

ABSTRACT

DEES, ERIC BRANDON. On the Use of Ultracold Neutrons for Exploration of Electroweak Interactions and 5th Force Contributions to Gravitation. (Under the direction of Dr. Albert Young).

Ultracold neutrons are neutrons that are low enough energy as to be contained completely within a material system. By exploiting their transportability, and their long wavelength yet neutral structure, we have directly measured angular correlations associated with their decay. These correlations allowed determination of fundamental constants characterizing the charged weak interactions: the vector and axial-vector coupling constants. Further, by restricting the neutrons vertical velocity phase space, eigenstates of the gravitational interaction can be created between material plates. We explore the theoretical reach of utilizing these eigenstates, through use of a separated oscillatory fields resonance method, coupled with thin film microfluidic arrays. We find the potential to allow exploration of perturbations to the Newtonian gravitational interaction at intermediate distance scales of order 1-100 microns. Microfluidic arrays for this procedure were developed and characterized as proof of principle, and as precursors to future experiments.

© Copyright 2020 by Eric Dees

All Rights Reserved

On the Use of Ultracold Neutrons for Exploration of Electroweak Interactions and 5th Force
Contributions to Gravitation

by
Eric Brandon Dees

A dissertation submitted to the Graduate Faculty of
North Carolina State University
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Physics

Raleigh, North Carolina
2020

APPROVED BY:

Albert Young
Committee Chair

John Thomas

Thomas Schaefer

Carl Koch

DEDICATION

To Isaac, my best bud, #1 Goober, who always tells me I can do it, and helps however he can. You can dedicate your dissertation to me, when it's your turn.

To Mommykins, for teaching me about Schrödinger's cat when I was 7, encouraging me in nearly all my goals, and probably like, 1 other thing, I guess. Love you.

To Albert, my sensei, and great friend. I've never met someone as excited about physics, and it's been contagious. You pushed me to the edge of my capabilities, let me struggle for answers and be independent to show I was capable, and were allowing, understanding, and permissive when I needed the help and was at the end of my rope. You're clearly smart, well liked by most, but maybe the most important: You care, more deeply and sincerely about your students, than any other advisor I've seen, or heard about from others. Your conception of time is warped, but you've set an amazing example to aspire towards, and you'll be a benchmark against which I'll hold myself for the rest of my life. Thanks.

To Robbie, for being the glue. The immovable stone of peace and tranquility, the complement to Albert's fervent existence, you helped me see that sometimes it's time to give up and move on, and sometimes you have to stay up until 6am until it works.

To Mike, Craig, Preston, Kory, and Saagi; if I ever needed a couch to sleep on, or someone to bail me out, I feel like I could have relied on any of you. Luckily it never came to that, and hopefully it won't.

To Andy and Mark, for becoming secondary advisers, and teaching me about the area, and the politics that goes into it, when you had no responsibility to do so whatsoever.

To Hannah, for the help and encouragement she provided when she was around.

To everybody else I forgot that I've worked with along the way; pretty much everyone has been kind, and helpful, and eager to contribute what they know. Science is a lot more fun when people are happy to do it, and enjoy helping one another.

BIOGRAPHY

I was born in Pensacola, Florida, to my parents Jerry and Lisa Dees, and later my only brother, Christopher Dees, was born. We moved to North Carolina when I was 9, to a small town called Aberdeen, for my father's work. 2 years later, we moved again, this time to the Raleigh area, where I've lived since, though in different houses and apartments along the way.

My father died when I was 16, from pancreatic cancer, which affected my life greatly, in good and bad ways. I placed in multiple math competitions, as well as in physics at Science Olympiad. In my senior year of high school, I'd run out of math classes to take, and thus enrolled here at NCSU halftime to take Calculus III and Differential Equations.

I obtained enough AP credits to begin undergrad as a sophomore, and subsequently completed my B.S. in physics in 3 years. While I was scheduled to double major in Applied Mathematics, I was 3 courses short, and was only awarded a minor. I graduated Summa Cum Laude, and as a member of the University Honors Program.

As I did not yet know in which field I'd like to specialize, and had a very limited time window for applications, I chose to stay at NCSU to continue my graduate career. My son Isaac was born that summer. I dabbled in computational nucleosynthesis, and expressed momentary interest in biophysics, but ultimately was drawn to experimental nuclear, both by the combination of particle physics, theory and experiment, and of course my advisor personally, Albert Young.

2 years after starting grad school, my ex-wife left my son and I, to move to a different state. I was awarded sole physical and legal custody. I've cared for him largely alone for the last 6 years. It has been difficult, yet rewarding, and has reminded me throughout what is truly important.

TABLE OF CONTENTS

LIST OF TABLES.....	vii
LIST OF FIGURES.....	viii
Chapter 1.....	1
1.1 A Preface on Context of Contributions.....	1
1.2 A Brief History of Neutrons.....	2
1.3 The Weak Force.....	3
1.4 Bilinear Covariants.....	5
1.5 Parity Violation.....	8
1.6 Angular Correlations.....	10
1.7 The Fermi Potential.....	14
Chapter 2.....	20
2.1 Neutron Production.....	20
2.2 Neutron Transport.....	28
2.3 Polarization.....	30
2.4 β -decay Measurements.....	31
2.5 Backscattering.....	39
Chapter 3.....	42
3.1 Depolarized Fraction.....	42
3.2 Spin Flipper.....	43
3.3 Pre-Shutter Polarimetry.....	54
3.4 Run Types.....	58
3.4.1 Depolarization (D).....	58
3.4.2 Shutter and Flipper Trapping.....	64
3.4.3 Polarization (P).....	65
3.4.4 Depolarization Calibration (DC).....	69
3.4.5 Spin Flipper Efficiency (SFE).....	73
3.4.6 Depolarization Evolution (DE).....	74
3.5 Data Quality Checks.....	75
3.5.1 Runs 21k-22k: Anomalously High Shutter Trapped Decay Rates.....	75
Chapter 4.....	83
4.1 Computational Resources.....	83

4.2 Experimental Reproduction	87
4.3 Systematic Corrections	97
4.2.1 Feeding	97
4.2.2 Leaking	102
4.2.3 Conditioning.....	102
4.4 Results.....	103
Chapter 5.....	109
5.1 Gravitational Eigenstates	109
5.2 Ramsay’s separated oscillatory fields	112
5.3 5 th force modifications	117
5.4 Microfluidic arrays	119
5.5 Detector Package	137
Chapter 6.....	143
6.1 Summary	143
6.2 Future Efforts	144
REFERENCES.....	149
APPENDICES	155
Appendix A.....	156
Appendix B	164
Appendix C	174
Appendix D.....	178
Appendix E	182

LIST OF TABLES

Table 1.1: Chart of Standard Model particles.[7]	5
Table 1.2: Example values for typical effective Fermi Potentials	17
Table 1.3: A shortened list of example energy ranges for neutron temperature classifications ..	18
Table 3.1: Various count rates during standard D-type run cycles. Includes Gate Valve loading counts (near the source), SCS decay trap counts for correlating transport efficiency, counts in the SCS monitor during the wrong spin draining portion, SCS monitor background rates, and the wrong-spin counts reaching the switcher. Includes both spin states, and both run years.	61
Table 3.2: Fits to lifetime for populations in both shutter and flipper trapped scenarios, for each run year and load state.	65
Table 3.3: Various counts for P type runs, for each load state and run year. These include the normalizing Gate Valve rate at the beginning of the system(near the source), the SCS decay trap counts for correlating transport efficiency, the SCS monitor counts during the draining period, the backgrounds for the SCS detectors, and the number of counts reaching the switcher detector.	66
Table 3.4: 0 th order calculations of the depolarized fractions ξ for each of the 2 run years, and loading states, determined from experimental data, with no systematic corrections included.	69
Table 4.1: Table of various independent parameters to be tuned, and ways to measure their effects, in reproducing the experimental behavior of the guide system. Note; \pm denotes those values which must be found separately for the on and off loaded spin states.	88
Table 4.2: Table of corrected polarization values, within 68% (1σ) confidence interval.....	106
Table 4.3: Table of corrected polarization values and respective uncertainties for each year and loading state, within 95% (2σ) confidence interval.	106
Table 5.1: Q_n, k values for first 5 eigenstates, from equation (5.10), in $m - 1$. These represent the degree of coupling between eigenstates, induced by an oscillating floor potential, sometimes referred to as the “overlap integrals”	115

LIST OF FIGURES

Figure 1.1: The Feynman Diagram for neutron β -decay.	4
Figure 1.2 Beta decay asymmetry variables for A, indicating the relative angle θ_e between σ , the initial polarization direction of the neutron, and p_e , the direction of the emitted electrons momentum.....	11
Figure 1.3 Measurements of A_0 , and the associated uncertainties, over time. [18].....	12
Figure 1.4 Current intersections of parameter constraints for V_{ud} , λ , and the neutron lifetime. Due to the large discrepancy between beam and bottle, the lifetimes from each experimental type are separated into distinct bands. The triple intersection of the UCN (bottled) lifetimes, the V_{ud} from $0^+ \rightarrow 0^+$ superallowed decays, and more recent determinations of λ shows internal consistency between these sets at the time of writing. [35].....	14
Figure 1.5 Cartoon of incoming wave (from left to right), and exponentially decaying into, while sampling, a finite potential barrier greater than its own energy.	15
Figure 2.1 Plot of UCN density at LANL source over time, indicating relative improvements of combined effects from increased particle incidence, as well as improved geometries and source preparations. Reproduced from [55].....	22
Figure 2.2 Illustration of a neutron population thermally distributed as Maxwell-Boltzmann. Because a specific amount of energy is required for the downscattering, only a small ribbon of available phase space is able to scatter into the UCN regime. Additionally, this transition is symmetric and reversible, such that longer residency times within the downscattering material will result in upscattering of UCN populations, demonstrating the requirement for rapid removal of UCN from their source after production. From [60].....	24
Figure 2.3 Cartoon of UCNA transport geometry, connecting UCN source to decay trap. Neutrons are produced in the bottom right, and proceed upwards and to the left.	25
Figure 2.4 Equilibrium loading rates at the entrance to the UCN transport system. Clusters demonstrate decreases in population densities over time, interspersed with discontinuous jumps to higher values, associated with so called “melt + refreeze” procedures to renew the source’s effectiveness.	27
Figure 2.5 Rendering of current experimental setup in “Area B” at Los Alamos National Laboratory. UCNA experiment consists of the purple and green components, along with the upstream guides connecting the UCN source. A UCN nEDM experiment (left) and the UCN lifetime experiment (right) are also depicted.....	30
Figure 2.6 Diagram of transport system and decay trap. From [35].....	33
Figure 2.7 Electron detector package diagram. Identical packages were located on either side of the superconducting solenoidal decay trap. Electrons (traveling from right to left) first transit the decay trap windows (which reflect neutrons), through the Multi-Wire Proportional Counters, and end in the plastic scintillator. Reproduced from [18].....	35
Figure 2.8 Position dependent energy calibration maps for each of the 4 quadrants in the PMT array (East package). Normalized via neutron activated Xe gas isotropically filling the system.....	36
Figure 2.9 Example energy distributions during sealed source calibration runs. From [35].....	37

