide, light generated must increasingly undergo diffuse reflection in order to reach the lightguide, decreasing the overall detection efficiency. Since the measurements of the detection efficiency as a function of position are so limited, we will model the position dependence of the relative detection efficiency ϵ_r as

$$\epsilon_r(x) = 1 - (0.03 \pm 0.01)|x| \quad (6.8)$$

where x is the distance in cm from the center and the slope has been taken as the average of the slopes implied by the two off-center datapoints, above. The total trap length is only 22 cm, so the maximum value of x is ± 11 cm.




Figure 6.11: Pulse height spectrum of the detector from a ^{113}Sn source mounted in the center of the detection region.

The average detection efficiency of a 360 keV electron generated at the center of our trapping region was measured immediately after the conclusion of our trapping measurements. Without changing any of the optics in the system, a ^{113}Sn conversion beta line source was placed in the center of the cell. The apparatus was subsequently cooled to less than 250 mK and light pulses were observed in the usual manner. The pulse height spectrum produced by the detector with this source in place and a single PMT mounted at the end of the 300 K lightguide (without the “Y” in place) is shown in figure 6.11. A gaussian curve fit to this measurement yields a peak at 10.7 ± 1.1 photoelectrons¹⁹. This corresponds to a quantum yield of 29.7 ± 3.1

¹⁹The uncertainty in the peak position is dominated by the uncertainty in the conversion of MCA channel number to photoelectrons. This conversion factor was determined from a separate measurement of the PMT’s “dark” spectrum with the same operating voltage and amplifier gains.




Figure 6.12: The energy spectra of recoil electrons from the decay of the free neutron and ^{113}Sn are shown. ^{113}Sn produces conversion electrons at a single energy of 360 keV. The neutron recoil spectrum peaks at 250 keV. The total released energy in neutron beta decay is 782 keV.

photoelectrons per MeV of electron kinetic energy. The total efficiency of the splitting “Y” was measured to be 94% (that is 47% of the light goes to each PMT).

The detection efficiency is measured using the tin line source in order to obtain the sharpest possible peak. Electrons produced in neutron beta decay have a range of energies. The relative spectra of the tin source and neutron beta decay are shown in figure 6.12. The number of EUV photons produced by recoil of the electron through the helium is linear in the electron’s energy.

The average detection efficiency for our apparatus can be calculated using a simple Monte Carlo simulation. This simulation takes into account the neutron beta-decay spectrum, our simple model of position dependence, finite “Y” efficiency, and considers the splitting of photons between the two PMTs to be a poisson process. The

Table 6.7: Detection efficiency as a function of the number of photoelectrons required per PMT. Above 8 photoelectrons per PMT the detection efficiency is 0.

Threshold [p.e.]	ϵ [%]
1	86 ± 4
2	65 ± 9
3	45 ± 11
4	28 ± 11
5	15 ± 9
8	1 ± 2

code for the simulation is found in Appendix F. The estimated detection efficiency, ϵ , as a function of threshold is shown in table 6.7. The dominant error is the uncertainty in the position dependence of the detection efficiency. The detection efficiency is estimated to be $65 \pm 9\%$ for a threshold of 2 photoelectrons per PMT.

We can compare our estimated detection efficiency as a function of threshold with the observed variation in the amplitude of the trapping signal as a function of threshold as a consistency check. We can calculate the implied initial decay rate $A_{weighted}$ as a function of threshold, $thresh$, from:

$$A_{weighted}(thresh) = \frac{A_{raw}(thresh)}{\epsilon_{calc}(thresh)} \quad (6.9)$$

where A_{raw} is the amplitude extracted from curve fitting, tabulated as a function of threshold in table 6.5, and ϵ_{calc} is the detection efficiency tabulated in table 6.7. The estimate of $A_{weighted}$ for 8 photoelectrons per PMT is not shown since the fractional error is enormous (due to the large uncertainty in $\epsilon_{calc}(8)$). At low photoelectron number $A_{weighted}$ is consistent with a constant value (see figure 6.13).




Figure 6.13: The weighted decay amplitude $A_{weighted}$ is shown as a function of the threshold. The fit to $A_{weighted}$ for thresholds of 1 to 3 photoelectrons per PMT is 1.35 ± 0.17 .

6.7 UCN Production Rate

The observed decay amplitude can be compared with the predicted UCN production rate to test the consistency of the theory of UCN production. As long as the trap is loaded for a time which is long compared with the trap lifetime, and neutrons are not lost by any mechanism which does not produce scintillation events, then the initial rate of neutron decays, $A_{weighted}$, must be equal to the production rate of UCN within the trap. In general, the number of neutrons in the trap, N_{trap} after loading for a time T_{fill} will be

$$N_{trap} = P\tau_{trap}(1 - e^{-T_{fill}/\tau_{trap}}) \quad (6.10)$$

where P is the UCN production rate and τ_{trap} is the lifetime of UCN in the trap. The trap lifetime is set by the neutron beta decay lifetime τ_n as well as any other loss

mechanisms τ_{other} where

$$\frac{1}{\tau_{trap}} = \frac{1}{\tau_n} + \frac{1}{\tau_{other}} \quad (6.11)$$

and in general the other loss mechanisms (such as UCN upscattering) will not produce a scintillation event²⁰. During the observation cycle, the number of neutrons remaining in the trap will decrease according to

$$N(t) = N_{trap} e^{-t/\tau_{trap}} \quad (6.12)$$

but the rate of decay events observed will be

$$R(t) = \frac{N(t)}{\tau_n} = \frac{P\tau_{trap}(1 - e^{-T_{fill}/\tau_{trap}})e^{-t/\tau_{trap}}}{\tau_n} \quad (6.13)$$

Therefore the initial observed decay rate $A_{weighted}$ can be used to calculate the production rate from

$$P = A_{weighted} \frac{\tau_n}{\tau_{trap}} \frac{1}{(1 - e^{-T_{fill}/\tau_{trap}})} \quad (6.14)$$

From the previous equation we see the difficulty with extracting the true production rate comes from our uncertainty regarding the trap lifetime, τ_{trap} . There are three plausible choices for τ_{trap} : 1) the value extracted from the Cold1 data analyzed with single photoelectron thresholds and double subtraction, 2) the value extracted from the Cold1 dataset analyzed requiring two photoelectrons per PMT (this value is consistent for thresholds of two to 10 photoelectrons per PMT), 3) the value extracted from the Cold2 dataset with two photoelectron threshold. Each of these choices has its advantages and disadvantages. The Cold2 dataset contains the most data. A possible temperature variation of the data (see table 6.6) calls into question the validity of combining data together to form the Cold2 dataset. Similarly, the

²⁰Capture of UCN on ^3He is a notable exception, where a UCN loss mechanism would result in an observed scintillation [121]. If capture on ^3He is the dominant loss mechanism, then equation (6.13) should be modified by replacing τ_n with τ_{trap} .

Cold1 dataset analyzed with thresholds from 2 to 10 photoelectrons per PMT yields the same lifetime estimate. However, we expect that no trapped neutrons should be detected at 10 photoelectrons per PMT, raising the serious question of a bias in the lifetime estimate due to imperfect background subtraction. The Cold1 data analyzed with single photoelectron thresholds has had the imperfect background subtraction effect removed, however the multiple subtractions performed may have added a bias of their own. We can compute the production rate for each of these scenarios, using

$$A_{weighted} = 1.35 \pm 0.17 \quad (6.15)$$

which is calculated from the weighted average of $A_{weighted}(thresh)$ for $thresh = 1-3$. Then the estimated production rate, in UCN s^{-1} is

$$P_1 = 1.5 \pm 1.2 \quad (6.16)$$

or

$$P_2 = 3.8 \pm 1.1 \quad (6.17)$$

or

$$P_3 = 2.3 \pm 0.5 \quad (6.18)$$

for the three values of τ_{trap} considered above.

The expected production rate is computed using a Monte Carlo simulation (see Appendix C). The Monte Carlo simulation predicts a production rate for the original UCN trapping experiment which is consistent with the observed production rate [42]. The simulation has since been extended to simulate diffraction of cold neutrons from the monochromator, and its operation has been validated through comparisons with measurements of the monochromatic beam. Using this simulation, for the conditions under which the experiment was operated the production rate is expected to be:

$$P_{expected} = 3.8 \pm 0.8 \quad (6.19)$$

This production rate has been corrected for air scatter over the 3.8 m cold neutron flight path.

Hence we see that we are in one of two possible situation. In one case the neutron lifetime in the cell is short, due to some unexplained loss mechanism, and the production rate is consistent with our expected value. Otherwise, if the neutron lifetime in the cell appears short due to imperfect background subtraction, and the true lifetime is consistent with the neutron lifetime, as seen in the single photoelectron data, then the measured production rate is 2-3 times lower than predicted by our Monte Carlo simulations.

6.8 Non-8.9Å production

One possible flaw in the Monte Carlo simulations is that UCN production is considered to be due only to the single phonon downscattering process in the region near 8.9 Å. Calculation of the production rate in this manner yielded good agreement with the measured production rate in our earlier work [42]. Nevertheless, multiple-phonon downscattering processes can scatter non-8.9 Å neutrons to produce UCN. If the lifetime in our trap is consistent with the neutron lifetime, and the production rate on the monochromatic beam is indeed low, then one reason could be the lack of non-8.9 Å neutrons in the monochromatic beam. The fraction of UCN production in the NG-6 “white” neutron beam due to non-8.9 Å neutrons is estimated to be no more than 0.5 [137].

Chapter 7

Conclusions and Future Directions

7.1 Magnetic Trapping of UCN

Magnetic trapping of UCN is demonstrated by the statistically significant difference between the observed difference signals with and without ^3He in the trapping region. As shown in table 6.6, the amplitudes of the exponential fit (with 312 s lifetime) are 0.86 ± 0.20 and -0.22 ± 0.19 for the Cold1 and ^3He datasets, respectively. This represents a statistically significant difference at the 3.9σ level. Even greater confidence can be obtained by comparing the ^3He dataset with the Cold2 dataset. When both datasets are fit to a constant offset and an exponential with lifetime 525 s, the amplitudes of the exponential are 0.76 ± 0.07 and -0.20 ± 0.14 for the Cold2 and ^3He datasets, respectively. Thus, magnetic trapping is established at the 6.1σ level. Since the error is dominated by the ^3He measurements, additional ^3He runs would be valuable in strengthening the verification of magnetic trapping. Additional ^3He runs would also, as discussed below, be extremely valuable in determining whether or not imperfect background subtraction is introducing a bias in the extracted fit lifetime for the trapping signal.

7.2 Time Dependent Background

Several important conclusions can be drawn from our neutron trapping measurements. First, we are trapping UCN produced from a monochromatic beam of 8.9 Å neutrons, as evidenced by the statistically significant difference between a decay signal seen in cold trapping runs with ultrapure ^4He , and a flat signal seen in cold trapping runs with natural abundance helium (containing ^3He at the 10^{-7} level). Beyond that, many of the most important characteristics of the system are determined, but with relatively large uncertainty, including both the trap lifetime and the UCN production rate. The difficulty in determining these characteristics is due to the presence of a time dependent background.

A time dependent background is seen at all detector thresholds. With single photoelectron thresholds, the observed time dependent backgrounds may be 30–100% due to luminescence. The persistence of the time dependent background at high photoelectron thresholds is inconsistent with the model of luminescence producing uncorrelated single photon events. The time dependent background for 2 or more photoelectrons per PMT is likely due to neutron activation of materials within the cryostat. It may be possible to determine which material or materials within the apparatus are the source of this activation. However, given the great care that has already gone into materials selection, it seems unlikely that significant reductions in the activation rate can be made in the current geometry.

In order to measure the neutron lifetime in the presence of a time dependent background it is necessary to either understand the background completely or to subtract it. Since we do not have an exact model for our time dependent backgrounds, we subtract negative and positive runs to isolate the signal due to decaying trapped neutrons. Any differences in the backgrounds between positive and negative runs will

lead to an imperfect background subtraction which will then bias the trapping data and the extracted neutron lifetime. Either luminescence or activation backgrounds may be different in positive and negative runs. Luminescence has been shown to vary with temperature and magnetic field for BN, and this dependence may well exist for other materials in the cell which may be luminescing [66]. Furthermore, the number of neutrons striking different parts of the apparatus may differ in positive and negative runs due to the effect of magnetic focusing. When the magnetic trap is on, one neutron spin state is focussed towards the center of the trap and another is defocussed. This leads to a different distribution of neutron absorptions among the various materials within the cell, and could result in different activation rates depending on which materials are being activated.

The presence or absence of imperfect background subtraction can be ascertained by analyzing data taken in ^3He runs. Since these runs do not contain trapped neutrons, any difference between positive and negative runs must be due to imperfect background subtraction. With single photoelectron thresholds, a large difference signal is observed in the ^3He data. This indicates that at least part of the time dependent background (probably, in large part the luminescence background) is imperfectly subtracted. Although luminescence is not expected to contribute significantly at higher thresholds, there is evidence for some kind of imperfect background subtraction. At thresholds of 8 photoelectrons per PMT and higher, the detection efficiency for trapped neutrons is zero. Nevertheless, a time varying decay signal is seen at thresholds of 10 and even 15 photoelectrons per PMT. Since no neutron decays should yield such large events, these time dependent difference rates must be due to imperfect background subtraction. The fact that some part of the difference threshold seen at lower thresholds is likely due to imperfect background subtrac-

tion throws into question the entire validity of that data. Adding to this concern is the difference in the extracted lifetimes for data analyzed by subtracting ^3He from Cold1 data analyzed with single photoelectrons and Cold1 data alone analyzed with higher thresholds. The single photoelectron data, corrected for imperfect background subtraction, fits to a longer neutron lifetime than the uncorrected data with higher thresholds, which gives consistent lifetimes for 2–10 photoelectrons per PMT. This uncertainty in the trap lifetime is exactly the sort of uncertainty that is caused by the presence of a time dependent background and the possibility of imperfect background subtraction.

In order to make a measurement of the neutron lifetime using the magnetic trapping technique, one must eliminate the time dependent backgrounds in the system, and be confident that they are eliminated. It may be possible to attain this goal by identification and replacement of the activated materials in the apparatus. Another approach, which would not require changing the materials in the cell, would be the use of separate production and detection regions.

7.3 Separate Production and Detection Regions

A possible solution to the problem of time dependent backgrounds comes in the form of a significant departure from the current technique —by producing UCN in a separate region of the apparatus from the region in which they are detected. A schematic of a possible design of an apparatus using separate detection and production regions is shown in figure 7.1. Although many specific implementations of this general idea may be possible, the general characteristics are the same: UCN are produced by the superthermal process and are stored within a bottle with material walls; a small diameter umbilical tube connects the UCN production region with the magnetic




Figure 7.1: A possible layout for an apparatus with separate production and detection regions is shown.

trapping and neutron decay region; a magnetic valve, consisting of a high field solenoid prevents neutrons in the low-field seeking (trappable) spin state from exiting the production region; the neutron beam might be able to be left on at all times so that activation and luminescence backgrounds are constant rather than time-dependent.

The most significant change in the operation of the experiment would be the ability to run with the neutron beam passing through the production region at all times. UCN would be continually produced within the storage region, and the density of stored UCN within the production region would slowly build up. In addition to the UCN losses inside of a magnetic trap (due to beta decay, upscattering, or absorption on ^3He), UCN would be lost from the production region due to interactions with the walls and leakage through the magnetic valve. A small magnetic holding field would necessarily surround the production region, to define an axis for spin-alignment and hence prevent Majorana flips. In that case, only high-field seeking neutrons would be lost through the magnetic valve. The constant leakage of high-field seeking UCN through the valve would lead to a constant background, which would be down by at

least $1/\tau_{bottle}$ compared with the amplitude of the neutron decay signal. Similarly, backgrounds produced by activation or neutron-induced luminescence of the materials in the production region, where the neutron flux is large, would be time-independent (once the beam had been on for a time long compared with the decay lifetime of the activated species). The physical separation between the production and detection regions could be quite large (several meters), allowing for reduction of these constant backgrounds by gamma shielding and light baffling¹. The simpler geometry of the setup should also allow more effective shielding of the external constant backgrounds in the detection region by increasing the coverage of lead shielding and muon vetoing paddles.

UCN would be loaded into the trap using the “trap door” method. In this case, the pinch solenoid on the side of the trap connected to the production region would be lowered to zero. The magnetic valve would then be opened and the UCN density would equalize between the production and trapping volumes. The pinch coil would then be raised and neutron decays observed. Because the number of neutrons interacting with materials in the trapping region would be dramatically reduced (by more than 9 orders of magnitude), time varying backgrounds inside the trapping region would be eliminated at any measurable level.

The main potential downside of using separate production and detection regions is a decrease in the number of trapped UCN. Several factors are important in determining how many UCN would be produced and trapped. Since UCN are produced within the production region, and then the UCN densities are equalized between production

¹The use of significant shielding between the production and detection region would allow a large reduction in time-dependent backgrounds even for the case where the (production) neutron beam could not be left on during the detection period. Hence the approach of separating the production and detection regions is attractive even in this case.

and trapping regions before raising the trap, the relative volumes of the two regions are important. If the volume of the production region could be made much larger than the trapping region, then the density in the trapping region would be essentially equal to the density in the production region with the magnetic valve closed. Practical considerations would probably limit the production region to be roughly the same size as the trapping region, and hence lead to about a factor of two loss in density². The number of UCN trapped would also be decreased, relative to our current setup, by the ratio of the neutron lifetime in the production region, τ_{bottle} , to the beta decay lifetime, τ_n . Details of the materials chosen for the walls and the geometry would ultimately determine τ_{bottle} . Numerous studies of potential materials have already been performed [30]. One very promising material is beryllium, which has a UCN potential of 252 neV and a theoretical loss rate per bounce of 2×10^{-6} [30, 138]. The best measured loss probability per UCN bounce which has been obtained for a beryllium film is $2f = 2.3 \pm 0.7 \times 10^{-5}$, where f is the ratio of real to imaginary components of the Fermi potential [138]. This is decidedly worse than the theoretical value, despite cooling the beryllium to 10–15 K to minimize inelastic upscattering due to hydrogen bound to the surface. Another attractive material is solid oxygen which has a lower UCN potential (80 neV) but from which coatings have been made with a measured $2f = 1.7 \pm 2.4 \times 10^{-6}$ [138].

²This is because the neutron beam must pass through the entire production region in order to maximize the UCN density. Hence for a cylindrical production region parallel to the neutron beam, the diameter should be about the beam diameter (as is the case for our trapping region). The length of the production region could then be as long as desired, but would likely be limited to not much more than the trap length of 22 cm in order to prevent the cryostat from being unreasonably large or complicated.

The storage lifetime due to wall losses can be calculated as:

$$\tau_{bottle} = \frac{4V}{S} \frac{1}{\mu(v)v} \quad (7.1)$$

where V and S are the volume and surface area of the bottle, v is the velocity of the UCN, and $\mu(v)$ is the loss probability per bounce, which can be calculated as

$$\mu(v) = 2f \left(\frac{\arcsin(y)}{y^2} - \frac{\sqrt{1-y^2}}{y} \right) \quad (7.2)$$

where $y = v/v_{crit}$ is the ratio of the UCN velocity to the maximum velocity which can be totally externally reflected by the Fermi potential of the material. For a production region 15 cm in diameter, and 25 cm in length, the lifetime for UCN with the maximum energy (68 neV due to the magnetic field + 17 neV from the Fermi potential of liquid helium) trappable in our current 1.1 T deep trap can be estimated using equation (7.1) as 3235 s with beryllium walls or $> 1.2 \times 10^4$ s with solid oxygen walls. Other wall materials such as diamond-like carbon coated by chemical vapor deposition [139] or laser ablation [140] may be even better both in terms of the Fermi potential and the loss probability per bounce. Thus, the number of UCN trapped should not be reduced by the bottle lifetime since it seems feasible to obtain bottle lifetimes in excess of the neutron lifetime.

An additional loss may occur as the UCN are transported from the production region to the trapping region. However, as long as the diffusion time within the umbilical region is short, the requirements for the storage lifetime in the transport region may be substantially relaxed.

Several statistical improvements also result from this arrangement. The number of UCN trapped will be relatively higher than in the trapping setup, since the production region can be placed much closer to the neutron guide. Furthermore, it should be possible to surround the production region with a tamper of several centimeters of

polyethylene, which is expected to increase the UCN production rate by a factor of 2 [141]. The ability to fill the production region while UCN decays are being observed will improve the duty cycle. The effective increase in running time would be approximately a factor of two. Thus it appears probable that the effective number of UCN trapped with separate production and detection regions could be maintained at about the current level. The ability to trap about the same number of UCN while entirely eliminating time varying backgrounds is extremely appealing and is likely the only possible avenue to a successful measurement of the neutron lifetime using magnetically trapped UCN.

An added benefit of using separate production and detection regions is the option to eliminate any possible systematic error due to marginally trapped neutrons. Marginally trapped neutrons are those UCN with energies greater than the trap depth but which are in semi-stable orbits that take a time comparable to the neutron lifetime to escape above the trap wall [142]. If the UCN potential of the wall material in the production region is lower than the trap depth, then by the time the trap is loaded there will be no UCN with energies above the trap depth. Hence, marginal trapping becomes impossible.

7.4 Permanent Monochromator Installation

A variety of improvements to the overall apparatus may be made, independent of the introduction of separate production and detection regions. The first such improvement, which is already underway, is the permanent installation of the monochromator as part of the NCNR infrastructure. The layout of the 8.9 Å beamline after the monochromator is permanently installed is shown in figure 7.2 (compare with the present layout, shown in figure 4.2). This permanent placement of the monochroma-

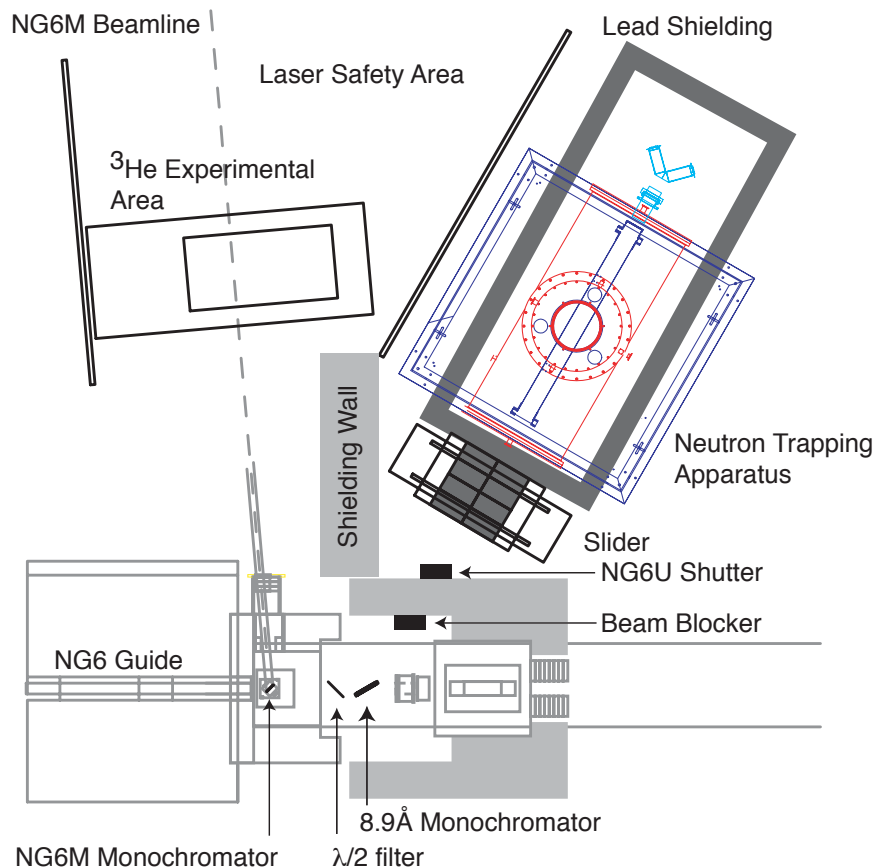


Figure 7.2: The layout of the end station at NG-6 is shown as it will be after permanent installation of the monochromator. The monochromator is installed inside of the NG-6 shielding walls. The apparatus is shown in its possible location as of March 2002.

tor will decrease the total neutron path length from the end of the neutron guide to the trapping region by about 1.5 m. The shorter neutron path length results in , a higher flux at the apparatus, due to less air scatter and a smaller effect from divergence. This repositioning of the monochromator will result in a factor of 1.7 increase in the number of trapped UCN in the current apparatus.

The construction of a separate 8.9 Å beam will allow operation of the UCN trapping experiment independent of other fundamental physics measurements being performed on the end station of the white beam. Since the cold neutron guide NG-6

is 6 cm by 15 cm, the monochromator will be placed in the lower half of the beam while the upper half is passed undisturbed to the end station. The ability to run two experiments concurrently will allow more fundamental physics measurements to be performed at the NCNR. Furthermore, the effective allowable running time for the UCN trapping experiment will be increased.

7.5 Source Improvements

The cold source used at the NCNR will be replaced by February of 2002. The new cold source, which uses an ellipsoidal rather than spherical shell of liquid hydrogen and incorporates other design enhancements, is expected to increase the total cold neutron flux by 1.68 and the 8.9 Å flux by 1.82 [118]. Assuming no other changes besides the cold source upgrade and monochromator installation, the total number of trapped neutrons should be a factor of 3 higher in February 2002 than it was in our August, 2001 measurements. Of course, the time dependent backgrounds should scale up by an identical factor.

In the longer term, higher intensity cold neutron sources could be used, such as the Spallation Neutron Source (SNS), under construction in Oak Ridge, TN, or the reactor at the Institut Laue Langevin (ILL) in Grenoble, France. A beamline for fundamental physics measurements is being proposed for construction at the SNS. Monte Carlo simulations of initial versions of the proposed design suggest that operating the experiment at the SNS should increase the number of trapped neutrons by at least a factor of 2. The SNS should have a higher 8.9 Å flux than the NCNR, due in large part to the use of supermirror guides with 2.6 times the Fermi potential of ^{58}Ni . The proposed guide is 10 cm by 12 cm so the potential gain factor for a larger magnetic trap may be even greater (see section 7.8).

7.6 Increased Detection Efficiency

Increasing the overall detection efficiency can both increase signal and decrease backgrounds. As the detection efficiency is slowly increased, the number of detected events will increase at each threshold. As the detection efficiency saturates near 100% for the low thresholds, then the threshold can be raised without loss of signal. The effect of certain backgrounds, especially luminescence, decreases at higher thresholds. Other backgrounds, such as activation and the constant background decrease with threshold for a fixed detection efficiency. As the detection efficiency is raised, the relative size of these events compared with neutron events should remain constant. Nevertheless, even keeping the Signal/Background ratio fixed, but raising both by increasing the detection efficiency will increase the overall statistical figure of merit which is $\sqrt{\text{Signal}^2/\text{Background}}$. Therefore, increasing the overall detection efficiency is very desirable and should be an important goal of future work.

It may be possible to raise the light collection efficiency (that is the fraction of photons produced within the cell reaching the detectors) by a factor of 3, since such an efficiency was measured earlier with the existing apparatus. Several changes which were made to improve the cryogenic performance of the apparatus led to decreased detection efficiency. These include the 4 K quartz windows (see section 4.2.4) and the 77 K lightguide (see section 4.2.4). A significant decrease in the detection efficiency is likely due to the quartz windows, which cracked during cooldown. Likewise, a 30% reduction in detection efficiency is due to the copper wires on the 77 K lightguide. Minor modifications to the quartz window design, to prevent cracking, and substitution of a sapphire window for the copper wires should lead to an increase of a factor of 2–3 in light collection efficiency.

Further increases in the light collection efficiency and decreases in the constant

background may be possible by switching from the current PMTs to different detectors. Several types of detectors, including Avalanche Photodiodes [143, 144], and MCP-based photomultipliers [145], exist which can be operated at low temperatures and high magnetic fields. Any such detector could be placed within the bore of the magnet, immediately following the beam stop (and perhaps a 4 K window). Acrylic lightguides would no longer be necessary for light transport. Their exclusion would reduce the constant backgrounds by up to a factor of 2, depending on detection threshold (see table 6.2). Light collection efficiency would also be increased by an estimated factor of 4 by the elimination of losses in the light guides and windows of the current collection optics.

7.7 Larger, Deeper Magnetic Trap

One possible way to significantly increase the statistical sensitivity of the experiment is to use a larger, deeper magnetic trap. The number of trapped neutrons scales roughly as $VB^{3/2}$ where V is the trap volume and B is the trap depth. One approach would be to improve the trap depth of the existing trap. This may be achievable simply by running the magnet at a higher current. The operating current of 170 A (far from the theoretical limit of 320 A) was set not by the limits of training, but by an assessment of the risk of catastrophic failure as a function of current, based upon the previous damage to the octupole trap (see section 4.3.3). Through better understanding of the quench dynamics and quench protection it may be possible to operate the existing magnet at a higher current.

Alternatively it may be possible to wind new coils to use on the existing magnet form. New coils would be wound with larger diameter wire, increasing the operating current and decreasing the magnet inductance. Since the operation of quench pro-

tection circuits for low inductance magnets is well studied, it should be possible to reduce the risk of a catastrophic failure. The training of a magnet wound from larger wire may also be better.

Perhaps much greater gains in trap depth can be achieved by building a magnetic trap using an existing, high field accelerator quadrupole. Quadrupole magnets are used for focusing and insertion of charged particles in high energy storage rings. Eight such accelerator quadrupoles, of an extremely attractive size and field gradient for this work, were constructed for use at the TRISTAN main ring at KEK [146]. We have borrowed and tested one of the TRISTAN quadrupoles (hereafter referred to as the “KEK magnet”). The KEK magnet has an inner diameter of 14 cm, an effective magnetic length of 1.14 m and an operating field gradient of 70 T/m when operated at the design current of 3450 A [146]. The KEK magnet is wound with individually insulated superconducting cables with 11% of the cable area open to liquid helium, which flows through it and provides cooling [147].

The performance of the magnet depends critically on the temperature and pressure of the helium. Typically it is operated using pressurized or superfluid helium to prevent bubbles in the liquid from limiting the efficiency of cooling. In unpressurized liquid helium, the load line of the quadrupole is reached at an operating current of 4090 A [148]. The field gradient can be maximized by operating in superfluid helium at a temperature of 1.8 K, at which an operating current of 5086 A was achieved [148]. With the KEK magnet operated at 5086 A, and for a reasonable cell design with a trap diameter of 11.5 cm, the field at the trap wall would be 5.9 T. Calculations of the magnetic field profile for a reasonable Ioffe trap geometry indicate that a trap depth of 4.4 T should be achievable for a trap of length 39 cm and 11.5 cm diameter. The number of UCN trappable in such a trap is dramatically higher than in our existing

trap, due largely to the factor of $4^{3/2} = 8$ gain from the increased trap depth. Monte Carlo simulations indicate that such a trap based on the KEK magnet would trap a factor of 20 times as many UCN as our current trap when operated on the permanent monochromatic beamline described above.

We have borrowed one of the spare KEK magnets and performed preliminary tests. We have constructed a quench protection circuit for the magnet and have operated the quadrupole at a current of 3000 A in 4.2 K unpressurized liquid helium. The magnet did not quench at this current, but the current leads used were only rated to operate at 3000 A. Although solenoids for constructing an Ioffe trap have been designed, they have not been constructed or tested. The performance of the trap assembly thus remains an open question. Calculations of the forces involved in such a trap assembly indicate that they are compatible with the construction of the KEK quadrupole. Furthermore, the KEK quadrupole has been run inside of a 1 T solenoid without incident [149]. In the long run, measurement of the neutron lifetime to an accuracy of less than one second will require use of a trap based on the KEK (or a similar) magnet.

7.8 Prospects for a Lifetime Measurement

In order to make a measurement of the neutron lifetime using the technique of magnetic trapping, statistical and systematic errors must be controlled to the desired level of accuracy. The known systematic errors, other than those related to the time-dependent background, should affect the measured neutron lifetime by less than one part in 10^5 [40, 42]. As discussed above, imperfect subtraction of the time-dependent background introduces a new and very troubling systematic error. In order to measure the neutron lifetime to sub-one second accuracy, it will be necessary

Table 7.1: The running time in days (T^*) to make a measurement of the neutron lifetime to a statistical accuracy of 10^{-3} is shown. For each of the experimental setups shown, the constant background rate, R_{bg} , and initial fluorine decay rate, R_F , are estimated from those observed in the current apparatus. The relative cold source brightness at 8.9 Å (compared to that of the NCNR cold source in 2001) is shown as ϵ_{CS} . The UCN production rate, P_{UCN} , is estimated from our Monte Carlo simulation (see appendix C). Detection efficiency ϵ_D is assumed to be 65% in the current apparatus (with 2 photoelectrons/PMT thresholds), while in all other apparatuses it is assumed to be 95% with 5 photoelectrons/detector thresholds.

Setup Considered	ϵ_{CS}	$P_{UCN}\epsilon_D$ [Hz]	R_{bg} [Hz]	R_F [Hz]	T^* [days]
Current Apparatus (8/01)	1.00	2.4	8.1	1300	4662
Current Apparatus (2/02)	1.82	10.7	5.4	2900	201
KEK Magnet	1.82	217.0	8.5	6500	2.0
Current Apparatus at SNS	N/A	13.1	5.4	6600	144
KEK Magnet at SNS	N/A	371.9	8.5	14300	0.9

to either eliminate or fully understand the imperfect subtraction of time dependent backgrounds. It seems that the preferable solution to this problem would be the elimination of the time-dependent backgrounds by using a separate production and detection region, as described in section 7.3.

Predictions of the achievable statistical uncertainty can be made when background subtraction works perfectly, using existing Monte Carlo code [134]. If we consider the existing apparatus as a baseline, for which the constant and time-dependent backgrounds are known as a function of threshold, we can estimate the backgrounds for various possible setups, as shown in table 7.1. We can also use our UCN production rate Monte Carlo (see appendix C) to estimate the number of trapped neutrons. We assume for all cases below (other than the present) that detection efficiencies of 95% can be achieved with thresholds of 5 photoelectrons per detector. The constant backgrounds are scaled based on the known constant background with thresholds of 5 photoelectrons/PMT and with helium present, and the estimated helium volume

of other traps. The time dependent background is assumed to be only that from activation of fluorine in the GoreTex and is scaled by the number of neutrons entering the cell. It is assumed in these simulations that the other time dependent backgrounds are eliminated through careful material selection and that separate production and detection regions are not used. The computed figure of merit is the number of days of running time (T^*) necessary to measure the neutron lifetime to a statistical accuracy of 10^{-3} .

Coupling the indicated magnets to an apparatus with separate production and detection regions should decrease T^* . In the case where the time-dependent background has been entirely eliminated, the experiment would measure a constant background and the neutron decay signal. In such a case, unless the neutron decay signal is substantially larger than the constant background, negative runs would not be needed³ [134]. Optimal fill and observation times are both longer, however, when background subtraction is not used [134, 136], partially offsetting the decrease in running time obtained by not taking negative runs. An additional duty cycle factor would be gained through the use of separate production and detection regions since the filling of the production region could occur simultaneously with the observation of decay events in the trapping region. In any event, it is expected that the values of T^* for an apparatus with separate production and detection regions should be at least as short as those shown in table 7.1.

³If the ratio of the constant background to the UCN production rate is less than 0.032 then it becomes optimal to take a small number of negative runs and perform background subtraction rather than fitting the data to a three parameter model.

Appendix A

Activation

When the amount, composition, and geometry of a material are all known and that material will be exposed to a known neutron field for a known amount of time, the resulting neutron activation may be easily calculated. In general, a specific model of the system must be generated. Herein, however, we will provide some general calculations which may be of use in many situations.

In general, one can calculate the absorption length for neutrons, l_{abs} , for a material as

$$\frac{1}{l_{abs}} = \rho \left[\frac{\text{g}}{\text{cm}^3} \right] \sum_i \sigma_i(\lambda) [\text{cm}^2] X_i \frac{6.022 \times 10^{23} \text{atoms}}{A_i [\text{g}]} \quad (\text{A.1})$$

where $\sigma_i(\lambda)$, X_i , and A_i are the cross section (as a function of neutron wavelength), mass fraction, and atomic mass, respectively of the i th isotope in the material of mass density ρ . The fraction of neutrons transmitted through the material, with thickness, d , can then be calculated as

$$e^{-d/l_{abs}} \quad (\text{A.2})$$

Let us now consider the case of a compound of known isotopic composition, which overall is weakly neutron absorbing. As described in Chapter 2, this requirement

characterizes all of the materials in our apparatus which are exposed to neutrons, other than the neutron shielding materials. In the case where the entire material has weak absorption, then we can make several approximations to (A.1) and (A.2). Since the overall absorption of the material is weak, then by definition $d \ll l_{abs}$ so we can write the fraction of neutrons absorbed by the material as

$$\text{fraction absorbed} = 1 - e^{-d/l_{abs}} \approx \frac{d}{l_{abs}} = \sum_i d\sigma_i(\lambda)\rho X_i \frac{6.022 \times 10^{23}}{A_i}. \quad (\text{A.3})$$

Once we have made this approximation, then it become reasonable also to consider the activation separately for the each of the isotopes. The inverse dependence of the neutron cross sections with neutron velocity can be simply taken into account by scaling the tabulated thermal neutron cross sections, σ_{th} , by the weighting factor appropriate for the mean neutron wavelength. For our monochromatic 12 K beam this weighting is just

$$\sigma(\lambda) = \sigma_{th} \sqrt{\frac{300}{12}} = 5\sigma_{th} \quad (\text{A.4})$$

Then we can simply consider that for each isotope, the rate of neutron absorptions, R , is

$$R = \Phi S 5\sigma_{th} \rho X \frac{6.022 \times 10^{23}}{A} d \quad (\text{A.5})$$

where Φ is the average flux and S the cross sectional area of material exposed to that flux (or alternatively ΦS is just the total rate of neutrons hitting the material). We can simplify this expression by noting that $Sd\rho = m$ where m is the mass of material exposed to the neutron beam. We also recall that after exposure to a neutron beam for a time t_{fill} , the number of activated atoms, N for an isotope whose decay product has lifetime τ is

$$N = R\tau(1 - e^{-t_{fill}/\tau}) \quad (\text{A.6})$$

Plugging everything in, we can write the initial decay rate, \dot{N} , per gram of material

per 10^{-6} fraction of isotope as

$$\frac{\dot{N}}{\mu\text{g}} = (3.011 \times 10^{18}) \frac{\Phi \sigma_{th}}{A} (1 - e^{-t_{fill}/\tau}) \quad (\text{A.7})$$

For the monochromatic neutron beam, the flux has been measured to be approximately 10^7 n/cm²/s (see section 3.9.2). Based upon that flux and the filling time of 2100 s which was used in the experiment, Table A.1 shows the expected initial decay rate per microgram of any given element attributable to each of its stable isotopes. That is, for an isotope which has a fractional natural abundance of Y , the tabulated initial decay rate is

$$\frac{\dot{N}}{\mu\text{g}} = (0.3011) \frac{\sigma_{th}[\text{barns}] Y [\%]}{A} (1 - e^{-2100/\tau}) \quad (\text{A.8})$$

Isotopes which produce another stable isotope when they absorb a neutron do not lead to activation.

Table A.1: Neutron activation properties of selected stable isotopes are tabulated. For each isotope, its natural abundance, Y , the thermal neutron cross section, σ_{th} , the decay lifetime, τ , of the activated product, the decay mode, and the initial decay rate per microgram calculated using (A.8) are shown. All of the data are taken from [67]. Isotopes which produce another stable isotope when they absorb a neutron are not included, since they do not result in activation. In the case of isotopes which have multiple possible products with different cross sections and lifetimes, each product is shown separately.

Isotope	Y [%]	σ_{th} [barns]	τ [s]	Decay Mode	$\dot{N}/\mu\text{g}$
H 2	0.02	5.20E-04	5.61E+08	18.6 keV β^-	8.79E-12
Be 9	100.00	8.00E-03	6.83E+13	556 keV β^-	8.22E-13
C 13	1.10	1.40E-03	2.60E+11	157 keV β^-	3.12E-13
N 15	0.37	2.00E-05	1.03E+01	β^- -s, γ s	1.59E-07
O 18	0.20	1.60E-04	3.88E+01	β^- -s, γ s	6.02E-07
F 19	100.00	9.40E-03	1.59E+01	5.4 MeV β^- , γ s	1.49E-02
Ne 22	9.25	5.00E-02	5.37E+01	4.4 MeV β^- , γ s	6.90E-03
Na 23	100.00	5.30E-01	7.77E+04	1.4 MeV β^- , γ s	1.85E-02

Isotope	Y [%]	σ_{th} [barns]	τ [s]	Decay Mode	$\dot{N}/\mu\text{g}$
Mg 26	11.01	3.82E-02	8.18E+02	1.8 MeV $\beta-$, γ s	4.81E-03
Al 27	100.00	2.31E-01	1.95E+02	2.9 MeV $\beta-$, γ	2.58E-01
Si 30	3.10	1.07E-01	1.36E+04	1.5 MeV $\beta-$, γ	5.08E-04
P 31	100.00	1.70E-01	1.78E+06	1.7 MeV $\beta-$	1.95E-04
S 34	4.21	2.90E-01	1.09E+07	167 keV $\beta-$	2.21E-06
36	0.02	2.00E-01	4.37E+02	1.8 MeV $\beta-$, γ s	3.72E-05
Cl 35	75.77	4.36E+01	1.37E+13	709 keV $\beta-$	4.30E-09
37	24.23	4.30E-01	3.22E+03	4.9 MeV $\beta-$, γ s	4.24E-02
Ar 36	0.34	5.20E+00	4.36E+06	electron capture	6.36E-06
38	0.06	8.00E-01	1.22E+10	565 keV $\beta-$	6.52E-11
40	99.60	6.50E-01	9.50E+03	1.2 MeV $\beta-$, γ s	9.67E-02
K 41	6.73	1.46E+00	6.42E+04	3.5 MeV $\beta-$, γ s	2.44E-03
Ca 40	96.94	4.10E-01	4.69E+12	electron capture	1.34E-10
44	2.09	8.40E-01	2.03E+07	358 keV $\beta-$, γ s	1.36E-06
46	0.004	7.00E-01	5.65E+05	694 keV $\beta-$, γ s	7.80E-08
48	0.19	1.10E+00	7.55E+02	2.2 MeV $\beta-$, γ s	1.45E-03
Sc 45	100.00	2.70E+01	1.04E+07	357 keV $\beta-$, γ s	3.63E-03
Ti 50	100.00	1.77E-01	4.99E+02	2.1 MeV $\beta-$, γ s	1.10E-01
V 51	99.75	4.90E+00	3.25E+02	2.5 MeV $\beta-$, γ s	2.88E+00
Cr 50	4.35	1.58E+01	3.45E+06	320 keV γ	2.42E-04
54	2.36	3.60E-01	3.03E+02	2.5 MeV $\beta-$, γ s	4.91E-03
Mn 55	100.00	1.33E+01	1.34E+04	2.8 MeV $\beta-$, γ s	1.06E+00
Fe 54	5.85	2.30E+00	1.24E+08	electron capture	1.23E-06
58	0.28	1.30E+00	5.55E+06	466 keV $\beta-$, γ s	7.43E-07
Co 59	100.00	2.10E+01	9.06E+02	1.6 MeV $\beta-$, γ s	9.67E+00
59	100.00	3.70E+01	2.40E+08	318 keV $\beta-$, γ s	1.66E-04
Ni 58	68.08	4.60E+00	3.46E+12	electron capture	9.76E-10
62	3.63	1.45E+01	4.55E+09	67 keV $\beta-$	1.25E-07
64	0.93	1.80E+00	1.31E+04	2.1 MeV $\beta-$, γ s	1.27E-03
Cu 63	69.17	4.50E+00	6.60E+04	578 keV $\beta-$, γ	4.62E-02
63	69.17	5.00E+00	6.60E+04	$\beta+$	5.13E-02
65	30.83	2.20E+00	3.24E+05	2.6 MeV $\beta-$, γ s	2.08E-03