Figure 2.10 Example of position detection signal during calibration run. Reproduced from [35].....	37
Figure 2.11 Example of energy uncertainty envelopes produced via sealed source calibration runs. Reproduced from [35].....	38
Figure 2.12 Energy reconstruction and fit value to A_0 parameter. Reproduced from [35].....	39
Figure 2.13 Depiction of various backscattering event types. Reproduced from [35]	41
Figure 2.14 Example of PMT signal comparison with and without coincidence requirement for MWPC detection. [35]	41
Figure 3.1 Diagram for the initial and rotated coordinate systems	45
Figure 3.2 Picture of water cooled Adiabatic Fast Passage Spin Flipper, from [63].....	50
Figure 3.3 Resonance spectra of spin flipper when coupled to network analyzer inductively. These produce much weaker minima of other modes, as the primary is selected preferentially.....	51
Figure 3.4 Resonance signal output from direct coupling of network analyzer to spin flipper, through end-driven mode. Multiple resonances can be observed as minima in impedance, along with discontinuities in phase angle.....	52
Figure 3.5 Examples of transmission efficiencies during frequency tuning in crossed polarizer analyzer configuration. Reproduced from [63]	53
Figure 3.6 (Top) Measured/Simulated values of holding field magnitude along propagation direction. (Bottom) Zoomed in view of holding / oscillating fields at the location of spin flip transition. Reproduced from [63].....	54
Figure 3.7 Example plots of switcher detector signals during complementary DT and R run types for determining polarization prior to decay trap shutter installation.....	57
Figure 3.8 Switcher detector signals during standard Depolarization run types, performed at the end of each β octet segment. These are analogous to the DT run types of the pre-shutter polarimetry.	63
Figure 3.9: An example of a D-type run signal in the switcher detector. The 5 phases are illustrated, though everything preceding the Wrong-Spin Unload is not used in calculations, as it represents a background reduction process.....	64
Figure 3.10 Switcher detector signals during the Polarization run type. These ex-situ runs are performed independent of standard β -octet run sequences, but due to the higher particle counts during these runs, fewer are required to be completed than for the D-type runs.	67
Figure 3.11 Depolarized Fraction measured during multiple Depolarization Calibration runs. The line represents the deviation from linear evolution of lower efficiencies. One also observes the normalization scheme by the gate valve becomes in worse and worse agreement, due to its inability to account for the changing loading rate efficiency in the trap.	71
Figure 3.12 On loaded case of depolarized fraction dependence during Depolarization Calibration run types. Exhibits similar behavior to that of off loaded (See Above)	71
Figure 3.13 Measured loading rate seen in the SCS monitor during depolarization calibration runs. Both the on and off case have been combined into 1 graph, on either side. An empirical curve of the form $x + 1 - x$ has been overlaid on the data, simply as a method for comparing to computation. It is now known that	

	<p>this is likely more accurately modeled as a Beta function, but sufficient data was not available to confirm or test this hypothesis.</p>	72
Figure 3.14	<p>Loading rates measured in the decay trap during the loading phase, while varying Spin Flipper power. 40W corresponds to typical input power during on-loading, and 0W corresponds to the simple off-loading case. Note; Due to the $1/B^4$ dependence of the efficiency, most of the interesting behavior actually lies between 0-1W, as can be seen by the fitted curves maximum being produced slightly past the y-axis. Data is fit to CDF of Beta function.....</p>	74
Figure 3.15	<p>Graph of wrong spin counts in Switcher detector for Depolarization Evolution runs, normalized to 1 for the standard D run holding time at 30s.</p>	75
Figure 3.16	<p>Decay rates fit to shutter trapped populations. Excessively high values between 21k-22k indicated a possible issue with the decay trap.</p>	76
Figure 3.17	<p>Initial observations of anomalous shutter trapped lifetimes. Fit lifetimes suggested experimental issues, but inclusion of spurious runs in analysis masked the degree of the effect at the time.....</p>	77
Figure 3.18	<p>Ratio of Iron-foil monitor to Gate-Valve monitor signals during loading.</p>	78
Figure 3.19	<p>Ratio of Iron-foil monitor to SCS monitor signals during loading.....</p>	79
Figure 3.20	<p>Background counts in SCS monitor vs run #.</p>	80
Figure 3.21	<p>Background counts in β monitors vs run #.</p>	80
Figure 3.22	<p>Signal ratios for P type runs for (left to right) 1) SCS monitor vs gate-valve monitor, 2) right-spin counts normalized by gate-valve monitor, and 3) right-spin counts normalized by scs monitor.....</p>	82
Figure 4.1	<p>Depolarized Fraction evolution inside decay trap over time, in the old "Depolarization Trapping" and "Reload" run types. As the Reload type injects more intentionally depolarized UCN, the increase in the depolarized fraction can be observed. However, this procedure relied on the assumption of re-establishing an equilibrium, producing double the depolarized population, which may not be the case.....</p>	85
Figure 4.2	<p>Depolarized fraction evolution over time, for the spin flipper on loading state. See Figure 4.1</p>	86
Figure 4.3	<p>Plot of UCN density inside decay trap vs. height. As one expects, the population is slightly denser at lower regions, from the gravitational potential energy being similar in scale to the kinetic energies.</p>	87
Figure 4.4	<p>Simple cartoon depicting 2 potential solution regions (circles), being inaccessible by exploration of 1-D scans from anywhere within one solution region (rectangles). It is these cases in which correlation studies are required.</p>	91
Figure 4.5	<p>Example of dependence of depolarized fraction versus depolarization per bounce. While the DPB will not significantly affect the lifetimes, it can control the produced polarization levels significantly. It is therefore a 2nd order tuning parameter, as it has effects on fewer observables, and should be tuned near last.</p>	93
Figure 4.6	<p>Comparison of changing parameters of the UCNA decay trap shutter. (Top) Varying the material potential demonstrates a stronger dependence on the flipper-on loaded state, compared to flipper off. (Bottom) Varying the loss per bounce of the shutter has little to no effect, on either population. This is attributable to the lower field at the shutter position, resulting in an acceleration of the on-loaded states. This also demonstrates the signature of selectively-</p>	

	differential component-parameter combinations, identification of which are essential to narrowing parameter scan scopes.	95
Figure 4.7	Example of behavior at an island of stability. Top and Bottom representing varied alternative parameter choices, such as rectangular guide material potential. Varying the Loss per Bounce (of the decay trap in this case) was able to bring a few lifetimes in near agreement, but some were fundamentally separated (Ex, top-middle vs. bottom-right). The regions of these discrepancies can hint towards the regions that may need adjusting, and subsequent compensation.	96
Figure 4.8	Qualitative plot, showing the potential increase in net depolarized population after sealing the decay trap.	101
Figure 4.9	Simulated depolarized population in decay trap over time. While being the majority population, or depolarization being generated primarily from inefficiency in state selection at time of loading, population behavior remains similar during the shutter trapped holding time. However, when fed primarily through depolarization interactions within the trap, the populations exhibit non exponential decay, causing significant overestimations of the truly present counts.	101
Figure 4.10	Plot of χ^2 for both load states simultaneously. Contours represent χ^2 of 10, and 100. Region of agreement is widened, as the requirement for a single parameter set to satisfy both spin conditions simultaneously is more rigorous.	107
Figure 4.11	Plots of χ^2 values for simulations versus experimental observations, for a grid of spin flipper efficiency and depolarization per bounce, for each of the 2 spin states. Contours represent χ^2 values of 10, and 100. These demonstrate the very narrow available parameter space to reproduce those values determined from experiment, and constrain the uncertainty for the requisite systematic corrections.	108
Figure 5.1	Reproduced from [68]. The first 5 eigenstates in earth's gravitational well for a neutron. Note that each is the same Airy function, only scaled for normalization, and translated to subsequent Airy Zeroes.	111
Figure 5.2	Sketch from [68], illustrating an apparatus for selecting $n=1$ eigenstates. Higher order states sample the absorber more completely and are removed from the ensemble rapidly.	112
Figure 5.3	Examples of transition probability widths for Rabi resonances (drift distance \rightarrow 0)	113
Figure 5.4	Proposed experimental geometry, reproduced from [70], for a Separated Oscillatory Fields measurement of gravitational eigenstates.	116
Figure 5.5	Transition probably for monochromatic neutron beam vs. driving frequency. Near perfect transition rates can be achieved, and the high frequency oscillations allow more precise determination of resonance widths.	116
Figure 5.6	Ramsey transition probability (as above), but with $v^2 dv$ velocity spectrum. Note that maximum transition probability has decreased from unity, and the resonance width is widened, allowing for less accurate determination than with a monochromatic source.	117
Figure 5.7	Example curves for energy shifts of 1 st eigenstate of neutron in gravitational potential, versus thickness of membrane dividing neutrons from attractor	

material, for varying values of the interaction range λ in a Yukawa-like potential. Values normalized for 0 thickness(optimum sensitivity). “1” corresponds to fundamental length scale of neutrons in earth’s gravitational field, as seen in eq. 5.2.....	118
Figure 5.8 Diagram of early design for winding array(grey), on quartz block(red), with silicon sidewalls(green).	123
Figure 5.9 [100] array with array walls etched out. The remaining silicon nitride strands coating the walls can still be seen vertically, and the yellow tinted film is still intact.....	124
Figure 5.10 Properly etched fluidic array, with compressive stress in silicon nitride film. Waves/ribs can be seen on the film between the side walls. Additionally, the [100] orientation over-etching can be seen most visibly as the small colored rectangles at the end of each silicon finger.....	126
Figure 5.11 SEM of [110] silicon wafer etching pattern producing deep grooves with high contrast ratio, reproduced from [82].	127
Figure 5.12 Improper bonding of fused silica wafer to quartz substrate via PDMS. The diffraction rings of varying thicknesses can be observed in the center, and the incomplete coverage is observable at the corners. Input ports were required to be blocked, to prevent flow of PDMS into input/output channels, to prevent clogging.	131
Figure 5.13 Completed microfluidic array with highest contrast ratio. In addition to have significantly more and smaller supporting silicon fingers, the relative area of channels to walls was increased by a factor of 10, permitted by use of the vertical sidewall etching.	135
Figure 5.14 Contact profilometry of wafer after bonding to quartz block substrate via activated PDMS. The large spike was simply physical debris. Oscillations show the variation across each channel only constituting total variations of $\sim 1\mu\text{m}$, over the span of 22mm.	136
Figure 5.15 Contact profilometry of bonded wafer across flat portion (excluding central array). This demonstrates that the total variation in height over 5cm distance is less than $0.5\mu\text{m}$	136
Figure 5.16 Mercury being flowed through completed and bonded array from top to bottom. Mercury was introduced through simple elevation of source input, and can be seen winding halfway across one of the channels from left to right.	137
Figure 5.17 Example of single detectors count rate increasing with voltage.	139
Figure 5.18 Count rates of triple coincidence enforced on detectors, in presence of Gd source.	139
Figure 5.19 Plot of normalized counts rates of 3 detectors, and the signal resulting from their coincidence, versus lower threshold turns on an analog to digital converter. While optimal ranges were available between ~ 5 -15 turns, the triple coincidence signal remains quite consistent even at the lower values, implying better independence from this factor’s exact value.	140
Figure 5.20 Background counts versus varying types of shielding, with and without presence of ZnS strip facing detectors. Note: data proceeds in time from left to right, as experiment was conducted during beam shut off.	142

Figure 6.1 Relative drain curves over time for the existing shutter (left), and a shutter coated in a neutron absorber, such as TPX(right). The greatly reduced time required to clean the system could lead to quicker turnaround for opening the shutter, reducing the degree of evolution that would be permitted. 145

Chapter 1

Ultracold Neutrons

1.1 A Preface on Context of Contributions

This thesis represents efforts on two separate but related experiments. The first is work concerned with the experimental collaboration “UCNA”. While UCNA had been in development for years prior to my arrival in 2012, I had the opportunity to gain operating knowledge of the system and measurements with ultracold neutrons (UCN) while characterizing the expected neutron polarization resident in the beta decay spectrometer. I was responsible for the 2011 to 2013 run cycle, with the beta asymmetry results published in 2018. This work involved implementing checks for consistency with previous years, characterizing system parameters, and conceptually developing future runs aimed at quantifying required systematic corrections.

The help from others in developing and operating UCNA cannot be overstated here, and I am indebted to the entire collaboration, but especially to the graduate students whose research focused on guide production and characterization: Mark Makela[1], Adam Holley[2] and Brittney Vorndick [3].

Lastly, my other involvement was more solitary. I worked nearly alone to develop and create microfluidic arrays, of targeted flatness scales, for use in an already constructed experimental apparatus, as well as the possibility as a substrate in future iterations as well. So, while I review the efforts and background associated with the required experimental apparatus, my direct work is that involved only with the microfluidic array component.

1.2 A Brief History of Neutrons

In the early 1900's, atoms were assumed to abide by the Rutherford model, in which the atom has a dense, positively charged nucleus surrounded by a diffuse cloud of negatively charged electrons. However, as the beginnings of quantum mechanics were formulated, and more experiments involving radioactivity were performed, it became increasingly clear that discrepancies with the Rutherford model were irreconcilable. In 1932 James Chadwick observed the effects of alpha radiation incident on beryllium, and was able to conclude, from the behavior of the subsequent decay of carbon-13, that a neutral particle of mass similar to the proton must exist[4]. In only 3 years' time from his publication, he was awarded the Nobel Prize in Physics for the discovery of the neutron.

Within a matter of months following Chadwick's article, Werner Heisenberg published multiple papers proposing models in which the nucleus consisted of both protons and neutrons[5]. Simultaneously, Enrico Fermi contextualized the decay of the free neutron into the language of quantum mechanics, and further posited a yet unseen particle, known now as the neutrino.

With the newfound understanding of the constituents making up the nucleus, a vast number of advancements became possible. Calculations of nuclei and their emergent symmetries became consistent in developing a framework for understanding. The nature of isotopes and their varying stabilities helped explain previously conflicting results. A cohesive description, uniting the various forms of radiation into a (comparatively) simple, shared origin, was created; the entire fields of nuclear and particle physics were born.

While the myriad practical and theoretical benefits provided by these fields can hardly be overstated, we shall direct our focus to the neutron itself. Although first described nearly a

century ago, free neutrons are still widely utilized and studied today. Their applications range from use as imaging tools for materials science, to activation of isotopes for use in medical imaging.

However, their primary importance (for the purposes of this dissertation) is in no less than exploring the fundamental nature of reality. Neutrons compose over half the mass of every stable isotope after helium, correctly suggesting that an understanding of neutrons is required to explain nearly all matter encountered in daily life. As neutrons are neutral, they interact minimally with the electromagnetic force, allowing high precision tests of the remaining strong, weak, and gravitational forces with which they interact. Additionally, neutrons are the 2nd lightest form of baryonic matter (after the proton), which makes them perfect candidates for probing the quark-quark interactions with minimal convolutions from nuclear structure.

1.3 The Weak Force

Neutrons bound within nuclei are stable, and indeed are required for maintaining stability of all nuclei above hydrogen. Free neutrons, however, decay with a characteristic lifetime of $\tau_n = 880.2 \pm 1.0\text{s}$ [6], or approximately 15 minutes. The decay occurs via the pathway



This reaction was directly responsible for inspiring the first postulation of the neutrino by Pauli [7], being required for maintaining energy conservation of the (then) observable decay products. It also did not go unnoticed that the masses of the proton and neutron were very nearly equal, and this observation, coupled with properties of many other nuclear reactions, led to the development of isospin as a quantum number, which states the neutron and proton are able to be

effectively treated as different states of a spin doublet. Ongoing tests are being performed to evaluate the extent to which this isospin symmetry is broken in superallowed $0^+ \rightarrow 0^+$ nuclear decays [8]. With the discovery of quarks, physicists were able to describe the neutron “decay” instead as a transition: a single down quark transforming to an up quark, via the charged current interaction, emitting a (virtual) W^- vector boson [see Figure 1.1]. This mechanism indeed describes the behavior of all β^- decays, and is encapsulated within the generalized $SU(2) \times U(1)$ electroweak theory [9].

The quark transformation possibilities are encapsulated within the Cabibbo-Kobayashi-Maskawa (CKM) matrix, a 3×3 matrix which represents the amplitudes for each generation of “up” type quarks (up, charmed, top) to transform into any of the 3 generations of “down” type quarks (down, strange, bottom). These mixing elements are a result of the eigenstates in the weak interaction being slightly different than the mass eigenstates; this then means that each quark is in a superposition of weak eigenstates, and that although each decay pathway within a single generation is highly preferred, decays through the additional 2 routes are still possible. Neutron decay is represented by the transition of a down to an up quark and is the 1st entry in the CKM matrix.

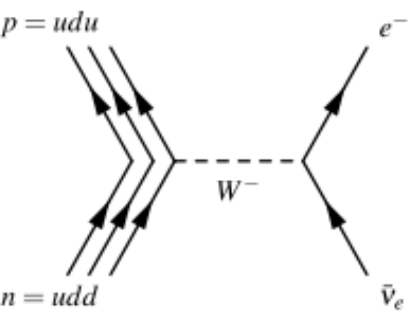
$$\begin{bmatrix} d' \\ s' \\ b' \end{bmatrix} = \begin{bmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{bmatrix} \begin{bmatrix} d \\ s \\ b \end{bmatrix} \quad (1.2)$$


Figure 1.1: The Feynman Diagram for neutron β^- -decay.

To begin to quantify β decay, and the interactions that are responsible for it, requires an analysis of all possible forms the interactions may take.

The Standard Model of physics includes 17 fundamental particles (**Error! Reference source not found.**). These consist of 12 fermions (spin $\frac{1}{2}$) (6 quarks, 3 leptons and their associated neutrinos), and 5 bosons (integer spin), responsible for the propagation of forces. Each fundamental fermion also has an anti- counterpart, with identical mass, but opposite charge.