Isotope	Y [%]	σ_{th} [barns]	τ [s]	Decay Mode	$\dot{N}/\mu\text{g}$	
Zn	64	48.60	7.60E-01	3.04E+07	1.1 MeV γ	1.18E-05
	64	48.60	1.50E+00	3.04E+07	$\beta+$	2.36E-05
	68	18.80	8.00E-01	4.81E+03	0.9 MeV $\beta-$, γs	2.45E-02
	70	0.60	8.30E-02	2.16E+02	2.6 MeV $\beta-$, γs	2.29E-04
Ga	69	60.11	1.70E+00	1.83E+03	1.7 MeV $\beta-$, γs	3.01E-01
	71	39.89	4.80E+00	7.48E+04	1.0 MeV $\beta-$, γs	2.29E-02
Ge	70	21.23	3.00E+00	1.42E+06	electron capture	3.90E-04
	74	35.94	1.40E-01	6.92E+01	$\beta-, \gamma$	2.09E-02
	74	35.94	4.40E-01	7.17E+03	1.2 MeV $\beta-, \gamma\text{s}$	1.67E-02
	76	7.44	9.00E-02	7.65E+01	2.9 MeV $\beta-, \gamma\text{s}$	2.78E-03
	76	7.44	1.50E-01	5.87E+04	2.2 MeV $\beta-, \gamma\text{s}$	1.63E-04
As	75	100.00	4.20E+00	1.37E+05	3.0 MeV $\beta-, \gamma\text{s}$	2.58E-02
Se	74	0.89	4.80E+01	1.49E+07	γs	2.29E-05
	80	49.60	5.00E-02	4.96E+03	$\beta-, \gamma\text{s}$	3.26E-03
	80	49.60	4.00E-01	1.60E+03	1.6 MeV $\beta-, \gamma\text{s}$	5.53E-02
	82	8.73	3.90E-02	1.01E+02	2.9 MeV $\beta-, \gamma\text{s}$	1.30E-03
	82	8.73	4.40E-02	1.93E+03	1.0 MeV $\beta-, \gamma\text{s}$	9.71E-04
Br	79	50.69	1.08E+01	1.53E+03	2.0 MeV $\beta-, \gamma\text{s}$	1.54E+00
	81	49.61	2.40E+00	5.28E+02	$\beta-, \gamma\text{s}$	4.40E-01
	81	49.61	2.65E+00	1.83E+05	444 keV $\beta-, \gamma\text{s}$	5.64E-03
Kr	78	0.35	6.17E+00	1.81E+05	$\beta+, \gamma\text{s}$	8.93E-05
	80	2.25	1.26E+01	9.52E+12	γ	2.22E-11
	84	57.00	9.00E-02	2.33E+04	840 keV $\beta-, \gamma\text{s}$	1.59E-03
	84	57.00	2.00E-02	4.90E+08	687 keV $\beta-, \gamma\text{s}$ 1	1.76E-08
	86	17.30	3.00E-03	6.60E+03	3.5 MeV $\beta-, \gamma\text{s}$	5.08E-05
Rb	85	72.16	5.00E-01	2.32E+06	1.8 MeV $\beta-, \gamma\text{s}$	1.15E-04
Sr	84	0.56	6.00E-01	5.85E+03	$\beta-, \gamma\text{s}$	3.48E-04
	84	0.56	8.00E-01	8.08E+06	γs	4.00E-07
	88	82.58	5.80E-03	6.30E+06	1.5 MeV $\beta-, \gamma\text{s}$	5.49E-07
Y	89	100.00	1.00E-03	1.66E+04	$\beta-, \gamma\text{s}$	4.03E-05
	89	100.00	1.28E+00	3.33E+05	2.3 MeV $\beta-, \gamma\text{s}$	2.73E-03
Zr	94	17.00	5.00E-02	7.98E+06	368 keV $\beta-, \gamma\text{s}$	7.38E-07

Isotope	Y [%]	σ_{th} [barns]	τ [s]	Decay Mode	$\dot{N}/\mu\text{g}$		
Nb	96	2.80	2.20E-02	9.36E+04	1.9 MeV $\beta-$, γs	4.51E-06	
	93	100.00	9.00E-01	5.39E+02	1.2 MeV $\beta-$, γs	2.86E-01	
Mo	98	24.13	1.27E-01	3.42E+05	1.2 MeV $\beta-$, γs	5.90E-05	
	100	9.63	1.90E-01	1.26E+03	0.7 MeV $\beta-$, γs	4.65E-03	
Ru	96	5.52	2.50E-01	3.60E+05	γs	2.39E-05	
	102	31.60	1.30E+00	4.89E+06	223 keV $\beta-$, γs	5.25E-05	
	104	18.70	4.70E-01	2.31E+04	1.2 MeV $\beta-$, γs	2.28E-03	
Rh	103	100.00	8.00E+01	6.10E+01	2.4 MeV $\beta-$, γs	2.34E+01	
Pd	102	1.02	3.00E+00	2.12E+06	γs	8.58E-06	
	108	26.46	8.00E+00	6.99E+04	1.0 MeV $\beta-$, γs	1.77E-02	
	110	11.72	3.00E-02	2.86E+04	350 keV $\beta-$, γs	7.05E-05	
	110	11.72	7.00E-01	2.03E+03	2.1 MeV $\beta-$, γs	1.50E-02	
	Ag	107	51.84	3.60E+01	2.07E+02	1.7 MeV $\beta-$, γs	5.21E+00
		109	48.16	1.38E+03	3.55E+01	1.0 MeV $\beta-$, γs	1.86E+02
Cd	106	1.25	2.00E-01	3.39E+04	$\beta+$, γs	4.03E-05	
	110	12.49	6.00E-02	4.20E+03	γs	7.90E-04	
	114	28.73	3.60E-01	2.78E+05	1.1 MeV $\beta-$, γs	2.09E-04	
	116	7.49	5.20E-02	1.29E+04	670 keV $\beta-$, γs	1.56E-04	
	In	113	4.29	5.00E+00	1.04E+02	2.0 MeV $\beta-$, γs	5.62E-02
		115	95.71	7.50E+01	4.69E+03	1.0 MeV $\beta-$, γs	6.79E+00
		115	95.71	4.10E+01	2.03E+01	3.3 MeV $\beta-$, γs	1.03E+01
Sn	112	0.97	5.60E-01	1.43E+07	γs	2.02E-07	
	120	32.59	1.60E-01	1.41E+05	383 keV $\beta-$	1.96E-04	
	122	4.63	1.50E-01	3.47E+03	1.3 MeV $\beta-$, γs	8.00E-04	
	124	5.79	1.30E-01	8.22E+02	2.0 MeV $\beta-$, γs	1.76E-03	
	124	5.79	1.30E-01	1.20E+06	2.4 MeV $\beta-$, γs	3.33E-06	
Sb	121	57.30	1.00E-01	3.63E+02	$\beta-$, γs	1.41E-02	
	121	57.30	6.00E+00	3.39E+05	1.4 MeV $\beta-$, γs	5.25E-03	
	123	42.70	4.00E-02	1.38E+02	1.2 MeV $\beta-$, γs	4.22E-03	
	123	42.70	4.00E+00	7.50E+06	610 keV $\beta-$, γs	1.18E-04	
Te	120	0.10	2.00E+00	2.09E+06	γs	4.54E-07	
	126	18.95	1.30E-01	1.36E+07	$\beta-$, γs	8.98E-07	
	126	18.95	9.00E-01	4.88E+04	690 keV $\beta-$, γs	1.69E-03	

Isotope	Y [%]	σ_{th} [barns]	τ [s]	Decay Mode	$\dot{N}/\mu\text{g}$
128	31.69	1.50E-02	4.19E+06	1.6 MeV $\beta-$, γs	5.62E-07
128	31.69	2.00E-01	6.02E+03	1.5 MeV $\beta-$, γs	4.40E-03
130	33.80	2.00E-02	1.68E+05	420 keV $\beta-$, γs	1.98E-05
130	33.80	1.80E-01	2.16E+03	2.1 MeV $\beta-$, γs	8.92E-03
I 127	100.00	6.20E+00	2.16E+03	2.1 MeV $\beta-$, γs	9.14E-01
Xe 124	0.10	1.68E+02	8.88E+04	$\beta+$, γs	9.00E-04
126	0.09	4.00E-01	9.95E+01	$\beta-$, γs	8.26E-05
126	0.09	3.00E+00	4.54E+06	γs	2.87E-07
132	26.90	4.00E-01	6.54E+05	346 keV $\beta-$, γs	7.92E-05
134	10.40	2.60E-01	4.73E+04	910 keV $\beta-$, γs	2.70E-04
136	8.90	2.60E-01	3.31E+02	4.1 MeV $\beta-$, γs	5.30E-03
Cs 133	100.00	2.70E+01	9.40E+07	658 keV $\beta-$, γs	1.37E-04
Ba 130	0.11	8.00E+00	1.44E+06	$\beta+$, γs	2.71E-06
132	0.10	8.00E-01	2.02E+05	$\beta-$, γs	1.83E-06
135	6.59	1.40E-02	4.44E-01	γs	2.02E-04
138	71.70	4.30E-01	7.25E+03	2.3 MeV $\beta-$, γs	1.70E-02
La 139	99.91	9.00E+00	2.09E+05	1.4 MeV $\beta-$, γs	1.95E-02
Hf 178	27.30	3.00E+01	2.69E+01	γs	1.38E+00
179	13.63	4.30E-01	2.86E+04	$\beta-$, γs	7.01E-04
180	35.10	1.30E+01	5.29E+06	405 keV $\beta-$, γs	3.06E-04
Ta 181	99.99	1.20E-02	1.37E+03	γs	1.57E-03
W 182	26.50	2.00E+01	7.43E+00	γs	8.68E-01
184	30.64	2.00E-03	1.42E+02	γs	1.00E-04
184	30.64	1.80E+00	9.32E+06	433 keV $\beta-$, γs	2.03E-05
186	28.43	3.80E+01	1.24E+05	622 keV $\beta-$, γs	2.97E-02
Re 185	37.40	1.12E+02	4.63E+05	1.1 MeV $\beta-$, γs	3.06E-02
187	62.60	2.80E+00	1.61E+03	$\beta-$, γs	2.07E-01
187	62.60	7.30E+01	8.80E+04	1.0 MeV $\beta-$, γs	1.74E-01
Os 184	0.02	3.00E+03	1.17E+07	γs	1.71E-05
189	16.10	2.60E-04	8.57E+02	γs	6.05E-06
190	26.40	1.30E+01	1.92E+06	143 keV $\beta-$, γs	5.94E-04
192	41.00	2.00E+00	1.58E+05	1.1 MeV $\beta-$, γs	1.71E-03

Isotope	Y [%]	σ_{th} [barns]	τ [s]	Decay Mode	$\dot{N}/\mu\text{g}$	
Ir	191	37.30	2.60E+02	1.25E+02	1.2 MeV $\beta-$, γ s	1.52E+01
	191	37.30	9.20E+02	9.20E+06	672 keV $\beta-$, γ s	1.23E-02
	193	62.70	1.06E+02	9.95E+04	2.2 MeV $\beta-$, γ s	2.18E-01
Pt	190	0.01	1.50E+02	3.69E+05	γ s	1.31E-05
	192	0.79	1.00E+00	5.40E+05	$\beta-$, γ s	4.74E-06
	194	32.90	1.00E-01	5.01E+05	$\beta-$, γ s	2.12E-05
	196	25.30	4.00E-02	8.26E+03	$\beta-$, γ s	3.51E-04
	196	25.30	6.00E-01	9.50E+04	642 keV $\beta-$, γ s	5.12E-04
	198	7.20	3.50E+00	2.67E+03	1.7 MeV $\beta-$, γ s	2.12E-02
Au	197	100.00	9.87E+01	3.36E+05	1.0 MeV $\beta-$, γ s	9.40E-02
Hg	196	0.15	1.20E+02	1.24E+05	$\beta-$, γ s	4.55E-04
	196	0.15	3.10E+03	3.33E+05	γ s	4.39E-03
	202	29.86	4.90E+00	5.81E+06	213 keV $\beta-$, γ s	7.94E-05
	204	6.87	4.00E-01	4.50E+02	1.6 MeV $\beta-$, γ s	4.09E-03
Tl	203	29.52	1.14E+01	1.72E+08	763 keV $\beta-$	6.05E-06
	205	70.48	1.00E-01	3.25E+02	1.5 MeV $\beta-$, γ s	1.04E-02
Pb	208	52.40	4.10E-04	1.69E+04	645 keV $\beta-$	3.65E-06
Bi	209	100.00	2.40E-02	6.24E+05	1.2 MeV $\beta-$, γ s	1.16E-05

Appendix B

Dewar Boiloff

High helium boiloff has been a persistent problem for our apparatus. Due to the high cost of liquid helium (~ 2.65 USD/L), helium consumption is a significant operating cost for the experiment. Additionally, the helium boiloff rate sets the time between helium transfers. When the helium level in the main bath drops below the level of the 1K pot pickup line then the 1K pot will no longer fill and the refrigerator cannot be operated. Hence the “hold time,” or time between helium fills, is equal to the time to boiloff the volume of helium (~ 60 L) above the minimum pickup level. Frequent helium transfers can decrease the duty cycle of the experiment as well as consume manpower.

During operation in 1998 and 1999 [42] helium consumption was approximately 150 L/day. In order to reduce the boiloff, a series of test measurements were performed. In these measurements, no windows were present in the horizontal section. From these measurements we can determine how much boiloff is due to various components of the apparatus (see Table B.1). The dewar was assembled with or without the vapor cooled current leads (and their associated copper/NbTi bus bar which extends down into the liquid helium) and with or without the dilution refrigerator.

Table B.1: The helium boiloff in liters/day is shown for various configurations. In all cases the dewar was assembled without any windows. When the refrigerator was present, the 1K pot was operated. Boiloff was calculated from the change in the helium level with time. Additional helium consumed due to inefficiencies during transfer are not considered. All boiloff rates are shown in L/day ± 3 L/day (due to uncertainties in the level measurement and the volume as a function of level).

Configuration	Transfer Tube	Boiloff
Dewar only	out	8
Dewar + current leads. Boiloff \rightarrow bypass	out	35
Dewar + current leads. Boiloff \rightarrow bypass + leads	out	30
Dewar + current leads. Boiloff \rightarrow leads	out	31
Dewar + fridge. Level above transfer tube	in	79
Dewar + fridge. Level below transfer tube	in	15
Dewar + fridge + current leads. Level above tube	in	101
Dewar + fridge + current leads. Level below tube	in	72
Dewar + fridge. Level above transfer tube	in	121
Dewar + fridge. Level below transfer tube	in	16
Dewar + fridge	out	19
Dewar + fridge + current leads	out	50

Boiloff could be routed either through the vapor cooled leads (if present) or through a bypass line. The end of the bypass line was equipped with a 1-way valve that maintained about 1 psi overpressure inside the main bath. The primary conclusion of these measurements was that at least 30 L/day of boiloff was due to the transfer tube.

Helium is transferred into the apparatus using a three part transfer tube. A long, flexible transfer tube inserts into liquid helium storage dewars and extends to the top of the main cryostat. A short L-shaped tube couples to the long section and inserts into the top of the cryostat, coupling to the main bath. Both of these transfer tubes are removable at any time. A third, fixed, stainless steel tube extends from the top of the IVC down into the horizontal section of the apparatus. Teflon tubing extends the stainless steel tube down to the bottom of the horizontal section, underneath




Figure B.1: The helium level within the main bath is shown as a function of time. The apparatus was assembled with the dilution refrigerator inside of the dewar. No current leads were present. The boiloff is higher when the level is above 65%, the position of the bottom of the helium transfer tube.

the magnet. When all three sections of the transfer tube are joined, helium may be transferred to the very bottom of the cryostat. This allows efficient cooling of the magnet using cold helium vapor as well as liquid. The long flexible section is decoupled from the other two sections when helium is not being transferred. The L-shaped section can either be left in place or removed between transfers.

As shown in Table B.1, there is a significant decrease in the boiloff when the L-shaped helium transfer tube is removed between transfers. In measurements where the transfer tube was left in place, two boiloff rates were observed. Figure B.1 shows the helium level as a function of time. A sharp change in the slope occurs at a level of 65% from a boiloff rate of 79 L/day to 15 L/day. This level corresponds to the position of the bottom end of the L-shaped helium transfer tube. Removal of the helium transfer

tube between transfers eliminated this change in boiloff rate, reducing the boiloff to at most what it had been below the transfer tube. It is believed that oscillations in the helium gas within the transfer tube conducted significant amounts of heat down into the helium bath, by the well-known mechanism of Taconis oscillations [150].

After the excess boiloff due to the transfer tube was identified and rectified, the dewar was assembled with all of the windows and neutron shielding in place (see chapter 4). With these components in place, the boiloff rose significantly from the 50 L/day measured without them. During April of 2001, with the acrylic lightguide clamped only by the copper clamp (see section 4.2.4) the total helium consumption averaged over 1 week of running was 158 L/day. After a second clamp (aluminum/nylon) was introduced to lower the temperature at the center of the lightguide further, so that the base temperature of the cell could be lowered, and a 1 mm thick copper shield was coupled to the 4 K end flange (to bring the edge of the 4 K acrylic window into thermal contact with the helium bath), the helium consumption rose to 225 liters per day.

After the horizontal nitrogen shield was directly cooled by liquid nitrogen and wires were embedded in the cold face of the 77 K lightguide, the average helium consumption over 14 days of running was 131 L/day. Hence, a huge source the helium boiloff was due to the blackbody radiation produced by the hot nitrogen end flange and lightguide. The importance of having liquid nitrogen in all parts of the nitrogen shield, when the shield may be subject to a large heat load, should not be forgotten. In constructing any future apparatus, care should be taken to ensure liquid nitrogen fills all parts of the nitrogen shield.

All of the helium consumption figures include helium which is lost in cooling the transfer tube (5–10 Liters/transfer). No attempt was made to optimize the temper-

ature of the 77 K flange on the neutron beam entrance side of the apparatus. Hence it may be possible to lower the helium consumption further by making additional changes to this end of the apparatus.

Appendix C

Neutron Beam Simulations

C.1 Diffraction Theory

The Bragg condition for diffraction of neutrons with a wavelength λ from a crystal with lattice spacing d is

$$n\lambda = 2d \sin \theta \quad (\text{C.1})$$

where n is the “order” of the diffraction peak and θ is the grazing angle. Differentiating both sides of (C.1) and then dividing through by (C.1) yields the relationship between the angular spread and wavelength spread of reflected neutrons:

$$\frac{d\lambda}{\lambda} = \frac{d\theta}{\tan \theta} \quad (\text{C.2})$$

The Bragg condition can also be written vectorially as:

$$\vec{k}_f = \vec{k}_i + \vec{G} \quad (\text{C.3})$$

where \vec{k}_i and \vec{k}_f are the initial and final wavevectors of the diffracted neutron and

$$\vec{G} = \frac{2\pi}{d} \hat{n}_M \quad (\text{C.4})$$

is the inverse lattice vector of the crystal with lattice spacing d and orientation \hat{n}_M , as defined by the inwardly pointing unit normal to the crystal. Since we are considering purely elastic scattering, conservation of momentum requires that $|\vec{k}_f| = |\vec{k}_i|$. The Bragg condition thus specifies which orientations \hat{n}_M of perfect crystals will allow diffraction. The Bragg condition can also be written as

$$-\frac{\pi}{d} = \hat{n}_M \cdot \vec{k}_i \quad (\text{C.5})$$

where d is the spacing between crystal planes, and \vec{k}_i is the initial wavevector of the neutron.

The above discussion describes the diffraction of neutrons from perfect crystals under the kinematic theory. Many crystals are in fact mosaic crystals, composed of a large number of small perfect crystal “mosaic blocks.” Each mosaic block has inwardly pointing normal \hat{n}_M (see figure C.5). The distribution of orientations of microcrystals within the mosaic crystal with overall orientation \hat{n}_0 , is a cylindrically symmetric Gaussian:

$$W(\Delta) = W_{norm} e^{-\Delta^2/2\eta^2} \quad (\text{C.6})$$

where $\cos \Delta = |\hat{n}_M \cdot \hat{n}_0|$ and η sets the width of the distribution (and is, by definition, 0.424 times the full width at half maximum of the Gaussian distribution, γ). It is important to note that γ is an intrinsic parameter of the crystal and is not usually measured directly by neutron scattering measurements. Because of the presence of extinction and absorption (see below) the “apparent mosaic,” β , which is measured in neutron rocking curve experiments is usually larger than γ .