Table 1.1 Chart of Standard Model particles.[10]

		three generations of matter (fermions)			interactions / force carriers (bosons)	
		I	II	III		
mass		$\approx 2.2 \text{ MeV}/c^2$	$\approx 1.28 \text{ GeV}/c^2$	$\approx 173.1 \text{ GeV}/c^2$	0	$\approx 125.09 \text{ GeV}/c^2$
charge		$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	0	0
spin		$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1	0
		u up	c charm	t top	g gluon	H higgs
	QUARKS	$\approx 4.7 \text{ MeV}/c^2$	$\approx 96 \text{ MeV}/c^2$	$\approx 4.18 \text{ GeV}/c^2$	0	
		$-\frac{1}{3}$	$-\frac{1}{3}$	$-\frac{1}{3}$	0	
		$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1	
		d down	s strange	b bottom	γ photon	
	LEPTONS	$\approx 0.511 \text{ MeV}/c^2$	$\approx 105.66 \text{ MeV}/c^2$	$\approx 1.7768 \text{ GeV}/c^2$	$\approx 91.19 \text{ GeV}/c^2$	
		-1	-1	-1	0	
		$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1	
		e electron	μ muon	τ tau	Z Z boson	
		$< 2.2 \text{ eV}/c^2$	$< 1.7 \text{ MeV}/c^2$	$< 15.5 \text{ MeV}/c^2$	$\approx 80.39 \text{ GeV}/c^2$	
		0	0	0	± 1	
		$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	1	
		ν_e electron neutrino	ν_μ muon neutrino	ν_τ tau neutrino	W W boson	
					GAUGE BOSONS VECTOR BOSONS	SCALAR BOSONS

1.4 Bilinear Covariants

The goal of Dirac was to produce a wave equation to describe the behavior of relativistic electrons; to do so required imposing the preservation of Lorentz invariance, which the Schrödinger equation did not. The first step, then, was to replace the energy with the relativistic

expression $E = \sqrt{p^2 c^2 + m^2 c^4}$, and to incorporate the operator formalism of quantum mechanics. The issue then was to develop a relation that would allow the energy operator squared to produce only 2nd order derivatives in each of the 4 relativistic coordinates (x,y,z,t). To exclude the cross terms one would ordinarily predict, Dirac realized that the only solution was for the coefficients of any 2 partial derivatives to anti-commute. Thus, it was required that the coefficients of the differential operators be *matrices*, and likewise the wavefunction necessitated multi-dimensional parameterization as well.

The coefficients for each of the partial derivatives must then satisfy the anticommutation property

$$\gamma^\mu \gamma^\nu + \gamma^\nu \gamma^\mu \equiv \{\gamma^\mu, \gamma^\nu\} = 2\eta^{\mu\nu} I_4 \quad (1.3)$$

where $\eta^{\mu\nu}$ is the Minkowski metric

$$\eta^{\mu\nu} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{bmatrix} \quad (1.4)$$

Eq. (1.3) expresses that all cross-derivative terms (such as $\partial_x \partial_y$ and $\partial_y \partial_x$) must cancel, while all 2nd order derivatives (such as ∂_x^2) should survive. The inclusion of the Minkowski metric indicates the distinction between the spacelike and time-like dimensions of the 4-vectors.

The smallest dimensionality γ matrices capable of satisfying this relationship span 4x4 matrices. Likewise, the solution for the wavefunction must form at least a 4-dimensional, possibly complex, vector space. Due to the transformation properties of the wavefunction, solutions require a class of complex vectors known as Spinors. Spinors behave similarly to vectors; however, a critical distinguishing feature is that they are NOT returned to an identical state after undergoing a 360-degree rotation, but instead are transformed into a completely inverse state, carrying a negative sign from the complex phase. This phase can only be observed

when pitted against an observation with an initially coherent state that foregoes this rotation, but nonetheless, represents something totally unlike anything seen in everyday life; fermions do not appear the same after being spun around once, and won't appear the same until 2 full circles have been completed.

The Dirac equation produces the solution for any spin 1/2 fermion, and thus is appropriate for the neutron, though it was derived specifically for the electron. The Dirac equation reads

$$i\hbar\gamma^\mu\partial_\mu\psi - mc\psi = 0 \quad (1.5)$$

Any set of γ matrices satisfying the commutation relation of Eq. (1.3) are valid for use in the Dirac equation, however it is customary to use the following definitions as the standard basis:

$$\gamma^0 = I \otimes \sigma_3 = \begin{bmatrix} I & 0 \\ 0 & -I \end{bmatrix} \quad (1.6)$$

$$\gamma^i = \sigma^i \otimes i\sigma_2 = \begin{bmatrix} 0 & \sigma^i \\ -\sigma^i & 0 \end{bmatrix}$$

With the traditional definition of the Pauli matrices

$$\sigma^1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad \sigma^2 = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \quad \sigma^3 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad (1.7)$$

This can also be referred to as the Dirac basis. Alternative constructions exist, such as the Weyl basis, and the Majorana basis, which can each be useful for simplifying expressions to specific problems, especially those involving chirality (a word describing whether a particle spins clockwise or counterclockwise, compared to the direction in which it is traveling). It is also the case that defining $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3 = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}$ turns out to be useful, specifically for systems in eigenstates of chirality, often termed left- and right-handed. In fact, it is the inclusion of this γ^5 term that permits the breaking of parity symmetry in all weak decays, of course including that of the free neutron.

Finally, we define a tensor element, $\sigma^{\mu\nu} = -\frac{i}{4}[\gamma^\mu\gamma^\nu - \gamma^\nu\gamma^\mu]$, which allows us to complete the set required to construct our Clifford algebra. With this $\sigma_{\mu\nu}$ and the γ matrices we define here, there are sufficient terms to form an irreducible representation of any Lorentz invariant Lagrangian. Since the 4x4 size of any interaction matrix is required, 16 linearly independent quantities are needed to wholly cover the basis. The following is the set of all possible forms of interactions in the Dirac theory, along with the number of unique terms each has;

$$\begin{aligned}
\bar{\psi}\psi: & \text{Scalar}(1) \\
\bar{\psi}\gamma^\mu\psi: & \text{Vector}(4) \\
\bar{\psi}\sigma^{\mu\nu}\psi: & \text{Tensor}(6) \\
\bar{\psi}\gamma^5\gamma^\mu\psi: & \text{Axial vector}(4) \\
\bar{\psi}\gamma^5\psi: & \text{Pseudoscalar}(1)
\end{aligned} \tag{1.8}$$

1.5 Parity Violation

In 1956 T.D. Lee and C. N. Yang published a literature review in which they identified that no tests of parity violation had been done with regards to any weak decays. As legend goes, C.S. Wu, an expert in both low temperature and polarization physics, was so excited by this untested prospect, she cancelled her imminent vacation to instead perform an experiment to support or refute the principle[11]. By observing the beta decay of ^{60}Co sources, Wu was able to demonstrate that, in fact, the formerly held principal of parity conservation was indeed broken in radioactive decays[12]. Further, and perhaps more profoundly, it was found that the violation was maximal; not only did left- and right- handed decays not occur at the same rate, but in fact NO right-handed decays were observed; this was corroborated by Goldhaber et al. at Brookhaven in the following year [13].

Given the possible forms of interaction from Eq. (1.8), one is able to construct the matrix decay amplitude for neutron decay, utilizing the transformation properties of each of the 5 separate quantities. To construct maximal parity violation requires an interaction of the so-called Vector minus Axial Vector (V-A) form, appearing in a form such as $\bar{\psi}_1 \gamma^\mu (1 - \gamma^5) \psi_2$ (omitting coupling constants and normalization) for leptons. Probing the limits on this V-A law has been shown to be possible up to $\mathcal{O}(10^{-4})$ through measurements of neutron decay observables [14]. While mathematically an interaction of the form obeying Scalar and Pseudoscalar decays is permissible to also generate parity violation, these decay paths would require new, unobserved gauge bosons, and no decays of this form have been positively detected [15], and limits can be placed from invoking neutrino mass constraints as well [16].

However, the decay is further complicated by quarks' participation in the strong interaction, and therefore may take a form that is not pure V-A in nature. Instead, additional coupling constants for each of the 2 terms must be included, and appear as $\bar{\psi}_1 \gamma^\mu (C_V + C_A \gamma^5) \psi_2$, where we have introduced the vector and axial-vector coupling constants, respectively. In the most general form possible, including each form from eq (1.8), the transition element may be expressed as

$$\begin{aligned} \langle p | J^\mu | n \rangle = \bar{u}_p & \left[g_V(q^2) \gamma^\mu - i \frac{g_{WM}(q^2)}{2M} \sigma^{\mu\nu} q_\nu + \frac{g_S(q^2)}{2M} q^\mu \right. \\ & \left. + g_A(q^2) \gamma^\mu \gamma^5 - i \frac{g_T(q^2)}{2M} \sigma^{\mu\nu} \gamma_5 q_\nu + \frac{g_P(q^2)}{M} \gamma_5 q^\mu \right] u_n \end{aligned} \quad (1.9)$$

With q the momentum transfer, M the nucleon mass, and the g 's are the form factors for the vector, weak magnetism, induced scalar, axial vector, induced tensor, and induced pseudoscalar interactions. As $q^2 \rightarrow 0$, however, only the vector and axial vector terms dominate, and as such we hereafter restrict ourselves only to these 2 interactions.

With the inclusion of the W boson as a propagator mediating the weak force, and the separation of the hadronic and leptonic transitions, we may then construct the decay amplitude

$$\mathcal{M} = i \overbrace{\bar{\psi}_p \gamma^\mu (g_V + g_A \gamma^5) \psi_n}^{\text{Hadronic}} \overbrace{\widehat{g}_{\mu\nu}}^{\text{Bosonic}} \overbrace{\bar{\psi}_e \gamma^\mu (1 - \gamma^5) \psi_{\nu_e}}^{\text{Leptonic}} \quad (1.10)$$

This expression, though seemingly convoluted, is actually just describing an empirical observation; that the amplitude for neutron decay is broken into 1) a hadronic transition from a neutron to a proton, with an unknown parity violation corresponding to experimentally determined coefficients g_V and g_A of the vector and axial-vector components, and requiring conservation of baryon number, 2) an intermediate gauge boson W to mediate the interaction with the leptonic current, and 3) the leptonic involvement of an electron and electron-neutrino, maximally violating parity in the V-A form, and requiring lepton number conservation.

1.6 Angular Correlations

In 1957 Jackson, Treiman and Wyld first parameterized the various possible correlations for any nuclear beta decay[17]:

$$\begin{aligned} \omega(\langle J \rangle | E_e, \Omega_e, \Omega_\nu) dE_e d\Omega_e d\Omega_\nu \\ = \frac{1}{(2\pi)^5} p_e E_e (E^0 - E_e)^2 dE_e d\Omega_e d\Omega_\nu \left\{ 1 + a \frac{\mathbf{p}_e \cdot \mathbf{p}_\nu}{E_e E_\nu} + b \frac{m}{E_e} \right. \\ \left. + c \left[\frac{1}{3} \frac{\mathbf{p}_e \cdot \mathbf{p}_\nu}{E_e E_\nu} - \frac{(\mathbf{p}_e \cdot \mathbf{j})(\mathbf{p}_\nu \cdot \mathbf{j})}{E_e E_\nu} \right] \left[\frac{J(J+1) - 3\langle \mathbf{J} \cdot \mathbf{j} \rangle^2}{J(2J-1)} \right] \right. \\ \left. + \frac{\langle J \rangle}{J} \cdot \left[A \frac{\mathbf{p}_e}{E_e} + B \frac{\mathbf{p}_\nu}{E_\nu} + D \frac{\mathbf{p}_e \times \mathbf{p}_\nu}{E_e E_\nu} \right] \right\} \end{aligned} \quad (1.11)$$

Following the reasoning laid out in [17], we may observe that for the case of the free neutron, we have multiple pieces of additional information, as well as simplifications, that allow us further insight into the potential of this unique system. First, we note that the neutron is a spin

$\frac{1}{2}$ system, and as such, the term $3\langle(\mathbf{J} \cdot \mathbf{j})^2\rangle = J(J + 1)$, thus cancelling all terms on the second line of eq. (1.11). Further, because of the transition rules for this reaction, and the neutron being a mirror decay (a decay in which the number of neutrons and protons is exchanged, in this case $1 \rightarrow 0$ and $0 \rightarrow 1$), we have that $|M_{GT}|^2 = 3$, and $|M_F|^2 = 1$, where the M's represent the Gamow-Teller and Fermi matrix transition amplitudes [17].

We now turn our attention to the A parameter of eq. (1.11). We may integrate over the corresponding solid angles, and neutrino momenta, to reduce the equation to a form which is only dependent on a single angular correlation coefficient,

$$\omega(E, \theta) \propto 1 + P \cdot A(E) \cdot \beta \cdot \cos \theta \quad (1.12)$$

Where P is the average polarization of the initial neutron population, β the relativistically normalized velocity, θ refers to the angle between the initial neutron polarization and the emitted electron's momentum, and A the (energy dependent) angular correlation coefficient (See **Figure 1.2**).




Figure 1.2 Beta decay asymmetry variables for A, indicating the relative angle θ_e between σ , the initial polarization direction of the neutron, and p_e , the direction of the emitted electrons momentum.

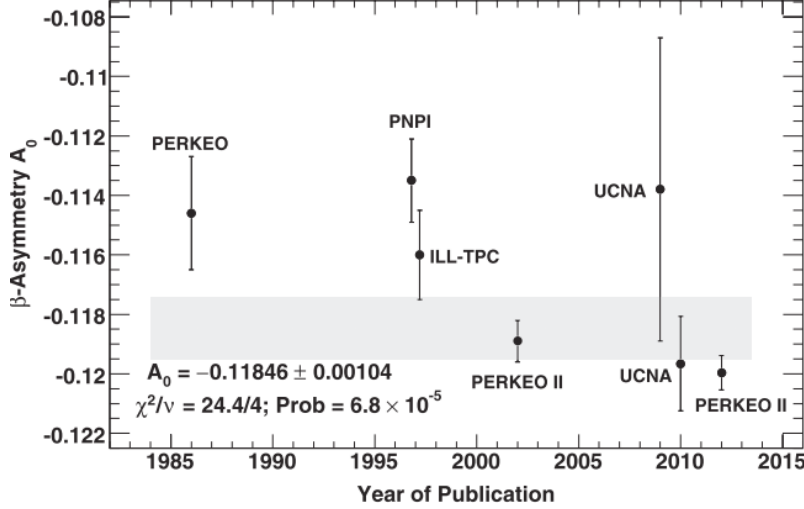


Figure 1.3 Measurements of A_0 , and the associated uncertainties, over time. [18]

By neglecting recoil-order corrections, such as weak magnetism, and V-A interference, one may express the β asymmetry in an energy independent form,

$$\begin{aligned}
 A_0 &= -2 * \left(\frac{\lambda(\lambda + 1)}{1 + 3\lambda^2} \right), \lambda \equiv \frac{g_A}{g_V} \\
 A &= A_0 + A_1 \frac{\epsilon}{Rx} + A_2 R + A_3 Rx \\
 R &= \frac{E_0}{M}, \quad \epsilon = \left(\frac{m_e}{M} \right)^2, \quad x = \frac{E_e}{E_0}
 \end{aligned} \tag{1.13}$$

Showing that a measurement of A alone is sufficient for determining the ratio of the coupling constants in weak decay, along with the expanded form representing the recoil order corrections. We now consider also an expression for the lifetime,

$$\frac{1}{\tau_n} = \frac{G_F^2 m_e^5}{2\pi^3} V_{ud}^2 (1 + 3\lambda^2) f (1 + RC) \tag{1.14}$$

With G_F the Fermi constant, f the phase space factor of $f = 1.6887$, and RC representing recoil order corrections. The neutron lifetime is also highly significant in determining big bang nucleosynthesis dynamics [19]–[21], and while contributions to the LANL UCN τ experiment have been made by the author, this will not be a primary focus for this thesis. Further discussion

of various means of determining the lifetime, and the systematics required therein, can be found elsewhere[22]–[29]. Improvements to, and details of, radiative corrections to neutron- and superallowed- β decays can be found discussed in [30], [31]. It is then easy to see that knowing any 2 between the lifetime, the CKM matrix element, or the β asymmetry, is sufficient for determining the 3rd. Thus, an important cross check is possible by careful determination of any 2 of the neutron-decay observables. As of this publication, current limits on V_{ud} are determined from observations of many $0^+ \rightarrow 0^+$ nuclear decays [32], and G_F is constrained most stringently by muon decay [33].

Prior to 2009, A had only been determined using neutron beam methods. A leading experiment in determining the A coefficient was the PERKEO II collaboration, who utilized cold neutron beams, polarized via a magnetic supermirror apparatus [34]. The leading sources of uncertainty came primarily from their neutron polarization of 99.7%, an extraordinarily good rate. One advantage of UCNA is specifically this aspect, in aiming to improve polarization rates, which will be discussed further in later chapters.

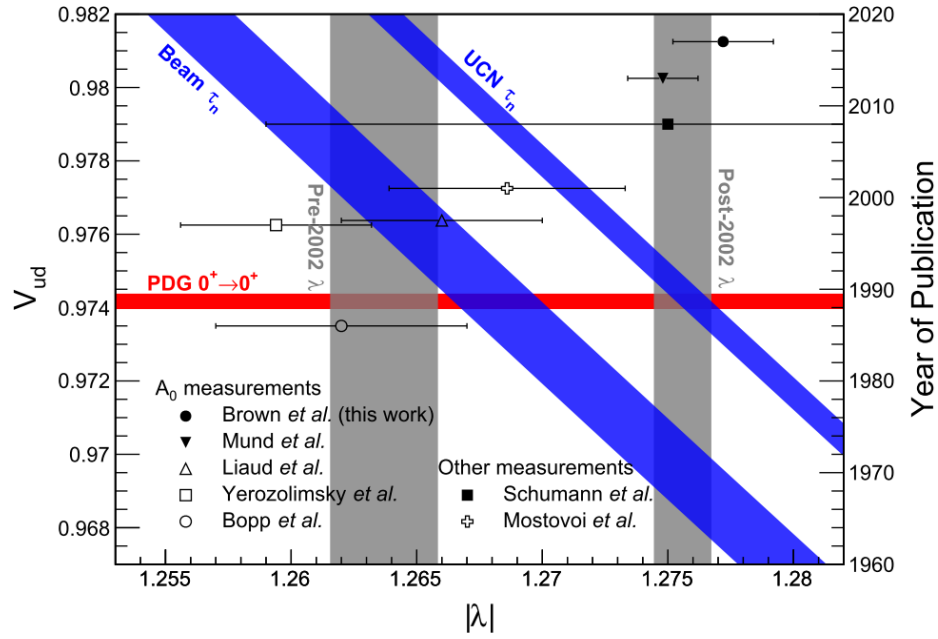


Figure 1.4 Current intersections of parameter constraints for V_{ud} , λ , and the neutron lifetime. Due to the large discrepancy between beam and bottle, the lifetimes from each experimental type are separated into distinct bands. The triple intersection of the UCN (bottled) lifetimes, the V_{ud} from $0^+ \rightarrow 0^+$ superallowed decays, and more recent determinations of λ shows internal consistency between these sets at the time of writing. [35]

1.7 The Fermi Potential

As particles are cooled, their wavelength proportionally increases, according to the de Broglie relation

$$\lambda = \frac{h}{p} \quad (1.15)$$

As this wavelength becomes long compared to a significantly shorter-range interaction, the system will approach the s-wave scattering limit, of angular momentum transfer 0. We also observe that at far separation distances of an incident particle to a scattering center, the scattering amplitude must be constant, as any angular dependence would create a nonzero angular momentum transfer.

It was these realizations that led Fermi to predict the possibility of UCN [36]. In keeping with the qualitative picture given by the de Broglie relation, the phenomenon is best thought of as in an optical model, with light waves being totally internally reflected.




Figure 1.5 Cartoon of incoming wave (from left to right), and exponentially decaying into, while sampling, a finite potential barrier greater than its own energy.

We start by writing down the Schrödinger equation for a finite potential square well

$$-\frac{\hbar^2}{2\mu}\nabla_r^2\psi + [E - V(\mathbf{r})]\psi = 0, \quad V(r) = \begin{cases} -V_0 & r < \rho \\ 0 & r > \rho \end{cases} \quad (1.16)$$

where μ is the reduced mass of the neutron-nuclei system, and r the separation distance. As mentioned, we seek a V such that the scattering amplitude is constant.