The normalization constant W_{norm} is chosen so that the total probability of finding \vec{n}_M is unity (integrating over all directions (Δ, χ) where Δ is the azimuthal angle and




Figure C.1: Geometric construction of the set of vectors \vec{G} which will allow diffraction of a neutron with incident momentum \vec{k}_i . The set of vectors \vec{G} which satisfy the Bragg condition is obtained by constructing lines from \vec{k}_i which are tangent to the sphere of radius $|\vec{G}|/2$.

χ is the other angle):

$$W_{norm} = \frac{1}{\int_0^{2\pi} d\chi \int_0^{\pi} e^{-\Delta^2/2\eta^2} \sin \Delta d\Delta} \quad (\text{C.7})$$

This integration is performed numerically for a given η by the Monte Carlo to yield W_{norm} .

Working in phase space, one can geometrically construct the set of crystal orientations, \hat{n}_M , for which an incident neutron with wavevector \vec{k}_i will diffract (see figure C.1). For diffraction of a given order n from a mosaic crystal with a given d -spacing, the set of all possible inverse lattice vectors, \vec{G} , lie on the surface of a sphere in momentum space. The set of vectors \vec{G} which satisfy the Bragg condition is obtained by constructing lines from \vec{k}_i which are tangent to the sphere of radius $|\vec{G}|/2$. The intersection of these lines with the sphere maps out a circle on the surface of the sphere. Orientations of $\vec{G}/2$ which fall on this circle satisfy the Bragg condition




Figure C.2: The crystal orientations \hat{n}_M in a mosaic crystal form a cone of opening angle α around the incident momentum \vec{k}_i .

and form a cone of opening angle

$$\alpha = \cos^{-1} \left(\frac{\lambda}{2d} \right) \quad (\text{C.8})$$

around the incident momentum \vec{k}_i . In a coordinate system (x''', y''', z''') where $\vec{k}_i/|\vec{k}_i| = \hat{z}'''$ the crystal orientations which satisfy the Bragg condition are

$$\hat{n}_m''' = (\sin \alpha \cos \phi, \sin \alpha \sin \phi, \cos \alpha) \quad (\text{C.9})$$

where ϕ parameterizes the circle (see figure C.2).

Following [90], section IV.3, we can consider the probability of a neutron reflecting from the mosaic crystal. We consider that a mosaic crystal is made up of many small, perfect crystals of thickness t_0 and orientation \hat{n}_M with a distribution W given by (C.6). We first consider the reflectivity of a single layer of perfect crystal blocks. Calling the reflectivity of each block $f(\vec{n}_M) = f(\Delta, \chi)$ then the reflecting power of a

layer which contains many blocks is

$$\int W(\Delta, \chi) f(\Delta, \chi) \sin \Delta d\Delta d\chi \quad (\text{C.10})$$

where the integration is performed over the surface of the \vec{n}_M sphere. Since the integration is being performed over the entire surface of the sphere, it can be done in any coordinate system we choose. Analysis of the problem becomes easier if we switch from the (Δ, χ) coordinate system to the α, ϕ coordinate system so that the z-axis is rotated and instead of $\hat{z}'' = \vec{n}_0$ we have $\hat{z}''' = \vec{k}_i/|\vec{k}_i|$. Then the integral (C.10) can be written

$$\int W(\alpha, \phi) f(\alpha, \phi) \sin \alpha d\alpha d\phi \quad (\text{C.11})$$

From the discussion above, $f(\alpha, \phi)$ is only non-zero for a circle on the surface of the sphere (these are the only orientations for which a small perfect crystal can Bragg reflect the incident neutron). At every point on the circle, $f(\alpha, \phi)$ should have the same value, since $f(\alpha, \phi)$ really depends only on $\vec{k}_i \cdot \vec{n}_m = \cos \alpha$. Since α is constant over the circle (and in fact defines the circle under consideration), then f has the same value at all points on the circle. Thus we can write

$$f(\alpha, \phi) = F \delta \left(\alpha - \cos^{-1} \left(\frac{\lambda}{2d} \right) \right) \quad (\text{C.12})$$

where F is a constant.

The integrated reflecting power R_H^θ is defined as the integral of the reflected power of a fixed small perfect crystal over all possible incident neutron orientations and over all outgoing directions [90]. Alternatively, R_H^θ is the reflecting power of a small perfect crystal as it is rocked through all values of α for a fixed incident neutron momentum. Hence

$$R_H^\theta = \int f(\alpha) d\alpha = F \quad (\text{C.13})$$

Thus instead of F we can write R_H^θ . Then we see that the integral (C.11) can be written as

$$R_H^\theta \int W(\alpha, \phi) \delta \left(\alpha - \cos^{-1} \left(\frac{\lambda}{2d} \right) \right) \sin \alpha d\alpha d\phi \quad (\text{C.14})$$

Now we see that the δ function reduces the integral to be over the circle on the surface of the sphere and the integral reduces to

$$R_H^\theta \cos \theta_B \int_0^{2\pi} W(\alpha(\phi)) d\phi \quad (\text{C.15})$$

where we have used the definition $\lambda = 2d \sin \theta_B$ to write $\sin \arccos \lambda/2d = \cos \theta_B$. The remaining integral is performed numerically by the Monte Carlo, as described below.

So, we have now calculated the reflecting power of a single layer, $R_H^\theta \cos \theta_B \bar{W}$ where \bar{W} is the numerical value of the integral of W over ϕ . We next consider the reflecting power of a thin layer of thickness $dT \gg t_0$ which contains many layers of perfect crystal blocks. We always assume that there is no phase relationship between the various mosaic blocks, so that intensities reflected from different blocks add. Then the reflecting power σ of a layer with thickness dT can be written (similar to [90], equation [4.20])

$$\sigma dT = \cos \theta_B \bar{W} \frac{R_H^\theta}{t_0} dT \quad (\text{C.16})$$

In the dynamical theory of neutron diffraction, one must consider the fact that as the neutron beam passes through the mosaic crystal and diffracts, the incoming beam is attenuated. Thus, the diffracted and incident beams at any depth are affected by the diffraction of the crystal at other depths. The effect of this multiple scattering of the neutron beam in a mosaic crystal leads to an effect called “secondary extinction” [151] which is the attenuation of the neutron beam due to scattering from many small perfect crystals. The process known as “primary extinction” is the attenuation of the

neutron beam within one of the small perfect crystals. Both of these extinctions lead to reductions in the power of the incident beam at any point in the crystal. Consideration of primary extinction affects the value calculated for R_H^θ while secondary extinction is important in determining the reflectivity from the bulk crystal.

The value of R_H^θ in the general non-absorbing case (no limit on thick or thin crystal) is

$$R_H^\theta = \frac{\Psi_B}{\sqrt{b} \sin 2\theta_B} \pi \tanh A \quad (\text{C.17})$$

where Ψ_B , \sqrt{b} and A are parameters of the theory of dynamical diffraction [90, 152], and π is the ratio of the circumference to the diameter of a circle. It is worth noting that A , in particular, is proportional to $1/\lambda$. In the case where $A \ll 1$, primary extinction is negligible and

$$R_H^\theta = Q \frac{t_0}{\gamma_0} \quad (\text{C.18})$$

The scattering power per unit length, Q , is defined as:

$$Q = \frac{\lambda^3 N_c^2 F^2}{\sin 2\theta_B} \quad (\text{C.19})$$

where N_c is the number of crystal cells per unit volume and F is the structure factor. From [90], Section IV.3C, the value of R_H^θ , in the case of primary extinction and small secondary extinction, is

$$R_H^\theta = Q' \frac{t_0}{\gamma_0} \equiv Q f(A_0) \frac{t_0}{\gamma_0} \quad (\text{C.20})$$

where $f(A_0)$ introduces the effect of primary extinction [90]. From the above equations we can see that

$$f(A_0) = \frac{\tanh A}{A} \quad (\text{C.21})$$

so that in the case of primary extinction and small secondary extinction the relevant value of R_H^θ is

$$R_H^\theta = Q \frac{t_0}{\gamma_0} \frac{\tanh A}{A} \quad (\text{C.22})$$

One can write ([90], equation 4.21) two differential equations which yield the differential change in the incident and diffracted powers (P_0 and P_H) from passing through a layer dT :

$$dP_0 = -\mu_0 P_0 \frac{dT}{\gamma_0} - \sigma P_0 dT + \sigma P_H dT \quad (\text{C.23})$$

and

$$dP_H = +\mu_0 P_H \frac{dT}{|\gamma_H|} + \sigma P_H dT - \sigma P_0 dT \quad (\text{C.24})$$

where μ_0 is the absorption length, and $\gamma_0 = |\hat{n}_0 \cdot \vec{k}_i/|k||$ and $\gamma_H = |\hat{n}_0 \cdot \vec{k}_f/|k||$ are direction cosines of the incident and diffracted wavevectors with respect to the overall crystal orientation. The first term in each of the differential equations represents loss of intensity due to absorption. The second term represents power loss due to diffraction. The third term represents power gained by diffraction of the other (incident or diffracted) beam. These two differential equations are solved, and two boundary conditions are imposed. The incident power $P_0(0)$ at $T = 0$ is assumed to be known while the power of the diffracted wave at the back of the crystal $P_H(T) = 0$ since that beam emerges from the front face and is generated inside the crystal. It is then found that the general solution to the above differential equations, with the specified boundary conditions allows one to write the reflected power ([90], equation 4.24):

$$\frac{P_H(0)}{P_0(0)} = \frac{\sigma + \frac{\mu_0}{\gamma_0} - U}{\sigma} - \frac{U \left[\sigma + \frac{\mu_0}{\gamma_0} - U \right] e^{-UT_0}}{\sigma \left[\left(\sigma + \frac{\mu_0}{\gamma_0} \right) \sinh UT_0 + U \cosh UT_0 \right]} \quad (\text{C.25})$$

where the symbol U is defined to be

$$U \equiv \sqrt{\left(\sigma + \frac{\mu_0}{\gamma_0} \right)^2 - \sigma^2}. \quad (\text{C.26})$$

In the case of negligible primary extinction (for very small t_0) and no absorption, the total reflecting power of the crystal reduces [151] to

$$R = \frac{\sigma T_0}{1 + \sigma T_0} \quad (\text{C.27})$$

where T_0 is the thickness of the entire crystal.

C.2 Monte Carlo Simulations

In order to study the production rate of UCN within the trapping region for various possible experimental configurations, a Monte Carlo simulation has been developed. This code simulates the neutron spectrum and divergence of NG-6, various wavelength filters (polycrystalline long-pass filters, monochromators, mechanical velocity selectors), neutron scattering from the helium in the trapping region, and production of UCN in the magnetic trap. Each part of the simulation will be described in turn.

C.2.1 Neutron Spectrum of NG-6

The neutron spectrum at the exit of the NG-6 guide is obtained from Monte Carlo simulations which include the luminosity of the cold source, the transmission of the ^{58}Ni coated waveguides, and the transmission of other materials in the beam between the reactor core and the experimental apparatus. These simulations were performed by Jeremy Cook of the NCNR and have been normalized to match experimental measurements of the neutron flux on NG-6 and other NCNR cold guides [155]. Relevant experimental measurements include absolute flux measurements and neutron spectra measured using time of flight techniques [42].

As described in Chapter 3, we have measured the absolute flux and time of flight (TOF) spectrum of NG-6. We can thus compare the flux and spectrum predicted by the Monte Carlo code with those which we have measured. In section 3.9.2, the capture flux at the exit of the NG-6 shielding, with 5 cm of polycrystalline beryllium and 10 cm of polycrystalline bismuth in the filter cryostat was measured to be $(9.0 \pm 0.5) \times 10^8 \text{ n cm}^{-2} \text{ s}^{-1}$. Using the Monte Carlo simulation, we predict a capture flux

of $(9.8 \pm 0.1) \times 10^8 \text{ n cm}^{-2} \text{ s}^{-1}$ at the exit of the NG-6 shielding ($z = 100 \text{ cm}$). In this calculation all neutrons with $\lambda < 3.95 \text{ \AA}$ are removed from the spectrum and measuring the flux over a $1 \text{ cm} \times 1 \text{ cm}$ square centered on the beamline. The calculated and measured fluxes at the end of the NG-6 shielding are therefore in reasonable agreement.

To convert the measured neutron TOF spectrum from counts into flux, we need to scale the measured spectrum by the flux divided by the total number of events in the measured spectrum. It is assumed for these calculations that the detector used for the TOF measurements has the same efficiency for detecting neutrons of all wavelengths¹. We therefore use the total flux for normalization, rather than the capture flux, which is weighted by $1/v$.

The measurement of the neutron TOF spectrum of NG-6 was done under slightly different circumstances than the flux measurement. First, the filter cryostat contained 10 cm of bismuth, but no beryllium. Second, the defining aperture of the chopper is at $z = 122 \text{ cm}$ beyond the end of the NG-6 guide. Taking into account these two changes, the Monte Carlo calculation predicts a total flux of $(4.3 \pm 0.1) \times 10^8 \text{ n cm}^{-2} \text{ s}^{-1}$ over the 3.8 mm diameter defining aperture for the TOF. It should be noted that this aperture was aligned with the NG-6 beam using a theodolite, and was centered on the beam to within 1 mm. The total neutron path length through air from the end of NG-6 to the TOF detector was 298 cm. Correcting for air scatter (estimated attenuation length of 30 m), the flux is reduced by a factor of 0.91 to $(3.9 \pm 0.1) \times 10^8 \text{ n cm}^{-2} \text{ s}^{-1}$. The neutron spectrum for the empty beam, measured using the TOF

¹This assumption is of questionable validity. Bicron specifies the detection efficiency of the BC-702 scintillator as 55% for 0.025 eV neutrons and 60% for 0.01 eV neutrons. The efficiency is unspecified for lower energy neutrons. Since our spectrum contains neutrons from 0.025 eV (1.8 Å) down to less than 0.001 eV (8.9 Å), it is not clear how the detection efficiency varies.




Figure C.3: The neutron spectra of NG-6 measured using neutron time of flight and simulated using the Monte Carlo code described in the text are compared. The spectra are for a 3.8 mm diameter aperture 122 cm from the end of the neutron guide. Neutrons pass through a filter cryostat containing 10 cm of 77 K polycrystalline bismuth prior to reaching the detector.

technique (see section 3.7 for details) and the spectrum predicted by the Monte Carlo simulations are compared in figure C.3. The spectra do not agree particularly well. The simulated spectrum has relatively more neutrons at longer wavelengths than the measured spectrum. This may be due in part to the fact that the simulation does not take into account the wavelength dependent transmission through the polycrystalline bismuth in the filter cryostat. Furthermore, the uncertainty in the efficiency of the TOF detector as a function of wavelength renders accurate comparisons between the two spectra difficult at best.

The flux at 8.9 Å implied by the measured spectrum is $9.7 \times 10^6 \text{ n cm}^{-2} \text{ s}^{-1} \text{ Å}^{-1}$ while the 8.9 Å flux calculated by the Monte Carlo is $1.6 \times 10^7 \text{ n cm}^{-2} \text{ s}^{-1} \text{ Å}^{-1}$. Hence the predicted production rate may be an overestimate by a factor of 1.65.

These Monte Carlo simulations do predict a UCN production rate for our original demonstration of magnetic trapping which is consistent with the measured production rate [42]. Since the simulations do not take into account the contribution of non-8.9 Å to UCN production, an overestimated 8.9 Å flux in the simulations might be compensated for by non-8.9 Å production in the trapping measurement. When using the monochromatic beam, such a happy coincidence would not be possible, due to the large reduction in 8.9 Å neutrons. The errors quoted on predicted UCN production rates do not take into account this uncertainty in the 8.9 Å flux, and rely upon the neutron spectrum generated by the Monte Carlo simulations. A more careful measurement of the 8.9 Å flux on the new, permanently installed monochromatic beam, should thus be a high priority. The Monte Carlo simulation must then be adjusted so that it is consistent with such a measurement.

C.2.2 Neutron Initial Conditions

Our simulations follow the trajectories of neutrons exiting from the end of NG-6 as they interact with various parts of the experimental apparatus. A large number of random trajectories are simulated using a Monte Carlo technique.

For each neutron under consideration we first determine its wavelength from the Monte Carlo neutron spectrum described above. The results of the Monte Carlo spectrum calculations described above were binned in 0.01 Å bins and provided to us [155]. This spectrum, $dJ(\lambda)$, is then normalized so that $\int_0^\infty dJ(\lambda)d\lambda = 1$. Our simulation code reads in a file containing $J(\lambda < x) = \int_0^x dJ(\lambda)d\lambda$. We can then generate random wavelengths with the proper spectrum by choosing a random number $y \in [0, 1]$ and setting the wavelength to x where $J(\lambda < x) = y$. Since $J(\lambda < x)$ is a discrete function rather than a continuous one, we use linear interpolation to obtain

a continuous spectrum of λ .

Once the neutron wavelength λ is fixed, all of the other conditions are also randomly determined. The neutron's initial position is chosen randomly over the 6 cm wide \times 15 cm tall face of the neutron guide. The neutron wavelength determines the magnitude of the wavevector

$$|\vec{k}| = \frac{2\pi}{\lambda} \quad (\text{C.28})$$

The direction of the wavevector is chosen within the constraints of the critical angle imposed by transmission down the guide tube. Only neutrons which graze the ^{58}Ni guides at angles less than

$$\theta_{max} = 0.002\lambda \quad (\text{C.29})$$

are transmitted down the neutron guide [154]. The neutron beam thus diverges from the end of NG-6 with maximum angle $\theta_{max}(\lambda)$. The coordinate system (x,y,z) is chosen so that \hat{z} is along the beam direction and \hat{y} is vertical, according to the Madison convention. The neutron is assigned a random direction up to the allowed divergence of the beam by randomly choosing the directional angles θ_x and θ_y (see figure C.4) within the range $[-\theta_{max}(\lambda), \theta_{max}(\lambda)]$. The directional components of the neutron wavevector are:

$$k_z = \frac{|\vec{k}|}{\sqrt{(\tan \theta_x)^2 + (\tan \theta_y)^2 + 1}} \quad (\text{C.30})$$

$$k_x = k_z \tan \theta_x \quad (\text{C.31})$$

$$k_y = k_z \tan \theta_y. \quad (\text{C.32})$$

This completely specifies the position and momentum of a random neutron at the end of the NG-6 neutron guide tube ($z = 0$).

Until the neutron interacts with some material, we can determine $x(z)$ and $y(z)$




Figure C.4: Momentum of a randomly chosen neutron in the (x, y, z) coordinate system.

from

$$x(z) = x(0) + \frac{k_x z}{k_z} z \quad (\text{C.33})$$

and

$$y(z) = y(0) + \frac{k_y z}{k_z} z. \quad (\text{C.34})$$

C.2.3 Reflection from a mosaic crystal

For simulations of setups containing monochromators, we must now handle diffraction from the monochromator. The geometrical arrangement of the crystal is shown in figure C.5. The crystal is set back a distance of z_{crys} from the end of NG-6, is aligned at an angle θ_0 with respect to the z axis and has size $x_{crys} \times y_{crys}$. The overall orientation of the crystal is chosen so that 8.9 Å neutrons with $\vec{k} = |\vec{k}|\hat{z}$ would Bragg reflect from a perfect crystal of that orientation. Thus the normal to the crystal




Figure C.5: Position and orientation of the monochromator with respect to the end of NG-6.

is $\hat{n}_0 = (\cos \theta_0, 0, -\sin \theta_0)$ where

$$\theta_0 = \sin^{-1} \left(\frac{8.9 \text{\AA}}{2d} \right). \quad (\text{C.35})$$

In the $x - z$ plane, the crystal defines a line

$$x(z) = -\frac{W_M}{2} + (z - z_{crys}) \tan \theta_0 \quad (\text{C.36})$$

where $W_M = x_{crys} \sin \theta_0$ is the width of the monochromator. We calculate $z = z_{cross}$ at which the neutron intersects the plane of the crystal as the intersection of the lines defined by (C.33) and (C.36) yielding

$$z_{cross} = \frac{x(0) + z_{crys} \tan \theta_0 + W_M/2}{\tan \theta_0 - \tan \theta_x} \quad (\text{C.37})$$

From z_{cross} , (C.33), and (C.34) we compute the position where the neutron intersects the monochromator, $(x_{cross}, y_{cross}, z_{cross})$. A neutron hits the crystal if it satisfies the three conditions $z_{cross} > z_{crys}$, $y_{cross} \in (-y_{crys}/2, y_{crys}/2)$ and $x \in (-W_M/2, W_M/2)$.




Figure C.6: Phase space diagram of reflection in the (x'', y'', z'') coordinate system. The (x'', y'', z'') coordinate system is obtained by rotating the (x, y, z) coordinate system by θ_0 about the y axis.

The Monte Carlo calculates the probability of reflection R and probabilistically determines the outgoing wavevector k_f for a given k_i . The function of the Monte Carlo simulation is then, following the discussion of section C.1, above, to numerically compute the integral of the gaussian probability distribution for the mosaic times the probability of reflection over the surface of the $\vec{G}/2$ sphere. If the neutron reflects, the direction into which it is reflected must then also be chosen from among the solutions to the Bragg equation, with the proper weighting.

For this calculation it is convenient to work in the coordinate system (x'', y'', z'') shown in figure C.6. The (x'', y'', z'') primed coordinate system is obtained by rotating the (x, y, z) coordinate system by θ_0 about the y axis. In the (x'', y'', z'') coordinate system, the orientation of the overall crystal is $\vec{n}_0 = (1, 0, 0)$. The mosaic orientations \vec{n}_M'' which satisfy the Bragg equation (C.5) form a cone of opening angle α around $-\vec{n}_k = -\vec{k}_i/|k|$ as described above.