Using the first-order Born approximation,

$$\begin{aligned} f(\theta) &= -a = -\frac{\mu}{2\pi\hbar^2}\langle\mathbf{k}_f|V|\mathbf{k}_i\rangle = -\frac{\mu}{2\pi\hbar^2}\int d^3r V(r)e^{i(\mathbf{k}_i-\mathbf{k}_f)\cdot\mathbf{r}} \\ &= \frac{2\mu}{3\hbar^2}V_0\rho^3 \end{aligned} \quad (1.17)$$

Where a represents the constant value of the scattering amplitude we require. It can be shown through a series of estimations that allowing the integral to be a simple δ function is valid for cases in which $kf \ll 1$. Combining this δ function with the reality of having a large number of these tiny scattering centers, we can reduce the triple integral to a simple, constant term representing the number density of the scatterers. This then let's us define an effective potential function, known as the Fermi Potential[37]:

$$V_F = \frac{2\pi\hbar^2}{m_n} Na \quad (1.18)$$

where N is the number density of the material, a is the (positive) neutron scattering length of the constituent nuclei, and we've used the assumption of $\mu \rightarrow m_n$ for nuclei with much greater mass than the neutron, which is nearly all of them. This then leads to the strange result that for large enough wavelengths (low enough energy) of incident neutrons, an originally *attractive* potential appears instead as a small *repulsive* potential in aggregate.

With this (constant) potential defined, we may revisit the traditional finite square well problem again. We first assume a 1-d system, in which an infinitely long potential step is encountered by a rightward propagating neutron. We may then simply guess our solutions,

$$\begin{aligned} \psi_1 &= e^{ik_1x} + R e^{-ik_1x} \\ \psi_2 &= T e^{ik_2x} \end{aligned} \quad (1.19)$$

With $E = \frac{\hbar^2 k_1^2}{2m}$, and $E - V = \frac{\hbar^2 k_2^2}{2m}$. The incident wave in region 1 is then reflected with probability $|R|^2$, and transmitted with probability $|T|^2$. We then impose boundary conditions for the continuity of both the wavefunctions, and derivatives of wavefunctions, at the potential boundary.

$$\begin{aligned}\psi_1(0) &= 1 + R = \psi_2(0) = T \\ \psi_1'(0) &= ik_1e^{ik_1x} - ik_1Re^{-ik_1x} = \psi_2'(0) = ik_2Te^{ik_2x}\end{aligned}\tag{1.20}$$

Substituting in the top line for the bottom to replace T, we solve for R and arrive at

$$R = \frac{k_1 - k_2}{k_1 + k_2} = \frac{\sqrt{E} - \sqrt{E - V}}{\sqrt{E} + \sqrt{E - V}}\tag{1.21}$$

If $E < V$, then k_2 becomes complex, leading to the initial assumption of a traveling wave solution, producing instead an exponentially decaying function which penetrates the potential region, but is ultimately reflected with $|R|^2 = 1$

$$\psi_2 = (1 + R)e^{ik_2x} = \left(\frac{2\sqrt{E}}{\sqrt{E} + i\sqrt{V - E}} \right) e^{-\sqrt{\frac{2m(V-E)}{\hbar^2}}x}\tag{1.22}$$

Table 1.2 Example values for typical effective Fermi Potentials

Material	V_F (neV)	σ_{tot} (barns)
Ni ⁵⁸	335	44
Be	252	1.4
Fe	210	30
Cu	168	43.5
PTFE	123	X
Pb	83	2.0
Al	54	2.8
H ₂ O	-14.7	X
Ti	-48	58

It is common practice to refer to the energy scales of particles by their temperature, and neutrons are no exception. To do so, we describe our ensemble with the Maxwell-Boltzmann distribution,

$$N_i = N_0 e^{-\frac{E_i}{kT}}\tag{1.23}$$

with E_i the energy of the i 'th state, T the temperature in kelvin, and k the Boltzmann constant

$k = 1.38 * 10^{-23} \frac{m^2kg}{s^2K}$. We can then begin to discuss the associated temperatures of the neutron

distributions as those associated with kinetic energies $\sqrt{p^2c^2 + m^2c^4} = kT$. For example, with the highest known Fermi Potential (Ni⁵⁸) of 335 neV, we define the upper limit of ultracold neutrons as those with a temperature of ~4mK.

Table 1.3 A shortened list of example energy ranges for neutron temperature classifications

Neutron energy	Associated Temperature (Kelvin)	Energy range
0 – 350 neV	0 - 0.004	Ultracold neutrons
0 – 0.025 eV	0 – 300	Cold neutrons
0.025 eV	300	Thermal neutrons
1 – 100 eV	1E+4 – 1E+6	Slow neutrons
100 eV – 1 MeV	1E+6 – 1E+10	Intermediate neutrons
≥ 1 MeV	$\geq 1E+10$	Fast neutrons

Further, we can get an estimate for the velocity distribution of the neutrons. Though produced superthermally (not allowed to reach thermal equilibrium), their repeated interaction with the guide walls at room temperature means that their quantum mechanical chance of tunneling must reduce to the low energy tail of the Maxwellian distribution.

$$f(v)d^3v = \left(\frac{m}{2\pi kT}\right)^{\frac{3}{2}} e^{-\frac{m|v|^2}{2kT}} d^3v \quad (1.24)$$

Which, under isotropic conditions, allows re-expression as

$$f(v)dv = \left(\frac{m}{2\pi kT}\right)^{\frac{3}{2}} 4\pi v^2 e^{-\frac{m|v|^2}{2kT}} dv \quad (1.25)$$

For sufficiently small velocities compared to the temperature of the surroundings, the exponential term can be neglected, and we obtain the so-called $v^2 dv$ distribution, describing the functional form of the range of velocities. The question then is; how does one produce sufficient UCN, when the low energy tail described by the velocity distribution occupies such a small fraction of the total phase space?

Chapter 2

UCNA

2.1 Neutron Production

Neutrons can be produced from multiple means, but the primary sources for research purposes are either from spallation, or fission, with spallation producing far more efficient neutron economy[38]. At production, neutrons are produced at MeV scale energies, yet are quickly cooled by thermal equilibration via interactions with any of a number of “moderators”: materials with large neutron scattering cross sections, often containing large fractions of hydrogen. To achieve “cold” neutron temperatures, repeated interaction with moderators cooled to LHe temperatures is sufficient. However, further cooling to the required millikelvin temperatures must be accomplished through superthermal processes. These can include being doppler shifted into lower energies via interaction with receding, neutron reflecting turbine blades, such as at the Institut Laue Langevin (ILL)[39], scattering from superfluid ^4He [40]–[48], or as will be discussed in this work, via SD_2 downscattering[49]–[51].

This work primarily focuses on analysis of depolarization in the UCNA experiment. The UCNA experiment aims to measure an angular asymmetry coefficient A of free neutron beta decay, by way of observing trapped UCN populations. The UCNA experiment is in Los Alamos, NM, on a beamline sourced from a linear accelerator hosted at the Los Alamos Neutron Science Center for Experimentation (LANSCE). This user facility utilizes a half-mile long linear accelerator to first accelerate, and then direct, a beam of 800 MeV protons to various beamlines, including uses such as neutron scattering, medical isotope production, and of course, our UCN source [52]. For UCNA, these protons are made to impinge upon a helium-cooled tungsten

target, 12cm in length. Because of the high energy of the beam, along with the large atomic number and high material density of tungsten, multiple neutrons (typically 10-20 per incident proton [38]) are liberated from the tungsten nuclei, in a process called spallation. As opposed to fission, in which unstable nuclei may fragment into typically 2-3 pieces, spallation involves the impingement of high energy particles which impart sufficient kinetic energy to blow apart the nucleus, splintering it into dozens of constituent particles. Additionally, this high incidence of total energy causes large amounts of heating, which if not cooled would lead to rapid physical degradation of the target. Therefore, a large bottleneck in the production of neutrons is the ability to cool this target sufficiently, permitting higher fluxes of neutrons to be produced.

A UCN source was developed for the UCNA experiment, which exploits UCN production in SD_2 . After initial prototype tests and R&D [53], a production source was developed and operated from 2006 to 2015 [54]. An upgrade was recently performed, increasing source output by roughly a factor of 5 [55]. The experimental data here were taken between 2011 and 2013.

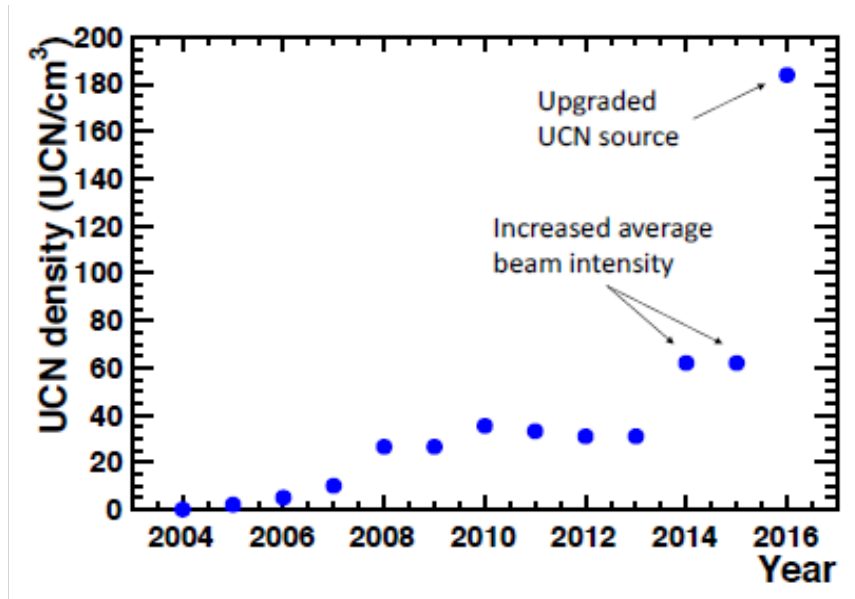


Figure 2.1 Plot of UCN density at LANL source over time, indicating relative improvements of combined effects from increased particle incidence, as well as improved geometries and source preparations. Reproduced from [55]

For use in this experiment, the neutrons must undergo multiple steps of “moderation”, or cooling of particles via scattering, to become usable. Some moderation occurs in the thermal beryllium reflector, which surrounds both the tungsten target, and the bottom of the UCN source. However, moderation to energies appropriate for UCN production in SD₂ occurs in helium-vapor cooled polyethylene beads, cooled to 150K [54]. The produced cold neutrons are roughly within thermal equilibrium with this ‘pre-moderator’. Repeated scattering from the polyethylene beads quickly chills the neutrons to the material’s temperature, corresponding to a reduction of energy by 11 orders of magnitude for the fastest of the population.