Once we have randomly generated a \vec{k}_i , we calculate \vec{k}_i'' using the defining rotation

for the (x'', y'', z'') coordinate system:

$$\vec{k}''_i = \begin{pmatrix} k''_{ix} \\ k''_{iy} \\ k''_{iz} \end{pmatrix} = \begin{pmatrix} \cos \theta_0 & 0 & -\sin \theta_0 \\ 0 & 1 & 0 \\ \sin \theta_0 & 0 & \cos \theta_0 \end{pmatrix} \begin{pmatrix} k_{ix} \\ k_{iy} \\ k_{iz} \end{pmatrix}. \quad (\text{C.38})$$

The angle θ_B is calculated in terms of the wavelength and orientation of the neutron from

$$\theta_B = \sin^{-1} \left(\frac{\lambda}{2d} \sqrt{\tan^2 \theta_x + \tan^2 \theta_y + 1} \cos \theta_x \right) + \theta_x \quad (\text{C.39})$$

Let us consider the unit vector $\hat{n}_p = \vec{p}_i/|\vec{p}_i|$ as the direction of the incident momentum. In the (x'', y'', z'') coordinate system we can write the unit momentum vector as:

$$-\hat{n}''_p = (\sin \beta \cos \gamma, \sin \beta \sin \gamma, \cos \beta) \quad (\text{C.40})$$

where

$$\beta = \cos^{-1}(-n''_{pz}) \quad (\text{C.41})$$

and

$$\gamma = \tan^{-1} \left(\frac{n''_{py}}{n''_{px}} \right). \quad (\text{C.42})$$

We previously defined the coordinate system (x''', y''', z''') as the one where the z''' axis is parallel to \vec{n}_p . In equation (C.9) we specified in that coordinate system the mosaic orientations \hat{n}'''_m which satisfy the Bragg condition. These solutions form a cone of opening angle α (defined in equation (C.8)) around the momentum vector. The angle around the cone is parameterized as ϕ . To obtain the vector \hat{n}''_m we can rotate the (x''', y''', z''') coordinate system by $-\beta$ about the y''' axis and then by $-\gamma$ about the new z axis:

$$\hat{n}''_m = \begin{pmatrix} \cos \gamma & -\sin \gamma & 0 \\ \sin \gamma & \cos \gamma & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \cos \beta & 0 & \sin \beta \\ 0 & 1 & 0 \\ -\sin \beta & 0 & \cos \beta \end{pmatrix} \hat{n}'''_m. \quad (\text{C.43})$$

or as a function of ϕ as:

$$\hat{n}_m''(\phi) = \begin{pmatrix} \cos \gamma \cos \beta \sin \alpha \cos \phi - \sin \gamma \sin \alpha \sin \phi + \cos \gamma \sin \beta \cos \alpha \\ \sin \gamma \cos \beta \sin \alpha \cos \phi + \cos \gamma \sin \alpha \sin \phi + \sin \gamma \sin \beta \cos \alpha \\ - \sin \beta \sin \alpha \cos \phi + \cos \beta \cos \alpha \end{pmatrix}. \quad (\text{C.44})$$

The (x'', y'', z'') coordinate system was defined so that the overall orientation of the mosaic crystal $\hat{n}_0'' = \hat{x}''$. Thus we can compute the angle Δ between the mosaic block and the overall crystal orientation as:

$$\Delta = \cos^{-1} \hat{n}_m'' \cdot \hat{n}_0''. \quad (\text{C.45})$$

When we substitute in for \hat{n}_m'' and \hat{n}_0'' we obtain:

$$\Delta(\phi) = \cos^{-1} [\cos \gamma \cos \beta \sin \alpha \cos \phi - \sin \gamma \sin \alpha \sin \phi + \cos \gamma \sin \beta \cos \alpha]. \quad (\text{C.46})$$

Once we have $\Delta(\phi)$ for the given initial momentum of the neutron we can numerically compute the integral over ϕ to obtain σ :

$$\sigma = \frac{Q}{\gamma_0} W_{norm} \int_0^{2\pi} e^{-\Delta(\phi)^2/2\eta^2} d\phi \quad (\text{C.47})$$

where Q is defined in (C.19), W_{norm} is computed numerically according to (C.7), and

$$\gamma_0 = |\hat{n}_0 \cdot \hat{n}_p| = |\hat{n}_0'' \cdot \hat{n}_p''| = |n_{p_x}''|. \quad (\text{C.48})$$

The probability R for a neutron to reflect from the mosaic crystal with thickness t_0 is then obtained from (C.27):

$$R = \frac{\sigma t_0}{1 + \sigma t_0} \quad (\text{C.49})$$

We thus randomly determine if the neutron reflects.

For a neutron which reflects, we must pick ϕ from the proper probability distribution in order to determine the momentum of the outgoing neutron. We then pick a random number $\xi \in [0, 1]$ and numerically determine the value of ϕ_{out} which solves:

$$\xi = \frac{\int_0^{\phi_{out}} e^{-\Delta(\phi)^2/2\eta^2} d\phi}{\int_0^{2\pi} e^{-\Delta(\phi)^2/2\eta^2} d\phi}. \quad (\text{C.50})$$




Figure C.7: Diffraction of a neutron from a crystal mosaic. Overall crystal (red) has orientation θ_0 . The Bragg condition sets the required orientation θ_B for a mosaic block (pink) to scatter a neutron. The final neutron velocity from (C.51) is shown.

Thus we have randomly determined ϕ_{out} which defines the orientation $\hat{n}_m(\phi_{out})$ of the mosaic block from which the neutron has diffracted.

A neutron which reflects from the crystal monochromator has final momentum

$$\vec{p}_f = \vec{p}_i - \frac{h}{d} \hat{n}_m = \vec{p}_i - 2(\hat{n}_m \cdot \vec{p}_i) \hat{n}_m \quad (\text{C.51})$$

as shown in Figure C.7. Since we can compute $\hat{n}_m(\phi_{out})$ from (C.44) and we have previously computed \vec{p}_i'' in (C.38), it is easy to compute \vec{p}_f'' using (C.51).

For further computation of the neutron trajectory it is convenient to rotate to a final coordinate system (x', y', z') shown in figure C.8. We rotate from (x'', y'', z'') to (x', y', z') by rotating by θ_0 about the $\hat{y}'' = \hat{y}$ axis (this is equivalent to applying the rotation (C.38) a second time).

We now define the “initial position” of the neutron after scattering off of the monochromator in the (x', y', z') reference frame in such a way that all further trajec-




Figure C.8: Primed (blue) and unprimed (black) coordinate systems. The primed coordinate system is rotated by $2\theta_0$ about the y axis. Thus \hat{z}' points in the same direction into which an 8.9 \AA neutron with initial momentum in the \hat{z} direction would scatter.

tory calculations are independent of the presence or absence of the monochromator. First we translate the original reference frame so that the center of the monochromatic crystal $z = z_{crys} + (W_M/2)/\tan\theta_0 \rightarrow z = 0$. Then we apply a rotation by $2\theta_0$ about \hat{y} to obtain the position \vec{x}'_{cross} in the primed coordinate system where the neutron scatters from the monochromator. We then call the neutron's “initial position” the position $(x'(0), y'(0), 0)$ which is consistent with \vec{x}'_{cross} and \vec{p}'_f . That is

$$x'(0) = x'_{cross} - \frac{p'_{fx}}{p'_{fz}} z'_{cross} \quad (\text{C.52})$$

$$y'(0) = y'_{cross} - \frac{p'_{fy}}{p'_{fz}} z'_{cross} \quad (\text{C.53})$$

With these definitions of $x'(0)$ and $y'(0)$ and the ratios $\frac{p'_{fx}}{p'_{fz}}$ and $\frac{p'_{fy}}{p'_{fz}}$ we can now calculate $x'(z')$ and $y'(z')$ for any values of z' using (C.33) and (C.34).

We have now defined the position and momentum of a neutron as it begins its flight toward the apparatus in the presence or absence of a monochromator. The geometry that we consider henceforth is shown in Figure C.9. We define the origin

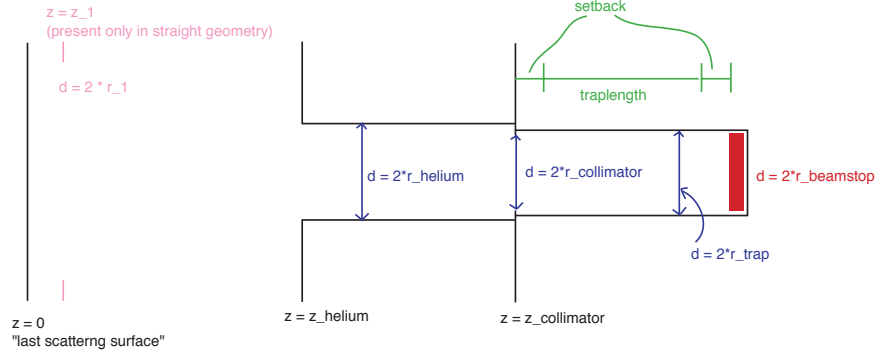


Figure C.9: Geometry of setup under consideration. The position $z = 0$ is taken as the center of the “last scattering surface” (either the end of NG-6 or the monochromator, if it is in the setup).

Table C.1: Nominal values of parameters describing possible 8.9 Å monochromator materials used in the Monte Carlo calculations. Values used were measured or calculated in the cited references.

Material	d [Å]	θ_B	F^2 [10^{-24} cm^{-2}]	N_c [10^{21} cm^{-3}]	References
KC ₈	5.35	56.28	24.48	8.93	[104]
KC ₂₄	8.74	30.61		3.64	
Fluorophlogopite	9.963	26.53	17.17	2.02	[94, 98, 93]

as the center of the “last scattering surface” which is either the end of NG-6 if no monochromator is present, or the monochromator itself. We have already calculated the momentum at this position as described above. The neutron now propagates along a straight path until it reaches the helium bath (see below).

C.2.4 Crystallographic Parameters

Nominal values of the crystallographic parameters are given in Table C.1 for materials of interest as 8.9 Å monochromators. Some of the parameters are obtained from the literature. Others were determined by fitting to our measurements at HMI (see Appendix D).




Figure C.10: The implied reflectivity from sample Z1 (computed as described in the text) is compared to the simulated neutron reflectivity spectrum using the parameters $F^2 = 7 \times 10^{24} \text{ cm}^2$ and $\gamma = 1.9^\circ$. Also shown are the measured and simulated incident neutron spectra.

The parameters for KC_{24} were initially set by using our measurements from HMI. A comparison was made between the measured transmission data through sample Z1 (see figure D.4) and the Monte Carlo simulation as the parameters F^2 and γ were modified. The implied reflectivity of sample Z1 was calculated as: “holder only” - “holder + sample Z1” - slope extracted from curve fit (see section D.3). A comparison of this implied reflectivity curve to the Monte Carlo with $F^2 = 7 \times 10^{24} \text{ cm}^2$ and $\gamma = 1.9^\circ$, using the HMI collimation constraints, is shown in figure C.10. These values were used in the simulations done to determine which monochromator material to use (see section 3.3.3), but were later modified based upon measurements of the final KC_{24} monochromator built and tested at NIST (see section 3.7). For the NIST monochromator, the best fit to the measured reflectivity was obtained with $F^2 =$

$$45 \times 10^{-24} \text{ cm}^{-2}.$$

C.2.5 Testing the Diffraction Calculations

In order to test the reliability of the Monte Carlo simulation for determining the reflectivity of mosaic crystals, an intercalated graphite crystal was simulated under conditions which had previously been measured [104]. Boeuf and his colleagues measured the rocking curve of a 2.08 mm thick sample of KC₈ illuminated with a 7.2 Å monochromatic beam obtained from a Mica monochromator with 0.2° mosaic spread. The measured rocking curve had a peak reflectivity of 70 % and an apparent mosaic width of $\beta = 2.7^\circ$. This experimental data was fit with the numerical parameters $v_0 = 112 \text{ Å}^3$, $\mu = 0.103 \text{ cm}^{-1}$, $\theta = 42.3^\circ$, $\gamma = 1.8^\circ$, $Q' = 0.4Q = 0.289 \text{ cm}^{-1}$. The value of γ (the intrinsic mosaic) was verified experimentally by x-ray diffraction. The fit to a value of $Q' = 0.4Q$ is due to the presence of primary extinction [104, 157]. From equation (C.21) we see that the extracted value of $Q' = 0.4Q$ requires that $0.4 = \tanh A(\lambda)/A(\lambda)$ for $\lambda = 7.2 \text{ Å}$. Since $A \propto 1/\lambda$ then we can solve $7.2 \tanh x/7.2 = 0.4x$ for x an unknown constant. Numerically solving this equation yields $A = 17.74/\lambda$. Thus primary extinction is accounted for by using

$$Q' = \frac{\lambda}{17.74} \tanh \frac{17.74}{\lambda} Q \quad (\text{C.54})$$

where Q is defined by equation (C.19).

Using the Monte Carlo described above, a simulated rocking curve was generated. Neutrons were generated with θ_x and θ_y in the range $\pm 0.1^\circ$ and with wavelength 7.2 Å. A value of $F^2 = 24.48 \times 10^{-24} \text{ cm}^2$ (calculated in [104]) was used and $0.4Q$ was substituted for Q to take primary extinction into account. The mosaic of the crystal was set in the Monte Carlo by $\eta = 0.424\gamma = 0.763^\circ$. The rocking curve generated by the Monte Carlo using these parameters is shown in figure C.11. The simulated




Figure C.11: Simulated rocking curve of 2.08 mm thick sample of KC_8 illuminated by 7.2 \AA neutrons. Parameters for the Monte Carlo are given in the text. Neutron absorption is not considered.

rocking curve has a peak reflectivity of 73 % and a full width half max of $\beta = 2.7^\circ$. The Monte Carlo does not take absorption into account. The average penetration depth should be $0.5T_0 = 1.04 \text{ mm}$ with an average pathlength for a reflected neutron of T_0 . Hence the effect of absorption can be approximated by multiplying the rocking curve by $e^{-\mu T_0} = e^{-0.02} = 0.979$. This reduces the peak reflectivity to 71%. Hence our simulation reproduces a rocking curve consistent with the measurements made by Boeuf and colleagues, and using their calculated parameters.

C.2.6 Primary Extinction

As described in section D.3, the measured mosaic of the KC_8 sample Z3 was $\beta = 3.9 \text{ deg}$. Using the Monte Carlo with primary extinction accounted for using Q' defined by equation (C.54), with beam divergence set to ± 0.005 radians in both θ_x and θ_y and $\lambda = 4.52 \text{ \AA}$ we simulated rocking curves to determine what value of γ would yield $\beta = 3.9 \text{ deg}$. In the simulated data, $\gamma = 3.45 \text{ deg}$ yielded $\beta = 3.9 \text{ deg}$




Figure C.12: Comparison of the simulated and measured transmission spectra through sample Z3. The incident beam was collimated to ± 0.1 deg in the measurements and the simulations.

under the conditions described.

The collimation constraint requires

$$\sqrt{(292 * \theta_x)^2 + (292 * \theta_y)^2} < 0.95/2 \quad (\text{C.55})$$

The Monte Carlo with primary extinction taken into account using (C.54), with $\gamma = 3.45$ deg and with the collimation condition (C.55) was used to simulate an expected transmission spectrum over the range of 8–10 Å. Although the measurements on Z3 show a peak “reflectivity” of $R_{peak} = 70\%$ and a full width half max $\Delta\lambda = 0.40$ Å, the simulation resulted in $R_{peak} = 78\%$ and $\Delta\lambda = 0.48$ Å. A much better fit to the data was obtained with a stronger primary extinction from

$$Q' = \frac{\lambda}{30.08} \tanh \frac{30.08}{\lambda} Q \quad (\text{C.56})$$

(corresponding to $Q' = 0.3Q$ at 9.05 Å, rather than $Q' = 0.49Q$ using equation (C.54). The measured transmission through sample Z3 is compared with the simulated curve using $\gamma = 3.3$ deg with primary extinction from equation (C.56) in figure C.12.

Primary extinction was not considered for Mica or stage 2 IHOPG.

C.2.7 Collimation

In order to minimize neutron induced luminescence and neutron activation, we have previously [41] operated the experiment with the neutron beam collimated so that all neutrons which pass the in-cell collimator are absorbed by the beam stop, except for those which scatter in the liquid helium. When the experiment is operated in the straight geometry (no monochromator), the collimation is defined by two apertures with radii r_1 and $r_{collimator}$ as shown in Figure C.9. The position of the first collimator, a lithium glass plate with a 5.97 cm diameter hole cut in it, is fixed at $z_1 = 17.48$ cm relative to the end of NG-6. Neutron guide NG-6 is 15 cm tall and 6 cm wide. We define $x = y = 0$ at the center of the lithium glass plate, so that the neutron guide covers the ranges $x \in [-3 \text{ cm}, 3 \text{ cm}]$ and $y \in [-11 \text{ cm}, 4 \text{ cm}]$. The value for $r_{collimator}$ is calculated as

$$r_{collimator} = \frac{r_{stop}L_1 - r_1L_2}{L_1 + L_2} \quad (\text{C.57})$$

where $L_1 = z_{collimator} - z_1$ is the distance between the two collimators and $L_2 = 2 * setback + traplength$ is the distance between the second collimator and the beam stop. This value for $r_{collimator}$ ensures the situation described above, where no neutron passing through the in-cell collimator can hit the cell walls unless it scatters in the helium. When the monochromator is used, no in-cell collimator is present.

C.2.8 Neutron Scattering from Liquid Helium

For those neutrons which enter the helium bath (that is those that have $r(z = z_{helium}) < r_{helium}$) we consider the effect of incoherent scattering from the helium bath. Transmission of neutrons with wavelengths of 4.5–14 Å through 69.1 cm of 1.25 K liquid helium was measured in [156]. The experimental transmission values

are fit to a sixth order polynomial transmission function, $f(\lambda)$,

$$f(\lambda) = 1.2261 - 0.81329\lambda + 0.24852\lambda^2 - 0.030131\lambda^3 + 0.0016593\lambda^4 - 3.4636 \times 10^{-5}\lambda^5 \quad (\text{C.58})$$

where λ is in units of Å. The scattering length $b(\lambda)$ is then calculated from $f(\lambda)$ by

$$b(\lambda)[\text{cm}] = \frac{-69.1}{\log f(\lambda)} \quad (\text{C.59})$$

For a simulated neutron of wavelength λ , we choose a random number y between 0 and 1. Then the path length L_{scatt} after which this neutron scatters in the helium is

$$L_{\text{scatt}} = -b(\lambda) \log(1 - y) \quad (\text{C.60})$$

We then examine the trajectory of the neutron through the helium bath. If the neutron's path length through the bath (ignoring scattering) would be longer than L_{scatt} , then we report that the neutron has scattered after travelling distance L_{scatt} . If the neutron's total path length through the helium is less than L_{scatt} , then scattering has no impact on this trajectory.

The neutron trajectory is followed through the helium bath until it hits something or scatters. If the neutron scatters before interacting with a surface, the location of the scattering event is recorded. A neutron is considered to hit a surface if $r(z = z_{\text{helium}}) \geq r_{\text{helium}}$, $r(z = z_{\text{collimator}}) \geq r_{\text{collimator}}$, $r(z \in [z_{\text{helium}}, z_{\text{collimator}}]) \geq r_{\text{helium}}$ or $r(z \in [z_{\text{collimator}}, z_{\text{collimator}} + 2\text{setback} + \text{traplength}]) \geq r_{\text{trap}}$ (see figure C.9). For any position z we can calculate

$$r(z) = \sqrt{\left(x(0) + \left(\frac{p_x}{p_z}\right)z\right)^2 + \left(y(0) + \left(\frac{p_y}{p_z}\right)z\right)^2} \quad (\text{C.61})$$

For the first two absorption conditions above (neutron hits the shielding at the entrance to the helium bath or it hits the collimator), we simply check the value of $r(z)$ at the appropriate position. To determine if the neutron hits the cylinder of

Boron Nitride shielding either in front of the collimator or behind the collimator, we calculate the value of z for which $r(z) = r_{escape}$ where $r_{escape} = r_{helium}$ or r_{trap} . By solving (C.61) for $z = z_{escape}$ we obtain

$$0 = \left(\left(\frac{p_x}{p_z} \right)^2 + \left(\frac{p_y}{p_z} \right)^2 \right) z_{escape}^2 + \left(2x(0) \left(\frac{p_x}{p_z} \right) + 2y(0) \left(\frac{p_y}{p_z} \right) \right) z_{escape} + (x(0)^2 + y(0)^2) \quad (C.62)$$

We then use the quadratic equation to solve for z_{escape} , always taking the larger value. If z_{escape} is less than z at the end of the region in question ($z_{collimator}$ or $z_{beamdump}$ respectively) then the neutron hits the wall at $z = z_{escape}$ (unless it scatters in the helium first, as described above).

C.2.9 Superthermal Production of UCN

If a neutron of the appropriate wavelength enters the trapping region before hitting a wall or scattering, then it contributes to UCN production. For this discussion we restrict ourselves to the single phonon downscattering process only.

The only neutrons which contribute to the scattering rate are those with $k = k^* \pm k_{UCN}$ [30] where $k^* = 2\pi/8.9 \text{ \AA} = 0.706 \text{ \AA}^{-1}$ and $k_{UCN} = \sqrt{\frac{2m_n E_{UCN}}{\hbar^2}}$. The mass of the neutron $m_n = 939.56 \text{ MeV}/c^2$ and $\hbar = 6.58 \times 10^{-22} \text{ MeV s}$. In our case the trap depth determines $E_{UCN} = k_B \mu_B B_{trap}$ where B_{trap} is the trap depth, $k_B = 8.617 \times 10^{-5} \text{ eV/K}$ and $\mu_B = 0.7 \text{ mK/T}$. So for a trap depth B_{trap} T, $k = 0.70598 \pm 0.0053933 \sqrt{B_{trap}} \text{ \AA}^{-1}$. So the superthermal production of UCN occurs only for neutrons with initial wavelengths in the band

$$\lambda = 8.9 \pm 0.068 \sqrt{B_{trap}} \quad (C.63)$$

The calculation of the production rate follows [30] and [42]. From Fermi's Golden Rule we calculate a transition probability by averaging over initial states and summing

over final states the individual transition probabilities $| \langle n_f | P | n_i \rangle |^2$. Thus we can write the production rate for UCN in a volume V as

$$P = \int_V dV \int dE_0 \Phi_0^*(E_0, \vec{r}) \Sigma(E_0 \rightarrow E_{UCN}) dE_{UCN} \quad (\text{C.64})$$

where $\Phi_0^*(E_0, \vec{r})$ is the flux [$\text{n cm}^{-2} \text{s}^{-1} \text{\AA}^{-1}$] of neutrons with energy E_0 as a function of position $vecv_r$ within the volume V . The flux density is normalized such that the flux $\Phi(\vec{r}) = \int_0^\infty \Phi_0^*(E_0, \vec{r}) dE_0$. The integral of the flux with respect to dE_0 provides the average over initial states while the integral over dE_{UCN} sums over the final states. The differential cross section per final energy $\omega = E_{UCN}/\hbar$ for single neutron scattering off of the collection of helium atoms of density n is $\Sigma(E_0 \rightarrow E_{UCN}) = n d\sigma/d\omega$.

As described above this downscattering cross section is negligible, except over a small range of energies near a critical energy $E^* = \hbar^2 k^{*2}/2m$, corresponding to the energy of 8.9 \AA neutrons. Since the flux varies only slightly over the narrow energy band of E_0 for which $\Sigma(E_0 \rightarrow E_{UCN})$ is significant we can approximate $\Phi_0^*(E_0, \vec{r}) = \Phi_0^*(E^*, \vec{r})$. Then we can write the production rate

$$P = \int_V n \Phi_0^*(E^*, \vec{r}) dV \int dE_0 \frac{d\sigma(E_0 \rightarrow E_{UCN})}{d\omega} dE_{UCN} \quad (\text{C.65})$$

Now we apply the principle of detailed balance

$$\int dE_0 \frac{d\sigma(E_0 \rightarrow E_{UCN})}{d\omega} = \int dE_0 \frac{E_{UCN}}{E_0} e^{E_0/kT} \frac{d\sigma(E_{UCN} \rightarrow E_0)}{d\omega} \quad (\text{C.66})$$

to write the unknown downscattering cross section in terms of the upscattering cross section for UCN. The term $\frac{d\sigma(E_{UCN} \rightarrow E_0)}{d\omega}$ implicitly contains a $\delta(E_0 - E^*)$ (see argument above about narrowness of superthermal process) so we can rewrite (C.66) as

$$\int dE_0 \frac{d\sigma(E_0 \rightarrow E_{UCN})}{d\omega} = \frac{E_{UCN}}{E^*} e^{E^*/kT} \int dE_0 \frac{d\sigma(E_{UCN} \rightarrow E_0)}{d\omega} \quad (\text{C.67})$$

We can now express the remaining cross section in terms of experimentally measurable quantities using (3.35) from [30]

$$\int dE_0 \frac{d\sigma(E_{UCN} \rightarrow E_0)}{d\omega} = \sigma_{COH} \frac{k^*}{k_{UCN}} \alpha S(k^*) e^{-E^*/kT} \quad (C.68)$$

where σ_{COH} is the coherent scattering cross section from a bound helium nucleus and $\alpha S(k^*)$ is the measured structure function for superfluid helium evaluated at the momentum transfer corresponding to a phonon of energy E^* . Plugging (C.68) back into C.67 gives

$$\int dE_0 \frac{d\sigma(E_0 \rightarrow E_{UCN})}{d\omega} = \sigma_{COH} \frac{k_{UCN}}{k^*} \alpha S(k^*) \quad (C.69)$$

which can then be substituted back into (C.65) to give the production rate

$$P = \frac{\sigma_{COH} n \alpha S(k^*)}{k^*} \int_V \Phi_0^*(E^*, \vec{r}) dV \int dE_{UCN} k_{UCN} \quad (C.70)$$

For the neutron $k_{UCN} = \sqrt{\frac{2m_n}{\hbar^2} E_{UCN}}$ where m_n is the neutron mass. Taking the energy integral over energies up to the maximum trappable energy as a function of position $E_T(\vec{r})$ yields a production rate for UCN within the trapping region of

$$P = \frac{2}{3} \sqrt{\frac{2m_n}{\hbar^2}} \frac{\sigma_{COH} n \alpha S(k^*)}{k^*} \int_V dV \Phi_0^*(E^*, \vec{r}) E_T^{3/2}(\vec{r}) \quad (C.71)$$

Only UCN in the appropriate spin state, produced with kinetic energy less than difference between the local magnetic potential energy and the “trap depth” (minimum magnetic potential energy along the edge of the trapping region) remain within the trap. The trap depth in energy E_T is due to the potential energy of the interaction between the neutron’s magnetic moment and the local magnetic field and hence can be written $E_T = \mu_n B_T$ where μ_n is the magnetic moment of the neutron. For a Ioffe-type magnetic trap such as ours, the magnetic field increases linearly from the center of the trap ($r = 0$) to the edge ($r = r_T$) where it has value B_T , also called the

trap depth. Hence the maximum energy of trappable neutrons at a given radius r is $E_T(r) = \mu_n E_T \left(1 - \frac{r}{r_T}\right)$.

The trapping rate (1/2 the production rate) is thus

$$P = \frac{1}{3} \sqrt{\frac{2m_n}{\hbar^2}} \frac{\sigma_{COH} n \alpha S(k^*)}{k^*} (\mu_n)^{3/2} (B_T)^{3/2} \int_V dV \Phi_0^*(E^*, \vec{r}) \left(1 - \frac{r}{r_T}\right)^{3/2} \quad (\text{C.72})$$

We use the following numerical values for the parameters:

$$\begin{aligned} \sigma_{COH} &= 1.3 \times 10^{-24} \text{cm}^2 \\ n &= 2.17 \times 10^{22} \text{atoms/cm}^3 \\ \alpha &= 1.45 \\ S_{COH}(k^*) &= 0.12 \\ \mu_n &= 7 \times 10^{-4} \text{K/T} \\ \mu_n^{1/2} &= \sqrt{9.66 \times 10^{-27} \text{J/T}} \\ m_n &= 1.675 \times 10^{-27} \text{kg} \\ \hbar &= 1.05 \times 10^{-34} \text{Js} \\ k^* &= \frac{2\pi}{\lambda} = \frac{2\pi}{8.9\text{\AA}} = 7.1 \times 10^9 \text{m}^{-1} \end{aligned}$$

This yields the trapping rate in the form which we can calculate with the Monte Carlo

$$P = (8.9 \times 10^{-9}) (B_T)^{3/2} \int_V dV \Phi_0^*(E^*, \vec{r}) \left(1 - \frac{r}{r_T}\right)^{3/2} \quad (\text{C.73})$$

We compute this final integral numerically using the Monte Carlo that has been described above. We divide the trapping region (see Figure C.9) into slices along the z-axis. Typically, simulations are done for 100 slices ($N_{slices} = 100$). For the i -th slice we calculate the average radius r_i for the neutron trajectory as it travels through that slice of the trap. We also keep track of the number of neutrons in each slice weighted by their average radius $\sum_i N_i \left(1 - \frac{r_i}{r_T}\right)^{3/2}$. For neutrons which exit the trap (due to

scattering or hitting a wall) we count the contribution to this weighted average as 0.

The numerical integration is performed by computing

$$P = (8.9 \times 10^{-9})(B_T)^{3/2}\Phi_0(A)\frac{N_B A_A}{N_A A_B}V\frac{\sum_i N_i \left(1 - \frac{r_i}{r_T}\right)^{3/2}}{\sum_i N_i} \quad (\text{C.74})$$

For any sufficiently large number of simulated trajectories, the appropriate normalization is achieved by using $\Phi_0(A) = 6.3 \times 10^6 \frac{\text{n}}{\text{cm}^2 \text{sK}}$, the flux at the end of NG-6 and V is the volume of the trapping region $\pi r_{trap}^2 \text{traplength}$ (see Figure C.9). The area A and total number of neutrons passing through that area N must also be kept track of for A (end of NG-6) and B ($z = z_{collimator}$). Since $\sum_i N_i = N_B$ we can write this expression more simply as:

$$P = (8.9 \times 10^{-9})(B_T)^{3/2}\Phi_0(A)\frac{A_A}{N_A}L_{trap}\frac{\sum_i N_i \left(1 - \frac{r_i}{r_T}\right)^{3/2}}{N_{slices}} \quad (\text{C.75})$$

Appendix D

Monochromator Materials Selection

We undertook a series of preliminary measurements to determine the mosaic and 8.9 Å reflectivity of each of the potential monochromator materials. From March 1 to March 26, 2000 a series of measurements were performed at the Hahn Meitner Institute in Berlin. The materials studied were fluorophlogopite, and stages 1 and 2 potassium-intercalated graphite. Measurements were made using two instruments: the membrane diffractometer on cold beam V1 and a two-circle diffractometer with time of flight spectrometry which we built on the fundamental physics beamline, V13.

D.1 Diffractometers

D.1.1 Membrane Diffractometer

The membrane diffractometer is a permanent installation at the HMI, primarily used for studying diffraction from biological membranes. The beam incident on the

sample is a 4.52 Å monochromatic beam reflected from a highly-oriented pyrolytic graphite (HOPG) monochromator with a mosaic of 0.6°. The horizontal divergence of 0.58° full width was defined by two slits 80 cm apart, each with width 4 mm. The instrument also has slits to control the vertical divergence of the beam. The width of these slits was not kept constant from run to run. This, along with the fact that different samples were studied in different sample holders (some pyrex tubes, others with aluminum windows of varying thicknesses) prevents determination of relative reflectivities from sample to sample.

The membrane diffractometer is a typical two-axis diffractometer, and was used to make rocking curve measurements for the purpose of determining the mosaic angles of the materials studied. Diffracted neutrons were detected using a two-dimensional position sensitive detector. The detector is 20 cm tall and has an angular width of 10°. Rocking curves were measured by rotating the sample angle, θ , and recording the intensity of the diffracted neutrons with the overall detector position held constant. The recorded data files sum the neutron counts vertically over the detector, so that the number of diffracted neutrons is known as a function of sample angle (θ) and detector position (ω). To compute rocking curves, these data may be summed over ω .

D.1.2 8.9 Å Diffractometer

The 8.9 Å reflectivity of the samples was measured using a custom-built 8.9 Å diffractometer. Rather than establishing the wavelength of the neutrons incident on the sample using a monochromator, as is traditionally done for shorter wavelength diffractometers, neutron time-of-flight (TOF) was used to determine the wavelength of neutrons incident on the sample as a function of time. The chopper used was

120 mm diameter, rotated at 100 Hz, and had a slit of width 3 mm. The total flight path from the chopper to the detector was 355 cm, yielding a wavelength resolution, set by the copper opening time, of 1% full width at 8.9 Å.

The neutron beam was collimated by two small apertures: a 1.8 mm aperture located 12 cm upstream from the chopper, and a 9.5 mm aperture located 10 cm upstream from the sample. The total separation between the two apertures was 282 cm, so the beam was effectively collimated to 0.1° . The sample and a detector arm were rotated by independent, coaxial rotation stages. Hence the sample angle, θ , and detector angle, ω , could be independently controlled. Neutrons were detected by a plastic scintillator coupled to a photomultiplier tube. An LED and a photodiode viewing the slit in the chopper produced a trigger which started a multi-channel scaler (MCS) at the same phase in the chopper rotation for each cycle. The MCS then counted neutron events seen by the detector as a function of their arrival time relative to the start pulse. The arrival time is proportional to the neutron wavelength (for details on the TOF technique, and a similar setup built at NIST for 8.9 Å diffraction measurements see section 3.7). TOF measurements of the beam spectrum were performed with and without samples, and with the detector in the transmission ($\omega = 0$) and reflection ($\omega = 2\theta_B$) positions.

D.2 Fluorophlogopite Characteristics

The fluorophlogopite studied was obtained from H.C. Materials Corporation¹. The material consisted of many pieces of fluorophlogopite 60×60 mm with a nominal thickness of 0.5 mm, though the thickness of individual pieces ranged from < 0.1 mm to about 0.6 mm. The pieces were initially segregated into “flat” and “wavy” pieces.

¹H.C. Materials Corporation, Urbana, IL, (<http://www.hcmat.com/>).




Figure D.1: A rocking curve for a 6.0 mm thick sample of “flat” fluorophlogopite is shown. The relative diffracted intensity as a function of sample angle θ is reported for a fixed position on the position-sensitive detector. The Gaussian curve fit yields a mosaic of 0.35 degrees.

When “flat” pieces were placed on a flat, smooth surface the corners and edges appeared to be touching the supporting surface at all points. “Wavy” pieces were visibly curved when placed on a flat smooth surface; at least one corner or part of one edge did not touch the supporting surface.

A rocking curve for the (001) reflection of 4.52 Å neutrons from a stack of 6.0 mm total thickness of the “flat” mica is shown in figure D.1. The countrate at a fixed angle (the angle ω at which the highest absolute intensity Bragg diffraction is observed) on the position sensitive detector of the membrane diffractometer is analyzed as function of the sample angle, θ . Fitting the rocking curve data to a Gaussian yields a mosaic angle of 0.35 degrees. Using the TOF setup, the peak reflectivity of the sample at 9.3 Å was measured to be 30%.

Several attempts were made to increase the mosaic of the mica. We constructed stacks of shimmed mica, where 160 μm aluminum sheets were inserted along one edge of the mica, between some of the layers. In all of the shimmed stacks, the peak reflectivity was less than that of the “flat” mica (since the amount of mica aligned with any given angle was thinner than the 6.0 mm thickness of the “flat” sample). Furthermore, the stacks did not produce symmetric, Gaussian, rocking curves, but instead complicated diffraction patterns. Misalignment of the different mica layers relative to the vertical direction of the beam (the tilt) could clearly be seen using the two-dimensional position-sensitive detector on the membrane diffractometer. Hence, different layers of the shimmed structure were diffracting neutrons to different vertical positions on the detector. Furthermore, the reflectivity of the individual layers varied strongly. In the end, no satisfactory samples were prepared in this manner.

The broadest observed mosaic for a fluorophlogopite sample was obtained by simply stacking all of the mica, including both the “flat” and “wavy” pieces in a single stack with minimal compression. The total thickness of the stack was 1.3 cm, consisting of 0.7 cm of “wavy” mica and 0.6 cm of “flat” mica. The “wavy” pieces were placed closest to the beam entrance side of the stack. A scan of the (001) diffraction peak at 4.52 Å using the membrane diffractometer is shown in figure D.2. Relative intensity contours are shown as a function of sample and detector angle. The presence of two peaks in the θ direction, indicates that there are two parts to the fluorophlogopite sample, with the same d spacing for the crystal) but in slightly different directions (due to relative misalignments of the mica crystals). The presence of two peaks in the ω direction (for $\theta \sim 12.9^\circ$) is indicative of two effective d -spacings for crystal planes with the same overall orientation. This is likely due to stress within some of the crystals (perhaps the same internal stresses that have caused the crystals




Figure D.2: The relative intensity contours are plotted as a function of sample angle (θ) and detector angle (ω) for a 1.3 cm thick sample of fluorophlogopite.

to be “wavy”). The effective mosaic of the double-peaked structure is 0.7° . The peak 8.9 \AA reflectivity of this sample was measured to be 31% using the TOF technique.

D.3 Intercalated Graphite Characteristics

Five samples of potassium graphite intercalation compounds (GICs) were studied in our preliminary measurements. All of the samples were prepared by the group of Hartmut Zabel of the Ruhr Universität Bochum during February and March of 2000. The comments herein are confined to the measured characteristics of sample Z3, the highest quality stage 1 sample, and sample Z1, the highest quality stage 2 sample. The pyrolytic graphite pieces intercalated were all grade ZYB HOPG². Sample Z1

²HOPG of grade ZYB is specified as having a mosaic of $0.8^\circ \pm 0.2^\circ$. Advanced Ceramics Corporation, Lakewood, OH, (<http://www.advceramics.com>).




Figure D.3: Rocking curve for GIC sample Z1. For each angle θ the counts have been summed over the entire detector. The apparent mosaic of the sample is 2.2° .

was 2 cm by 5 cm by 2 mm before intercalation. The substrate for sample Z3 was obtained by cleaving one piece of graphite with those same dimensions roughly in half (so that the nominal thickness was 1 mm prior to intercalation).

Sample Z1 split (parallel to the crystal planes) over about a 1×2 cm are at one end of the sample. For diffraction measurements, the split area was masked off. The split occurred during intercalation as the sample apparently reached stage 2, as determined by the sample color³). Intercalation was stopped after the crack was observed, and sample Z1 was assumed in these measurements to be a pure stage 2 sample. The mosaic of sample Z1 was measured on the membrane diffractometer to be 2.2° for the stage 2, $n=1$ reflection (see figure D.3).

³While unintercalated HOPG has a gunmetal gray appearance, stage 1 potassium GICs appear gold, stage 2 IHOPG appears a dark metallic blue, and higher stages appear a lighter metallic blue. The stage of an intercalated compound can thus be determined reasonably reliably from the sample color. Diffraction measurements are necessary to make definitive stage determinations.




Figure D.4: Transmission curves through the Z1 sample holder in the presence and absence of sample Z1 are shown. The difference between the two curves is also shown. When the difference is fit to a Gaussian (the diffracted neutrons) plus a slope (a background loss process), the implied peak reflectivity is 0.51 ± 0.02 at 8.85 \AA .

Sample Z1 was studied on the 9 \AA diffractometer in a sample holder made entirely of aluminum. Both windows in this sample holder were 0.5 mm thick aluminum. The neutron spectrum transmitted through the sample holder with and without sample Z1 inside is shown in figure D.4. The difference between the spectrum transmitted with and without the sample is due to on-Bragg diffraction plus any other losses caused by the presence of the sample (such as absorption or incoherent scattering). As is shown, the difference curve does not have a flat background, but rather a background that fits to a linear function of λ . The difference curve was thus fit to a linear slope plus a Gaussian. The data fit to the function:

$$(769 \pm 35) - (69.3 \pm 3.5)\lambda + (417 \pm 14)e^{((\lambda - 8.85 \pm 0.01)/2(0.269 \pm 0.012))^2} \quad (\text{D.1})$$




Figure D.5: Rocking curve for GIC sample Z3. For each angle θ the counts have been summed over the entire position sensitive detector. The apparent mosaic of the sample is 3.9° .

with a χ^2 per degree of freedom of 1.4. The peak reflectivity at 8.85 \AA calculated as the ratio of the amplitude of the Gaussian to the number of counts seen with the sample holder only, is $51 \pm 2\%$. This data was also used to determine the effective structure function and inherent mosaic, as described in section C.2.3. Although sample Z1 was determined not to contain any significant amount of stage 1 contamination, the presence of stage 3 in the sample was untested.

Sample Z3 was apparently an almost entirely stage 1 sample. For a given θ , dips in the transmission spectra on the the 9 \AA diffractometer were much stronger at the stage 2 wavelengths than at the stage 1 wavelengths. The mosaic of sample Z3 was measured on the membrane diffractometer to be 3.9° (see figure D.5).

On the TOF diffractometer, sample Z3 was held in a sample holder with two 0.1 mm thick windows. The sample was rotated so that the peak of the stage 1




Figure D.6: Comparison of the transmission spectrum through sample Z3 with the beam spectrum in the absence of a sample. The incident beam was collimated to ± 0.1 deg. The transmitted beam is uncorrected for absorption by the windows of the sample holder.

reflection was expected to be near 8.9 \AA . A transmission spectrum through the sample made with the collimator 9.5 mm in diameter is compared with a spectrum taken in the absence of any sample in figure D.6. The peak reflectivity, at 9.06 \AA is $70 \pm 4\%$. These measurements were also used to tune the way in which extinction was handled for the KC_8 material in the Monte Carlo simulations (see section C.2.3).

Appendix E

The Parser

The purpose and operational principles of the parser is described in section 5.3.1. This appendix contains the code for the parsing routine. The code and the comments within it demonstrate how the parameters `startcross` and `finishcross` are determined for each pulse. The method for determining which points are summed to calculate the pulse area is also shown.

Certain external references bear explanation in order to make the code understandable. The parsing threshold `thresh` is determined manually and passed to the parser. The `averageevent()` function computes the “average event” (as described in section 5.3.1). It also determines the location at which the main event is expected, which is the most negative point of the average event. This trigger position is stored in the `pt.off` variable from which the pointer `trigpos` is computed for each event.