The second phase of slowing the neutrons is accomplished by superthermal down-scattering from a solid ortho-deuterium (SD₂) source. The ortho orientation is the ground state for D₂, and consists of total nuclear spin 0 or 2, with rotational angular momentum value of 0. The ortho- orientation is required to minimize up-scattering caused by the transition from J=1→0, which is related to phonon modes within the crystal [56]. At room temperature in its

equilibrium distribution, D_2 exists with roughly 33% para- state, with the remaining 2/3 in the ortho- state. Conversion to the ortho- orientation is thus required for significant UCN production. This transformation is induced by exposing the gas to a ferric chloride salt, which preferentially selects for the low energy state [57]. The reduction of the para- state typically achieved is to roughly 1% of the total deuterium population and has a relaxation time of multiple months before needing re-exposure [58].

After production of the UCN, they are then directed upwards through a 1m vertical guide, to further remove kinetic energy from the population, which received a ~ 100 neV velocity boost upon exiting the SD_2 source. As the accelerator is pulsed, it has been found the optimal method for injection of the UCN into the system is through exploiting a so called “flapper”. This plate, coated in Ni^{58} , alternates between being parallel and perpendicular to the upward direction, allowing neutrons from below past the plate, and then closing to serve as a floor for the neutrons to then bounce from, due to the high likelihood of upscattering or absorption once returned to the D_2 [59]. This then allows a higher density to be concentrated in the upper levels of the containment, and thus more neutrons are able to enter the horizontal guide section that leads to the remainder of the experiment.

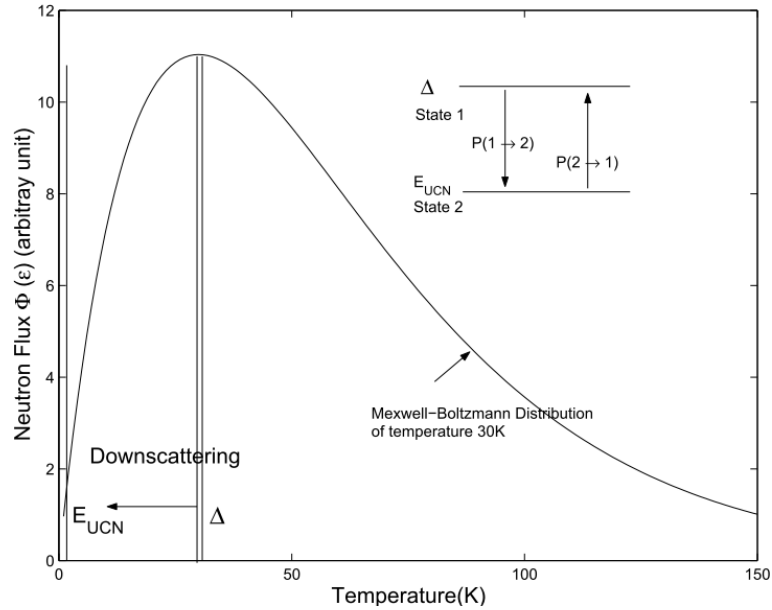


Figure 2.2 Illustration of a neutron population thermally distributed as Maxwell-Boltzmann. Because a specific amount of energy is required for the downscattering, only a small ribbon of available phase space is able to scatter into the UCN regime. Additionally, this transition is symmetric and reversible, such that longer residency times within the downscattering material will result in upscattering of UCN populations, demonstrating the requirement for rapid removal of UCN from their source after production. From [60]

After entering the horizontal guide section, the neutrons then encounter a thin foil of zirconium. This foil permits the separation of the vacuum present in the remaining sections of guide, from the deuterium source. The foil is also placed inside a strong magnetic field of 6T (PPM), which accelerates one spin state of the neutrons through the foil, for higher penetration and lower losses. Then, the neutrons encounter a 30° bend in the guide system, which permits only those neutrons of sufficiently low energy to continue towards the decay trap, with the incidentally higher energy population exiting through the copper bend.




Figure 2.3 Cartoon of UCNA transport geometry, connecting UCN source to decay trap. Neutrons are produced in the bottom right, and proceed upwards and to the left.

Following this bend, the neutrons transit a switcher mechanism, which connects them to the downstream portion of the guide. In the alternate mode, the switcher mechanism disconnects all upstream components, and reorients to connect the downstream components, including the decay trap, to a neutron detector, labeled as the switcher monitor. Next, the neutrons enter the AFP magnet, with a peak field of 7T. The guides in this section are uniquely composed of diamond-like carbon (DLC) coating of a quartz tube, to permit the transmission of radio-frequency oscillations to be induced inside the magnet. The 7T field serves as a kinematic trap for the high-energy spin state, which will be discussed in detail later.

After the strong peak field of the AFP magnet, neutrons encounter a highly tailored portion of the magnetic field, utilizing 10 shimming coils to achieve a highly uniform, and shallow, magnetic field gradient. This area is located at a region of 1T, where an adiabatic fast passage (AFP) spin flipper is utilized. This spin flipper is tuned to the resonance energy for the transition between spin states of the neutrons, similar to NMR methods, using a birdcage RF

resonator. It can be chosen to be energized or not during loading, allowing control of the spin state introduced to the rest of the system.

After spin selection, the neutrons make their way to the decay trap, a 4.4m long superconducting solenoid (SCS), set to the same magnetic field magnitude at which the spin flip occurred, 1T. This ensures that each spin state receives no preferential boost or deceleration from the difference in magnetic fields, partially mitigating a systematic correction required for transport differences between the 2 states. The neutrons are incident at 90° to the SCS orientation, impinging on a small region of the decay trap as they exit the narrow guides, and begin to populate the system towards equilibrium. Each end of the decay trap is also capped with foils, to separate the detector gases from the vacuum of the guide system. The foils also serve to prevent loss of the neutrons, as well as preventing decay in a field expansion region. However, these foils also come with their own systematic corrections, accounting for leading order terms with respect to energy loss and backscattering due to the electron's interaction with the materials. As such, these foils are made as thin as possible, to minimize these contributions.

One interesting effect observed during runs was the steady decline in UCN production. An identified solution was a melt and refreeze procedure, in which the SD_2 was allowed to warm to its gaseous state, and then repressurized and cooled in a specific and controlled fashion. In doing so, the structure of the crystalline D_2 ice was controlled as well, and multiple procedures for the ideal formation pattern were tested and optimized [61]. This process also identifies the “source” of the problem. During running, the impinging high energy neutrons affect the quality of the SD_2 crystal through heating, producing so called D_2 frost layers, and the deformation of the lattice resulted in higher loss rates within the crystal. This procedure was repeated every few octets, as needed.

Additionally, source degradation over the long term also occurred. This was attributable partially to the slow accumulation of contaminant gases. Due to the gases higher freezing point, and the requirement that the D₂ be an isolated system for safety concerns, as well as expense, complete removal of the contamination required moderate levels of deconstruction, which were not feasible during run periods. It is postulated that the introduction of these gases could have occurred through minute injections of gas during actuation of the gate valve, or by minute flow through possible pinhole(s) in the PPM foil, responsible for separating the D₂ containment volume.

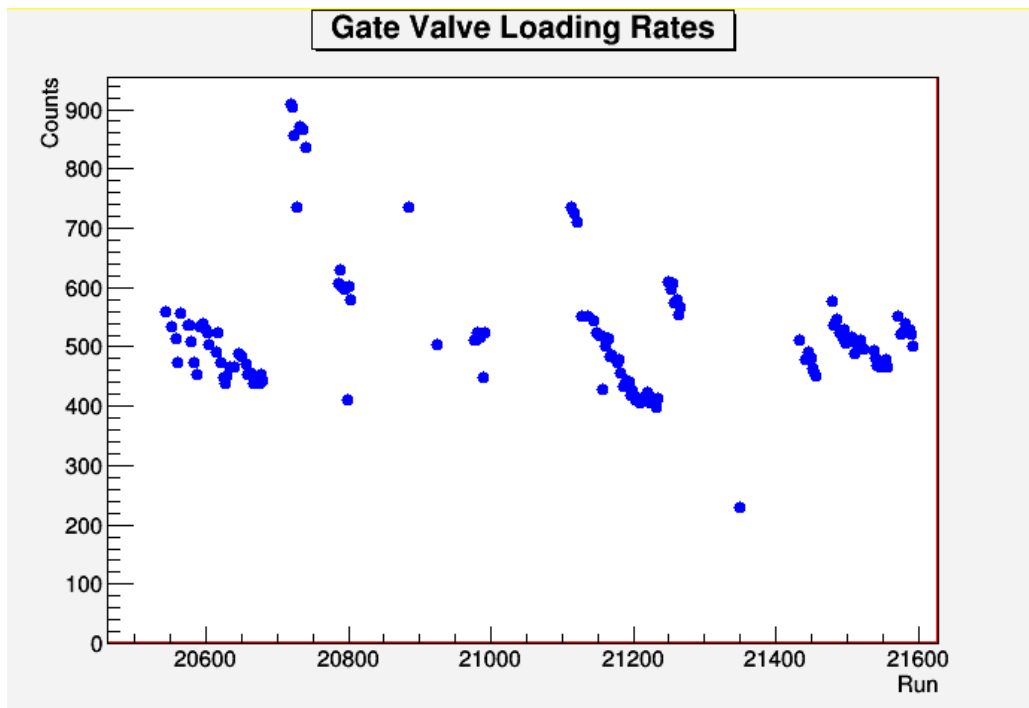


Figure 2.4 Equilibrium loading rates at the entrance to the UCN transport system. Clusters demonstrate decreases in population densities over time, interspersed with discontinuous jumps to higher values, associated with so called “melt + refreeze” procedures to renew the source’s effectiveness.