The results of analysis by the parser are stored in `frame`, which is an array of `Frame` structures. The `Frame` structure stores the main pulse area and a list of afterpulses and is declared as:

```
struct Frame {  
    double    area;                               /*main pulse area*/  
    AfterPulseList *afterpulses;
```

```
};
typedef struct Frame Frame;
```

The afterpulses are stored in a linked list declared as:

```
struct AfterPulseList {
    double ptime;
    double parea;
    struct AfterPulseList *next;
};
typedef struct AfterPulseList AfterPulseList;
```

The code of the parser follows:

```
#include <profiler.h>
#include <stdio.h>
#include <stdlib.h>
#include <math.h>
#include <timer.h>
#include "sequencer.h"
#include "parse.h"
#include "parserutilities.h"
#include "params.h"
#include "CTVByteblock.h"

struct Pulse {
    UInt32    ptime;           /*point in the event of startcross*/
    double    parea;           /*pulse area*/
    UInt8     pheight;         /*peak height*/
};
typedef struct Pulse Pulse;

/*parse:  Takes the following arguments:
data:      a pointer a block of raw data from the Signatec card
numframes: the number of events stored within *data
frames:    an array of Frame structures in which the
           parsed data is stored.
thresh:    the parsing threshold*/

void parse (UInt8 *data, long numframes, Frame *frames, UInt8 thresh)
{
/*BEGIN VARIABLE DECLARATION*/
```

```

long i,j,k,l; /*index variables*/
UInt8 crossed; /*boolean: has the threshold been crossed?*/
long nerr,perr,terr,serr; /*count up various errors*/

/*for computing background average*/
UInt8 *backgroundpt;
double backgroundave, bgsum, bgarray[NUMFRAMESTOAVE];

double peaksum; /*integral of peak*/

/*dataend points to the end of the *data block.
event and eventend point to the first and last + 1 byte
of the current event, respectively*/
UInt8 *dataend, *event, *eventend;

UInt8 *trigpos;      /*the expected trigger position*/
UInt8 *startcross;   /*first point below threshold in peak*/
UInt8 *finishcross;  /*first point above threshold after peak*/
UInt8 *safebefore;   /*beginning of "safe" region for integral*/
UInt8 *safeafter;    /*end of "safe" region for integral*/
UInt8 *nextpeak;     /*startcross for the next peak*/
UInt8 *currpt;       /*incremented through file*/

UInt8 *findpeak, *mostneg;

/*data structures for storing pulse characteristics
for the current event*/
Pulse allpulses[MAX_AFTERPULSES];
Pulse *mainpulse;
int mainpulsenum;
int lastpulsenum;

/*pointers to the data structures used for storing the parsed
data*/
Frame *currframe;
AfterPulseList *currafterpulse;
/*END VARIABLE DECLARATION*/

nerr=0; perr=0; terr=0; serr=0;
dataend = data + numframes * framelength;

/*generate the average event.  Fix byte ordering if necessary.
```

```

set pt_off to the position of the peak of the (neg) pulse in
the average event*/
averageevent(data,(float) numframes);

/*correct for missing 4 bytes at beginning of some blocks*/
if(pt_off <= 14)
{
    data -= 4;
    fprintf(ferr,"data shifted by 4 points\n");
    pt_off += 4;
}

l=0;
bgsum=0.0;

/*parse events*/
for(i = 0; i<numframes; i++)
{
    lastpulsenum = 0;
    currframe = frames + i;
    event = data + i*framelength;
    eventend = event + framelength;
    trigpos = event + pt_off;
    mainpulse=NULL;

    /*calculate the background average by averaging the last
    16 points in the event. Then average the backgrounds from
    NUMFRAMESTOAVE events*/
    backgrounddave = 0.0;
    for(backgroundpt = (eventend-16);
        backgroundpt < eventend; backgroundpt++)
        backgrounddave += *backgroundpt;
    backgrounddave /= 16.0;
    if (i >= NUMFRAMESTOAVE)
        bgsum -= bgarray[i % NUMFRAMESTOAVE];
    bgarray[i % NUMFRAMESTOAVE] = backgrounddave;
    bgsum += backgrounddave;
    if(i < NUMFRAMESTOAVE)
        backgrounddave = bgsum / (double) (i+1);
    else
        backgrounddave = bgsum / (1.0 * NUMFRAMESTOAVE);
}

```

```

crossed=0;
/*make sure that if the file starts with a peak we skip it*/
for(startcross = event;
    (*startcross < thresh) && ((trigpos - startcross) >1);
    startcross++);

/*initialize safeafter to the first point in the file
above the threshold,for use by the peak finding routine*/
safeafter=startcross;

/*advance currpt either to just before the first peak or
just before the expected trigger position*/
for(currpt = startcross;
    ((trigpos - currpt) > 1) && (crossed == 0);
    currpt++)
{
    if (*currpt < thresh)
        crossed = 1;
}
if (crossed == 1)
    currpt--;
if (currpt <= startcross)
    serr++;

/*Go through the file and write into an array the starting
point and integrated area of every peak, all determined the
same way. Write a special marker to keep track which one is
the "main peak" (the one that contains the trigger
position). Once all peaks are integrated, go back and dump
everything into the Frames data structure. If there is no
main peak, need to mark it anyway so that timestamps still
line up.*/
for(j=0, nextpeak=currpt, startcross=currpt;
    (nextpeak < eventend) && (j< MAX_AFTERPULSES);
    j++, startcross=nextpeak)
{
    /*move forward until we find a crossing, or pass out
of the event region*/
    for(;(*startcross > thresh) && (startcross < eventend);
        startcross++);

    /*find beginning of safe region. This is the later

```

```

of (SAFEBEFORE before startcross and the location
of the previous peak's safeafter)*/
safebefore = startcross - SAFEBEFORE;
if (safebefore < safeafter)
safebefore=safeafter;

/*find finishcross; note require 3 consecutive
points to be above the threshold for the peak to
be considered finished*/
for(finishcross = startcross;
    (((*finishcross <= thresh)
    || ((*finishcross+1) <= thresh)
    && (finishcross + 1 < eventend))
    || ((*finishcross+2) <= thresh)
    && (finishcross + 2 < eventend)))
    && (finishcross < eventend));
finishcross++);

/*if trigpos is located within 5 points of the
region between startcross and finish cross, then
set the mainpeak pointer to this event*/
if ((trigpos >= (startcross-5))
    && (trigpos <= (finishcross+5)))
    mainpulse = allpulses + j;

/*as long as the current peak finished before the
end of the event region, write down its information.
find next peak (if there is one within SAFEAFTER,
and the event)*/
for(nextpeak = finishcross;
    (nextpeak - finishcross < SAFEAFTER)
    && (nextpeak<eventend) && (*nextpeak > thresh);
    nextpeak++);

/*if we found a next peak, then put the end of the
safe region at the minimum point between the peaks*/
if((nextpeak - finishcross < SAFEAFTER)
    && (nextpeak < eventend))
    safeafter = findmax(finishcross,nextpeak);
else
    safeafter = nextpeak;

```

```
/*if safeafter isn't the end of the file, then
compute the integral and save the information
about the pulse in the array. if safeafter is the
end of the file, but this is the first pulse in
the file, then integrate it anyway. figure out
what region to integrate over as follows: start
from startcross to finishcross. Extend startcross
back until either it reaches safebefore or its
value crosses the background. Then extend
finishcross forward until it reaches safeafter
or its value crosses the background*/
if ((safeafter < eventend) || (j == 0))
{
    /*Assign each peak's time as the time of its
    smallest (i.e. most negative voltage for PMT
    pulses) point*/
    for(findpeak = startcross, mostneg = startcross;
        findpeak < finishcross; findpeak++)
    {
        if(*findpeak < *mostneg)
            mostneg = findpeak;
    }
    allpulses[j].ptime = mostneg - event;
    allpulses[j].pheight = *mostneg;

    /*move startcross until it reaches the
    background or safebefore*/
    for(;(startcross > safebefore)
        && (*startcross < backgrounddave);
        startcross--);
    startcross++;

    /*move finishcross until it reaches the
    background or safebefore*/
    for(;(finishcross < safeafter)
        && (*finishcross < backgrounddave);
        finishcross++);

    peaksum=0;
    for(currpt = startcross;
        currpt < finishcross; currpt++)
        peaksum+=((double) *currpt)-backgrounddave;
```

```

        /*save the peak sum*/
        allpulses[j].parea=peaksum;
        if(peaksum > 0)
        {
            perr++;
        }
        lastpulsenum = j+1;

    } /*if safeafter*/

} /*for each peak*/

/*load the results of parsing the current event into the
appropriate Frame datastructure (unless the maximum number
of pulses was exceeded, in which case the event must be
all noise).*/
if(lastpulsenum > maxpulsenum)
{
    fprintf(ferr,"%dt%d\n",lastpulsenum,i+1);
    terrr++;
    currframe->area = 0;
    currframe->afterpulses = NULL;
}
else
{
    /*if there was a main pulse, print out the
    information about each pulse, otherwise print out
    0.0 for the main pulse size and print out all
    other pulses as afterpulses*/
    if (mainpulse == NULL) /*no main pulse*/
    {
        nerrr++; /*count events with no main pulse*/

        currframe->area=0;
        if (lastpulsenum != 0) /*found >= 1 afterpulse*/
        {
            currframe->afterpulses =
            (AfterPulseList *) malloc(sizeof(AfterPulseList));
            if (currframe->afterpulses == NULL)
            {
                printf("parser: out of memory\n");
            }
        }
    }
}

```

```

        exit(-506);
    }
    currafterpulse = currframe->afterpulses;
    currafterpulse->next = NULL;
    for(j=0;j<lastpulsenum;j++)
    {
        currafterpulse->ptime =
(allpulses[j].ptime - pt_off)*xincr;
        currafterpulse->parea =
            allpulses[j].parea;

        /*If there is at least one more
        afterpulse left then allocate another
        afterpulselist entry*/
        if( (j+1) < lastpulsenum)
        {
            currafterpulse->next =
(AfterPulseList *) malloc(sizeof(AfterPulseList));
            if (currafterpulse->next == NULL)
            {
                printf("out of memory\n");
                exit(-506);
            }
            currafterpulse =
                currafterpulse->next;
        }
        else
            currafterpulse->next = NULL;
    }
}
else /*there is a main pulse*/
{
    mainpulsenum = mainpulse - allpulses;
    currframe->area=mainpulse->parea;

    if (lastpulsenum > 1) /*found >= 1 afterpulse*/
    {
        currframe->afterpulses =
(AfterPulseList *) malloc(sizeof(AfterPulseList));
        if (currframe->afterpulses == NULL)
        {

```

```

        printf("parser: out of memory\n");
        exit(-506);
    }
    currafterpulse = currframe->afterpulses;
    currafterpulse->next = NULL;
    for(j=0;j<lastpulsenum;j++)
    {
        if (j != mainpulsenum)
        {
            currafterpulse->ptime =
(allpulses[j].ptime - pt_off)*xincr;
            currafterpulse->parea =
                allpulses[j].parea;

            /*If there is at least one more
            pulse left which isn't the main
            pulse then allocate another
            afterpulselist entry*/
            if (((j+1) < lastpulsenum)
                && !((j+1 == mainpulsenum)
                    && (j+2 == lastpulsenum)))
            {
                currafterpulse->next =
(AfterPulseList *) malloc(sizeof(AfterPulseList));
                if (currafterpulse->next
                    == NULL)
                {
                    printf("out of memory");
                    exit(-506);
                }
                currafterpulse =
                    currafterpulse->next;
            }
            else
                currafterpulse->next = NULL;
        }
    }
} /*if lastpulsenum > 1*/
else
    currframe->afterpulses = NULL;
} /*else there is a main pulse*/
} /*else not more than max num afterpulses*/

```

```
    } /*for each event*/

    printf("For this set: %ld\tp: %ld\tn: %ld\ts: %ld\n",
           terr,perr,nerr,serr);
    fprintf(ferr,"\nfor this block: %ld\tp: %ld\tn: %ld\ts: %ld\n",
           terr,perr,nerr,serr);
}
```


Appendix F

Detection Efficiency Calculation

As described in section 6.6, a simple Monte Carlo simulation is used to estimate the average probability of detecting the decay of a trapped neutron in our apparatus.

The simulation operates by considering a large number of neutron decays within the apparatus. Each event is assigned a random z position within the trap volume (from a distribution). The energy of the beta-decay electron is chosen from the known beta-decay spectrum. The total number of photoelectrons detected by both detectors is computed as

$$\text{YEFFICIENCY} \times \text{reldetectioneff}(z) \times \text{pepermeV}(EMeV) \quad (\text{F.1})$$

where **YEFFICIENCY** is the total efficiency of the "Y" splitter, $\text{reldetectioneff}(z)$ is a function which returns the relative detection efficiency as a function of the particle's position (z) where $z = 0$ is the center of the trap and $\text{pepermeV}(EMeV) = 29.7 \times EMeV$ is the total number of photoelectrons for a decay event with $EMeV$ MeV of energy. The detection efficiency is assumed to be independent of the r position of the decay. The position z can vary from **+ZMAX** to **-ZMAX**.

The code of the calculation follows:

```

/*
 * Detection Efficiency Monte Carlo
 *
 * ran1 is a random number generator from numerical recipes in
 * C which returns a random float between 0 and 1.0
 */

#include <stdio.h>
#include <math.h>
#include "random.h"

#define YEFFICIENCY 0.94
#define ZMAX 11.0
#define PEPERMEV 29.7 /*Determined from tin calibration*/
#define SLOPE 0.0312 /*from position measurement with alpha
                      source*/

#define NEVENTS 1000000
#define E_max (1.2926) /* Maximum electron energy in MeV */
#define M_e (0.511) /* electron mass in MeV */

/* Integrated Beta spectrum (From Clint Brome's PhD Thesis)*/
/* for an input energy E in MeV, returns the fraction of the
spectrum with energy less than E. So iBeta(0) = 0 and
iBeta(E_max) = 1.*/
double iBeta(double E)
{
    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);
}

/* number of photoelectrons for an E keV decay event in the

```

```
center of the apparatus*/
double pepermev(double E)
{
    return E * PEPERMEV;
}

/*relative detection efficiency as a function of position*/
double reldetectionefficiency(double z)
{
    return (1.0 - SLOPE * fabs(z));
}

int main(void)
{
    long thresh;          /*threshold in p.e. per PMT*/
    double E;             /*energy in MeV of decay electron*/
    double z;             /*position in cm from center of trap*/
    double pe;            /*total number of p.e. reaching PMTs
                           for this decay*/
    double p,step,q,j; /*beta spectrum calculation parameters*/
    long pei;             /*total number of p.e. reaching
                           PMTs (integer)*/
    long A,B;            /*total number of PMT reaching
                           PMT A and B respectively*/
    long i;               /*index variable*/
    long idum;            /*random number generator seed*/
    long n_events;
    long n_detected;

    seed_random();        /*seed the random number generator*/

    printf("This is the detection efficiency Monte Carlo\n");

    for(thresh = 1; thresh < 9; thresh++)
    {
        n_events = NEVENTS;
        n_detected = 0;

        for(i=0;i<n_events;i++)
        {
            /*randomly choose the decay energy.
```

```

        randomly choosing a number p between 0 and 1
        do binary search through the integrated beta
decay spectrum to find
        the energy for which p of the spectrum has
less energy.
        Must get final energy to within 10 eV.*/
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;

/*randomly choose the decay position*/
z = ZMAX * 2.0 * (ran1(&idum) - 0.5);

/*total number of photoelectrons for this
decay event*/
pe = YEFFICIENCY * pepermev(E)
* reldetectionefficiency(z);

pei = (long) floor(pe);

/*adjust for any fractional photoelectrons*/
if (ran1(&idum) < (pe - pei))
    pei++;

/*divide up pei photons between
PMTs A and B, randomly*/
A=0;B=0;
for(j=0;j<pei;j++)
{
    if (ran1(&idum) < 0.5)
        A += 1;

```

```
        else
            B += 1;
    }

    if ((A >= thresh) && (B >= thresh))
        n_detected++;
    }

    printf("%ld events of %ld were detected\n",
n_detected,n_events);
    printf("Efficiency is %lf for %ld p.e. per PMT\n",
n_detected / ((double) n_events),thresh);
}
return 0;
}
```


Bibliography

- [1] J. Chadwick. Possible existence of a neutron. *Nature* **129**, 312 (1932).
- [2] E. Fermi. In search of a theory of β -rays. *Zeitschrift für Physik* **88**, 161 (1934).
- [3] K. Grotz and H. V. Klapdor. *The weak interaction in nuclear, particle and astrophysics*. Adam Hilger, 1990. Translated from the German original by S. S. Wilson.
- [4] J. D. Jackson, S. B. Treiman, and H. W. Wyld. Possible tests of time reversal invariance in beta decay. *Physical Review* **106**, 517 (1957).
- [5] D. E. Groom et al. The review of particle physics. *The European Physical Journal C* **15**, 1 (2000).
- [6] 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).
- [7] H. Abele, S. Baessler, D. Dubbers, J. Last, U. Mayerhofer, C. Metz, T. M. Mueller, V. Nesvizhevsky, C. Raven, O. Scheaerpf, and O. Zimmer. A measurement of the beta asymmetry a in the decay of free neutrons. *Physics Letters B* **407**, 212 (1997).
- [8] J. Reich, H. Abele, M. Astruc Hoffman, S. Baeßler, P. V. Bülow, D. Dubbers, V. Nesvizhevsky, U. Peschke, and O. Zimmer. A measurement of the beta asymmetry in neutron decay with PERKEO II. *Nuclear Instruments and Methods in Physics Research A* **440**, 535–538 (2000).
- [9] I. S. Towner and J. C. Hardy. The current status of V_{ud} . In *Physics Beyond the Standard Model*. World Scientific, 1999.
- [10] P. Slocum. *The Panofsky Ratio, and the Response of the PIBETA Calorimeter to Photons and Positrons*. PhD thesis, University of Virginia, 1999.
- [11] K. Saito and A. W. Thomas. Effect of nucleon structure variation in super-allowed Fermi beta-decay. *Physics Letters B* **363**, 157–161 (1995).

- [12] D. H. Wilkinson. Analysis of neutron β -decay. *Nuclear Physics A* **377**, 474–504 (1982).
- [13] W. J. Marciano and A. Sirlin. Radiative corrections to β decay and the possibility of a fourth generation. *Physical Review Letters* **56**, 22–25 (1986).
- [14] A. Garcia, J. L. Garcia-Luna, and G. Lopez Castro. Neutron beta decay and the current determination of V_{ud} . *Physics Letters B* **500**, 66–74 (2001).
- [15] R. Carr et al. Technical review report for an accurate measurement of the neutron spin - electron angular correlation in polarized neutron beta decay with ultra-cold neutrons. NSF Proposal (Unpublished). Available at <http://www.krl.caltech.edu/ucn/UCNSD2v12.doc>, 2000.
- [16] E. W. Kolb and M. S. Turner. *The Early Universe*. Perseus Press, 1990.
- [17] S. Burles, K. M. Nollett, J. W. Turan, and M. S. Turner. Sharpening the predictions of big-bang nucleosynthesis. *Physical Review Letters* **82**, 4176–4179 (1999).
- [18] D. N. Schramm and L. Kawano. Cosmology and the neutron lifetime. *Nuclear Instruments and Methods in Physics Research A* **284**, 84–88 (1989).
- [19] R. E. Lopez and M. S. Turner. Precision prediction for the big-bang abundance of primordial ^4He . *Physical Review D* **59**, 103502 (1999).
- [20] T. P. Walker, G. Steigman, D. N. Schramm, K. A. Olive, and H. Kang. Primordial nucleosynthesis redux. *The Astrophysical Journal* **376**, 51–69 (1991).
- [21] G. Stasinska and Y. Izotov. Collisional excitation of hydrogen and the determination of the primordial helium abundance from H II regions. *Astronomy and Astrophysics* **378**, 817–825 (2001).
- [22] K. Olive, E. Skillman, and G. Steigman. The primordial abundance of ^4He : an update. *Astrophysical Journal* **483**, 788–797 (1997).
- [23] Y. I. Izotov and T. X. Thuan. Reexamining the helium abundance of I Zw 18. *Astrophysical Journal* **497**, 227–237 (1998).
- [24] 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).
- [25] 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).

- [26] W. M. Snow, Z. Chowdhuri, M. S. Dewey, X. Fei, D. M. Gilliam, G. L. Greene, J. S. Nico, and F. E. Weitfeldt. A measurement of the neutron lifetime by counting trapped protons. *Nuclear Instruments and Methods in Physics Research A* **440**, 528–534 (2000).
- [27] M. S. Dewey, D. M. Gilliam, J. S. Nico, F. E. Wietfeldt, W. M. Snow, and G. L. Greene. Measurement of the free neutron lifetime. Technical Report NISTIR 6798, NIST, 2001.
- [28] Z. Chowdhuri. *Absolute neutron measurements in neutron decay*. PhD thesis, Indiana University, 2000.
- [29] J. S. Nico. NIST. Personal communication.
- [30] R. Golub, D. Richardson, and S. K. Lamoreaux. *Ultra-Cold Neutrons*. Adam Hilger, 1991.
- [31] 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).
- [32] W. Mampe, P. Ageron, C. Bates, J. M. Pendlebury, and A. Steyerl. Neutron lifetime measured with stored ultracold neutrons. *Physical Review Letters* **63**, 593–596 (1989).
- [33] 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).
- [34] S. Arzumanov, L. Bondarenko, S. Chernyavsky, W. Drexel, A. Fomin, P. Geltenbort, V. Morozov, Yu. Panin, J. Pendlebury, and K. Schreckenbach. Neutron life time value measured by storing ultracold neutrons with detection of inelastically scattered neutrons. *Physics Letters B* **483**, 15–22 (2000).
- [35] V. V. Vladimirkii. Magnetic mirrors, channels and bottles for cold neutrons. *Soviet Physics JETP* **12**, 740–746 (1961).
- [36] G. Abov Yu, S. P. Borovl'ev, 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).
- [37] 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).

-
- [38] 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).
- [39] N. Niehues. *Untersuchungen an einer magnetischen Flasche zur Speicherung von Neutronen (Investigations on a magnetic bottle for the storage of neutrons)*. PhD thesis, Rheinischen Friedrich-Wilhelms-Universität, Bonn, 1983.
- [40] 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).
- [41] 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).
- [42] C. R. Brome. *Magnetic Trapping of Ultracold Neutrons*. PhD thesis, Harvard University, 2000.
- [43] C. R. Brome, J. S. Butterworth, K. J. Coakley, M. S. Dewey, S. N. Dzhosyuk, D. M. Gilliam, R. Golub, G. L. Greene, K. Habicht, S. K. Lamoreaux, C. E. H. Mattoni, D. N. McKinsey, P. R. Huffman, F. E. Wietfeldt, and J. M. Doyle. Magnetic trapping of ultracold neutrons. *Physical Review C* **63**, 055502 (2001).
- [44] Y. V. Gott, M. S. Ioffe, and V. G. Tel'kovskii. *Nuclear Fusion* **3**, 1045 (1962).
- [45] H. F. Hess et al. Magnetic trapping of spin-polarized atomic hydrogen. *Physical Review Letters* **59**, 672–675 (1987).
- [46] R. van Roijen, J. J. Berkhout, S. Jaakkola, and J. T. M. Walraven. Experiments with atomic hydrogen in a magnetic trapping field. *Physical Review Letters* **61**, 931–934 (1988).
- [47] J. M. Doyle, B. Friedrich, J. Kim, and D. Patterson. Buffer-gas loading of atoms and molecules into a magnetic trap. *Physical Review A* **52**, R2515–R2518 (1995).
- [48] W. D. Phillips. Laser cooling and trapping of neutral atoms. In *Laser Manipulation of Atoms and Ions*. North Holland, 1992.
- [49] D. G. Henshaw and A. D. B. Woods. Modes of atomic motions in liquid helium by inelastic scattering of neutrons. *Physical Review* **121**, 1266 (1961).
- [50] R. Golub and J. M. Pendlebury. Super-thermal sources of ultra-cold neutrons. *Physics Letters* **53A**, 133–135 (1975).

- [51] R. Golub and J. M. Pendlebury. The interaction of ultra-cold neutrons (UCN) with liquid helium and a superthermal UCN source. *Physics Letters* **62A**, 337–339 (1977).
- [52] R. Golub, C. Jewell, P. Ageron, W. Mampe, B. Heckel, and I. Kilvington. Operation of a superthermal ultra-cold neutron source and storage of ultra-cold neutrons in superfluid helium-4. *Zeitschrift für Physik B* **51**, 187–193 (1983).
- [53] 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).
- [54] R. Golub. On the storage of neutrons in superfluid ^4He . *Physics Letters A* **72**, 387–390 (1979).
- [55] F. Pobell. *Matter and Methods at Low Temperatures*. second edition. Springer, 1996.
- [56] P. V. E. McClintock. An apparatus for preparing isotopically pure He^4 . *Cryogenics*, 201–208 (1978).
- [57] P. C. Hendry. The isotopic purification of ^4He . Master’s thesis, University of Lancaster, 1985.
- [58] Richard C. Pardo. Argonne national laboratory. Personal communication.
- [59] D. N. McKinsey, C. R. Brome, J. S. Butterworth, S. N. Dzhosyuk, P. R. Huffman, C. E. H. Mattoni, and J. M. Doyle. Radiative decay of the metastable $\text{He}_2(a^3\Sigma_u^+)$ molecule in liquid helium. *Physical Review A* **59**, 200–204 (1999).
- [60] M. Stockton, J. W. Keto, and W. A. Fitzsimmons. Ultraviolet emission spectrum of electron-bombarded superfluid helium. *Physical Review Letters*, 654–657 (1970).
- [61] G. M. Seidel, R. E. Lanou, and W. Yao. Rayleigh scattering in rare-gas liquids. hep-ex/0111054, 2001.
- [62] 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. Fluorescence efficiencies of thin scintillating films in the extreme ultraviolet spectral region. *Nuclear Instruments and Methods in Physics Research B* **132**, 351–358 (1997).
- [63] W. M. Burton and B. A. Powell. Fluorescence of tetraphenyl-butadiene in the vacuum ultraviolet. *Applied Optics* **12**, 87–89 (1973).

- [64] E. C. Bruner Jr. Absolute quantum efficiency of sodium salicylate for excitation by extreme ultraviolet. *Journal of the Optical Society of America* **59**, 204 (1969).
- [65] D. N. McKinsey. *Detecting Magnetically Trapped Neutrons: Liquid Helium As A Scintillator*. PhD thesis, Harvard University, 2002.
- [66] P. R. Huffman, C. R. Brome, J. S. Butterworth, S. N. Dzhosyuk, R. Golub, K. Habicht, S. K. Lamoreaux, C. E. H. Mattoni, D. N. McKinsey, and J. M. Doyle. Magnetically stabilized luminescence of boron nitride. *Journal of Luminescence* **92**, 291–296 (2000).
- [67] Josef R. Parrington, Harold D. Knox, Susan L. Breneman, Edward M. Baum, and Frank Feiner, editors. *Nuclides and Isotopes: Chart of the Nuclides*. 15th edition. General Electric Co. and KAPL, Inc., 1996.
- [68] Brush Wellman Engineered Materials. IF-1/99.8% beryllium metal x-ray windows, batch 0000186007. Material certification provided with windows.
- [69] H. Kudo and K. Tanaka. Recoil ranges of 2.73 MeV tritons and yields of ^{18}F produced by the $^{16}\text{O}(\text{t},\text{n})^{18}\text{F}$ reaction in neutron-irradiated lithium compounds containing oxygen. *Journal of Chemical Physics* **72**, 3049–3052 (1980).
- [70] G. F. Knoll. *Radiation Detection and Measurement*. second edition. Wiley, 1989.
- [71] R. L. Walker. The (α, n) cross section of boron. *Physical Review* **76**, 244–247 (1949).
- [72] T. W. Bonner and L. M. Mott-Smith. The energy spectra of the neutrons from the disintegration of fluorine, boron and beryllium by alpha-particles. *Physical Review* **46**, 258–268 (1934).
- [73] J. R. Cameron, N. Suntharalingam, and G. N. Kenney. *Thermoluminescent Dosimetry*. University of Wisconsin Press, 1968.
- [74] R. Golub and S. K. Lamoreaux. Neutron electric-dipole moment, ultracold neutrons and polarized ^3He . *Physics Reports* **237**, 1 (1994).
- [75] A. R. Reddy, K. Ayyangar, and G. L. Brownell. Thermoluminescence response of LiF to reactor neutrons. *Radiation Research* **40**, 552–562 (1969).
- [76] G. A. Lubyanskii, V. V. Styrov, and V. A. Sokolov. Dosimetric properties of boron nitride. *Atomnaya Energiya (translated in Soviet Atomic Energy)* **31**, 119–122 (1971).
- [77] S. Larach and R. E. Shrader. Multiband luminescence in boron nitride. *Physical Review* **104**, 68–73 (1956).

-
- [78] A. Katzir, J. T. Suss, A. Zunger, and A. Halperin. Point defects in hexagonal boron nitride. I. EPR, thermoluminescence, and theramilly-stimulated-current measurements. *Physical Review B* **11**, 2370–2377 (1975).
- [79] H. Kobayashi, H. Shibata, and S. Tagawa. Continuum optical radiation produced by H^+ and He^+ ion bombardment of pyrolytic boron nitride. *Nuclear Instruments and Methods in Physics Research B* **90**, 556–559 (1994).
- [80] S. Radhakrishna and M. Rama Krishna Murthy. Thermoluminescence of irradiated poly(methyl methacrylate). *Journal of Polymer Science* **15**, 987–993 (1977).
- [81] A. Tomita. Thermoluminescence of γ -irradiated polytetrafluoroethylene. *Journal of the Physical Society of Japan* **28**, 731–738 (1970).
- [82] 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**, 3998–3999 (1998).
- [83] Magda L.G.Leite, Luciana Maia, Ana C.M. Rodrigues, and Edgar D. Zanotto. Vitreous Materials Laboratory, Federal University of Sao Carlos. Personal communication.
- [84] V. V. Lopatin and F. V. Konusov. Energetic states in the boron nitride band gap. *Journal of the Physical Chemistry of Solids* **53**, 847–854 (1992).
- [85] V. Wagner, H. Friedrich, and P. Wille. Performance of a high-tech neutron velocity selector. *Physica B* **180–181**, 938–940 (1992).
- [86] B. Hammouda. Multidisk neutron velocity selectors. *Nuclear Instruments and Methods in Physics Research A* **321**, 275–283 (1992).
- [87] J. R. D. Copley. Acceptance diagram analysis of the performance of multi-disk neutron velocity selectors. *Nuclear Instruments and Methods in Physics Research A* **332**, 511–520 (1993).
- [88] P. A. Egelstaff and R. S. Pease. The design of cold neutron filters. *Journal of Scientific Instruments* **31**, 207–212 (1954).
- [89] A. K. Freund. Cross-sections of materials used as neutron monochromators and filters. *Nuclear Instruments and Methods* **213**, 495–501 (1983).
- [90] W. H. Zachariasen. *Theory of X-Ray Diffraction in Crystals*. John Wiley and Sons, 1945.
- [91] C. G. Darwin. *Philosophical Magazine* **27**, 315–333 (1914).

-
- [92] C. G. Darwin. *Philosophical Magazine* **43**, 800–829 (1922).
- [93] M. L. Crow. Characterization of mica crystals as reflectors for an ultracold neutron doppler converter. *Physica B* **241–243**, 110–112 (1998).
- [94] A. N. Goland, J. H. Sondericker, Jr., and J. J. Antal. Synthetic mica as a monochromator for long wavelength neutrons. *Review of Scientific Instruments* **30**, 269–271 (1959).
- [95] T. W. Dombeck, J. W. Lynn, S. A. Werner, T. Brun, J. Carpenter, V. Krohn, and R. Ringo. Production of ultra-cold neutrons using doppler-shifted bragg scattering and an intense pulsed neutron spallation source. *Nuclear Instruments and Methods* **165**, 139–155 (1979).
- [96] 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 using doppler-shifted bragg reflection at a pulsed-neutron source. *Physics Letters A* **75**, 223–224 (1980).
- [97] J. W. Lynn, W. A. Miller, T. W. Dombeck, G. R. Ringo, V. E. Krohn, and M. S. Freedman. Time-of-flight spectrometer for ultra-cold neutrons. *Physica B* **120**, 114–117 (1983).
- [98] K. T. Forstner, L. Passel, and C. F. Majkrzak. NIST center for neutron research. Personal communication.
- [99] F. M. Zelenyuk, K. N. Zaitsev, A. V. Timakov, I. N. Anikin, S. I. Matveev, V. E. Zhitarev, and S. V. Stepanov. Fluorophlogopite single crystal as a long-wave neutron monochromator. *Instruments and Experimental Techniques* **16**, 399–401 (1973).
- [100] K. Inoue, T. Kanaya, Y. Kiyanagi, S. Ikeda, K. Shibata, H. Iwasa, T. Kamiyama, N. Watanabe, and Y. Izumi. A high-resolution neutron spectrometer using mica analyzers and the pulsed cold source. *Nuclear Instruments and Methods in Physics Research A* **309**, 294–302 (1991).
- [101] C. Shafthäutl. *Journal für Praktische Chemie* **21**, 129 (1841).
- [102] J. E. Fischer and T. E. Thompson. Graphite intercalation compounds. *Physics Today*, 36–45 July (1978).
- [103] M. S. Dresselhaus and G. Dresselhaus. Intercalation compounds of graphite. *Advances in Physics* **30**, 139–326 (1981).
- [104] A. Boeuf, A. Freund, R. Caciuffo, A. Hamwi, and Ph. Touzain. Intercalated pyrolytic graphite for neutron monochromatisation. *Synthetic Metals* **8**, 307–312 (1983).

- [105] DB21. Web page. (<http://www.ill.fr/YellowBook/DB21/>).
- [106] N. Akuzawa, T. Fukisawa, T. Amamiya, and Y. Takahashi. Stability of potassium-graphite intercalation compounds in an oxygen atmosphere. *Synthetic Metals* **7**, 57–63 (1983).
- [107] N. Akuzawa, T. Seino, A. Yugeta, T. Amemiya, and Y. Takahashi. Potassium-graphite intercalation compounds prepared from carbon materials of different heat-treatment temperature and their reaction with water vapor. *Carbon* **25**, 691–696 (1987).
- [108] P. C. Eklund and H. G. Smith. Possible application of graphite-SbCl₅ intercalation compounds as long-wavelength neutron monochromators. *Journal of Applied Crystallography* **17**, 400–402 (1984).
- [109] J. E. Fischer, J. W. Milliken, A. Magerl, H. J. Kim, and A. W. Moore. Intercalated graphite neutron optical components: low-mosaic acceptor compounds with large d-spacings. *Nuclear Instruments and Methods in Physics Research A* **300**, 207–209 (1991).
- [110] H. Zabel. Ruhr universität bochum. Personal communication.
- [111] H. Zabel and P. C. Chow. Intercalated graphite. *Comments on Condensed Matter Physics* **12**, 225–251 (1986).
- [112] M. E. Misenheimer and H. Zabel. Stage transformation and staging disorder in graphite intercalation compounds. *Physical Review Letters* **54**, 2521–2524 (1985).
- [113] Ph. Touzain and A. Hamwi. De-intercalation and second intercalation of potassium into a highly oriented pyrolytic graphite. *Synthetic Metals* **23**, 127–132 (1988).
- [114] A. Hérold. Recherches sur les composés d’insertion du graphite. *Bulletin de la Société Chimique de France* **187**, 999–1012 (1954).
- [115] J. R. Anderson. NIST optical shop. personal communication.
- [116] R. Stedman. *Review of Scientific Instruments* **31**, 1156 (1956).
- [117] NIST Center for Neutron Research. <http://www.ncnr.nist.gov>.
- [118] Robert Williams. NIST Center for Neutron Research. Personal communication.
- [119] Y. T. Cheng and D. F. R. Mildner. A sensitive high resolution neutron imaging detector. *Nuclear Instruments and Methods in Physics Research A* **454**, 452–459 (2000).

-
- [120] G. M. Seidel. Brown university. Personal communication.
- [121] Klaus Habicht. *Szintillationen in flüssigem Helium - ein Detektor für ultrakalte Neutronen (Scintillations in liquid helium - a detector for ultracold neutrons)*. PhD thesis, Technischen Universität Berlin, 1998.
- [122] L. Duband, L. Hui, and A. Lange. Thermal isolation of large loads at low temperature using kevlar rope. *Cryogenics* **33**, 643–647 (1993).
- [123] G. Hartwig and S. Knaak. Fibre-epoxy composites at low temperatures. *Cryogenics* **24**, 639–647 (1984).
- [124] G. E. Childs, L. J. Ericks, and R. L. Powell. Thermal conductivity of solids at room temperature and below. Technical Report NBS Monograph 131, US DoC, 1973.
- [125] A. Visser. A 10,000A 1000 VDC Solid State Dump Switch. Technical Report FERMILAB-Conf-91/114, Fermilab, 1991.
- [126] V. Weidner and J. Hsia. Reflection properties of pressed polytetrafluoroethylene powder. *Journal of the Optical Society of America* **71**, 856–861 (1981).
- [127] K. C. Peng, S. K. Sahu, H. C. Huang, K. Ueno, Y. H. Chang, C. H. Wang, and W. S. Hou. Measurement of radiation damage on an optical reflector. *Nuclear Instruments and Methods in Physics Research A* **384**, 544–546 (1997).
- [128] G. A. Morton, H. M. Smith, and R. Wasserman. Afterpulses in photomultipliers. *IEEE Transactions on Nuclear Science* **14**, 443–448 (1967).
- [129] P. B. Coates. The origins of afterpulses in photomultipliers. *Journal of Physics D* **6**, 1159–1166 (1973).
- [130] P. B. Coates. A theory of afterpulse formation in photomultipliers and the prepulse height distribution. *Journal of Physics D* **6**, 1862–1869 (1973).
- [131] M. Yamashita, O. Yura, and Y. Kawada. Probability and time distribution of afterpulses in GaP first dynode photomultiplier tubes. *Nuclear Instruments and methods* **196**, 199–202 (1982).
- [132] C. C. Lo and B. Leskovar. Afterpulse time spectrum measurement of RCA 8850 photomultiplier. *IEEE Transactions on Nuclear Science* **30**, 445–450 (1983).
- [133] Saint-Gobain Industrial Ceramics Carborundum. Certificates of analysis. Lot numbers X50021, X50022, X50024, X50034, X50036, and X50038.

-
- [134] K. J. Coakley. Neutron lifetime experiments using magnetically trapped neutrons: optimal background correction strategies. *Nuclear Instruments and Methods in Physics Research A* **469**, 354 (1998).
- [135] G. L. Yang and K. J. Coakley. Likelihood models for two-stage neutron lifetime experiments. *Physical Review C* **63**, 014602 (2000).
- [136] K. J. Coakley. Statistical planning for a neutron lifetime experiment using magnetically trapped neutrons. *Nuclear Instruments and Methods in Physics Research A* **406**, 451–463 (1998).
- [137] R. Golub. HMI. Personal communication.
- [138] A. G. Kharitonov, V. V. Nesvizhevsky, A. P. Serebrov, R. R. Taldaev, V. E. Varlamov, A. V. Vasilyev, V. P. Alfimenkov, V. I. Lushchikov, V. N. Shvetsov, and A. V. Strelkov. Preliminary results of neutron lifetime measurements with gravitational UCN trap. *Nuclear Instruments and Methods in Physics Research A* **284**, 98–100 (1989).
- [139] M. G. D. van der Grinten, J. M. Pendlebury, D. Shiers, C. A. Baker, K. Green, P. G. Harris, P. S. Iaydjiev, S. N. Ivanov, and P. Geltenbort. Characterization and development of diamond-like carbon coatings for storing ultracold neutrons. *Nuclear Instruments and Methods in Physics Research A* **423**, 421–427 (1999).
- [140] M. F. Makela. *Polarized Ultra-Cold Neutron Transport for Neutron Beta Decay Asymmetry Measurements*. PhD thesis, Virginia Polytechnic Institute and State University, 2002.
- [141] P. D. Bangert, M. D. Cooper, and S. K. Lamoreaux. Enhancement of superthermal ultracold neutron production by trapping cold neutrons. *Nuclear Instruments and Methods in Physics Research A* **410**, 264–272 (1998).
- [142] Carlo Egon Heinrich Mattoni. The precision measurement of the neutron lifetime using magnetically trapped neutrons: Marginally trapped neutrons and fluorescent time constants, 1995. Undergraduate Thesis.
- [143] Nathan G. Woodard, Eric G. Hufstедler, and Gregory P. Lafyatis. Photon counting using a large area avalanche photodiode cooled to 100 K. *Applied Physics Letters* **64**, 1177–1179 (1994).
- [144] Jeffrey J. Fox, Nathan Woodard, and Gregory P. Lafyatis. Characterization of cooled large-area silicon avalanche photodiodes. *Review of Scientific Instruments* **70**, 1951–1956 (1999).
- [145] Paul Hink. Burle electro-optics, inc., sturbridge, ma. Personal communication.

- [146] K. Tsuchiya, K. Egawa, K. Endo, Y. Morita, N. Ohuchi, and K. Asano. Performance of the eight superconducting quadrupole magnets for the TRISTAN low-beta insertions. *IEEE Transactions on Magnetics* **27**, 1940–1943 (1991).
- [147] K. Tsuchiya, N. Ohuchi, A. Terashima, K. Nakayama, K. Egawa, S. Takasaki, K. Asano, R. Hoshi, and T. Yamagiwa. Superconducting quadrupole magnets for the TRISTAN low-beta insertion. In *11th International Conference on Magnet Technology (MT-11)*, Volume 1, 171–175, 1990.
- [148] K. Sasaki, K. Tsuchiya, T. Ogitsu, N. Ohuchi, Y. Doi, T. Haruyama, K. Kawamata, S. W. Kim, T. Nakamoto, N. Kimura, T. Shintomi, K. Tanaka, S. Kato, A. Terashima, A. Yamamoto, G. A. Kirby, and V. T. Kovachev. Test results of spare TRISTAN insertion quadrupole magnet with yoke added in superfluid helium. *IEEE Transactions on Applied Superconductivity* **7**, 570–573 (1997).
- [149] K. Tsuchiya. KEK. Personal communication.
- [150] Hoare, Jackson, and Kurti, editors. *Experimental Cryophysics*, Chapter 7 by A. Wexler. Butterworths 1961.
- [151] G. E. Bacon and R. D. Lowde. Secondary extinction and neutron crystallography. *Acta Crystallographica* **1**, 303–314 (1948).
- [152] R. Golub. I hate diffraction... always have, always will. unpublished notes on dynamical diffraction, 2000.
- [153] K. Coakley, Z. Chowdhuri, M. S. Dewey, and J. Richardson. Monte carlo modelling of neutron scattering rocking curve data. To be submitted to the *Journal of Applied Crystallography*.
- [154] J. Richardson. *Accurate Determination of Thermal Neutron Flux: A critical Step in the Measurement of the Neutron Lifetime*. PhD thesis, Harvard University, 1993.
- [155] M. S. Dewey. NIST. Personal communication.
- [156] H. S. Sommers, J. Gregory Dash, and L. Goldstein. Transmission of slow neutrons by liquid helium. *Physical Review* **97**, 855–862 (1955).
- [157] A. K. Freund. Reflectivity of pyrolytic graphite monochromators. Technical Report 80FR33S, Institute Laue-Langevin, 1980.