2.2 Neutron Transport

Neutrons in the ultracold regime offer an advantage in that they can be transported via guide walls. However, there are further methods of controlling the trajectories of the neutrons, namely through both the gravitational and magnetic interactions.

Although neutrons are neutral, they do have a nonzero magnetic moment, allowing manipulation via magnetic field gradients. This interaction is commonly characterized by the interaction Hamiltonian $\mathcal{H} = \mu \cdot B$, with μ the magnetic moment of any given particle. In the case of the neutron,

$$\mu = -1.913 \mu_N = -9.66 \times 10^{-27} \text{J} \cdot \text{T}^{-1} = 60.3 \text{ neV} \cdot \text{T}^{-1} \quad (2.1)$$

with μ_N the nuclear magneton, equal to $\mu_N = \frac{e\hbar}{2m_p}$. This magnetic moment expresses that for a given magnetic field, we have a latent energy of alignment or antialignment of 60 neV for each 1T field strength. The negative sign indicates that, contrary to ordinary magnets, the low energy state of the neutron is when the spin axis is *anti*-aligned to the magnetic field. This can lead to confusing use of standard terminology, and will generally be taken that the $|-\rangle$ state is this low energy, anti- aligned state, and likewise $|+\rangle$ will refer to the higher energy, aligned state, in keeping with standard practices of referencing 2-state systems.

When compared to the average energy values of various Fermi Potentials, one finds a magnetic field strength of $\sim 3\text{T}$ is sufficient to match that of copper (168 neV), and $\sim 6\text{T}$ sufficient for matching even the highest Fermi potential of Ni^{58} (335 neV). This then was the motivating factor for including various pre-polarizing magnets, of peak field strengths 6T and 7T. By doing so, all neutrons of the high energy state are then kinematically forbidden from passing through the field. Furthermore, as the neutrons must first pass a 30° bend in their guide system, it is

ensured that all higher energy neutrons will have exited the system. Therefore, UCNA is able to obtain initially 100% polarization of incoming UCN.

Next, we can observe the classical gravitational potential energy U of a mass m in a given gravitational acceleration g , $U = mgh$. For the neutron on Earth's surface, one finds that the energy corresponds to roughly $100 \text{ neV}\cdot\text{m}^{-1}$. By coincidence, this means that the energy range of the gravitational interaction on the scale of a few meters, the magnetic interaction on the scale of a few T, and the fermi potential of common materials are all the same order of magnitude.

It may be considered whether or not the quantum effects of the gravitational interaction would cause the neutron to behave in a non-classical manner. This will be addressed in Chapter 5, but it is worth noting that the energy scale for the ground state in earth's gravity is of order 1 peV. This means that UCN energies are 3 orders of magnitude above this regime, per cm, and thus that the quantum mechanical behavior at this height scale can be neglected; in this way, the neutrons evolve with classical, parabolic trajectories during their transport. Of course, their interactions with walls necessitate by their very nature the inclusion of quantum mechanical interactions, and the essence of existing as a 2-state, spin $\frac{1}{2}$ system is wholly based in quantum mechanics as well; however, these effects can be regarded separately, allowing a mixed treatment of the neutron as simultaneously a classical and quantum object.




Figure 2.5 Rendering of current experimental setup in “Area B” at Los Alamos National Laboratory. UCNA experiment consists of the purple and green components, along with the upstream guides connecting the UCN source. A UCN nEDM experiment (left) and the UCN lifetime experiment (right) are also depicted.

2.3 Polarization

For the asymmetry to be constructed, we require careful control of the polarization state of the loaded neutron populations. To do so, an adiabatic fast passage (AFP) spin flipper is implemented [2]. Although all other guide components are composed of copper, and possibly coated with higher Fermi potential materials, these traditional materials would instead shield the radio frequencies required to induce a spin flip transition. As such, a special quartz guide, coated in a thin layer of diamond-like carbon (DLC) was developed [1], as an alternative to the use of the toxic beryllium typically used [62]. Use of these guides is generally not preferable, as they suffer multiple drawbacks; the cost and difficulty of machining is increased, and the risk of

breakage/cracking was significant. Additionally, as this component was required to span the inside of the magnet, it was required to be composed of a single, 1-meter long piece. This increased further the risk of damage to the guide during any installations or modifications to the physical system. After the quartz guide was placed, the AFP spin flipper was fit snugly around the outside, centered at the 1T region of the magnetic field.

For calibration of the spin flipper, a series of polarizing fields was constructed. Instead of the eventual path to the decay trap, the neutrons were instead directed past another high peak B field, of order 6T, beyond which was placed a neutron detector. By slowly sweeping the driving frequency of the resonator, we were able to observe a minimum in transmission, corresponding to the maximum spin flip efficiency [63]. Although driving power was also a method of controlling this efficiency, physical limitations from the ability to cool the spin flipper sufficiently imposed constraints as to the maximum deliverable energy. Also, careful analysis of the AFP transition efficiency reveals that as the rotating field increases in magnitude, becoming less and less perturbative, the transition becomes less likely. However, power scans of the response of the neutrons displayed no such behavior up to the maximum input power of 60W, nor significant increase beyond 40W, which is where the system was chosen to be maintained. For further details on the effect and use of driving power, see section 3.4.5.

2.4 β -decay Measurements

The details of the UCNA β asymmetry measurement are exhaustively treated in [18]. We will briefly address a few key details here, for completeness. To measure the angular asymmetry A requires, at minimum, either 2 detectors or 2 polarization states. The asymmetry can then be expressed, in either case, as

$$A_{meas}(E_e) = \frac{r_1(E_e) - r_2(E_e)}{r_1(E_e) + r_2(E_e)} = P_n A \beta \langle \cos\theta \rangle, \quad (2.2)$$

With $r_{1,2}$ representing the count rates in either each detector, or each polarization state, P_n the polarization of the neutron, and β the electron velocity in units of c .

However, UCNA utilizes 2 detectors along with 2 polarization states, allowing utilization of a super-ratio technique. For detectors 1/2, in polarization loading states $-/+$, one may instead construct a ratio of ratios,

$$R = \frac{r_1^-(E_e)r_2^+(E_e)}{r_1^+(E_e)r_2^-(E_e)} \quad (2.3)$$

$$A_{meas}(E_e) = \frac{1 - \sqrt{R}}{1 + \sqrt{R}}$$

In determining the asymmetry using this method, both transport and detector efficiencies cancel to first order, removing what would otherwise be the dominant contributions to the uncertainty. A critical assumption made here is that there are no polarization-dependent contributions to the efficiency. Although the symmetry of our geometry suppresses some sources of systematic error, special attention must be paid to any systematic effects coupled with polarization.

The raw asymmetry was measured by utilizing a series of detectors and vetoes to determine the counts in each of 2 directions; emission along the magnetic field direction, or emission opposite it. Due to the orientation of the lab and magnet setup, these 2 directions will be referred to as the East and West detector setups.




Figure 2.6 Diagram of transport system and decay trap. From [35]

The electrons emitted in the beta decay of the neutrons located within the decay trap are confined by the ambient magnetic holding field of 1T. At this field strength, electrons emitted even at a 90° angle to the detectors, and with the maximum 782 keV endpoint of energy, are trapped within a maximum radius of 3.88mm for their Larmor rotation. This results in an effectively full 4π steradians of coverage for the detector system. Each end of the decay trap was capped with thin mylar foils, coated with between 150-300nm of Be, to increase the storage time, and thus density, of UCN within the trap. Additionally, the foils prevented the neutrons from sampling a field expansion region, in which the magnetic field drops to $\sim 0.6T$. This region serves 2 purposes; first, that electrons incident towards the detectors will become more forward directed, reducing the chance of backscatter events. For example, this field drop would cause a pitch angle of 90° in 1T to be lowered to 51° in the expansion region. 2nd, if backscatters do occur, electrons must traverse a magnetic mirror moving from the 0.6T to 1T region. Adiabatic invariance of the magnetic moment implies the Larmor radius shrinks, and symmetrically the

pitch angle increases as the electron approaches the 1T region. Therefore, any electrons emitted with pitch angles larger than the critical 51° value, will be totally reflected by the magnetic field, and will prevent the electron from being detected in the opposite detector as a false event.

The detector system consists first of a Multiwire Proportional Chamber (MWPC), composed of 64 anode and 128 cathode wires, with each cathode plane strung perpendicularly to one another. This configuration creates a grid pattern, with separations of 2.54mm. The anode wires were constructed with 10- μm diameter gold-plated tungsten, while the cathodes consisted of 50- μm diameter gold-plated aluminum. With this grid, it is then possible to reconstruct the position of the particle. This then allows for position dependence to be calculated and corrected for the secondary portion of the detector; the scintillator. Additionally, the MWPC is highly insensitive to gamma rays, which drastically reduces background. Although the face of the MWPC spanned $16.3 \times 16.3 \text{ cm}^2$, the magnetic field difference corresponds instead to only $\sqrt{0.6}$ of that area. However, this maps to $12.6 \times 12.6 \text{ cm}^2$, offering full coverage of the 12.4-cm diameter decay trap guide.

The MWPC also requires a “fill gas” to be used for generating sufficient signal voltage gain. For this case, neopentane was chosen. Because of the requirement of a thin foil for neutron containment, and to separate this fill gas, it was important to minimize the pressure of this fill gas within the MWPC system, to decrease the pressure differential supported by the foil. A pressure of 100 Torr was settled on, which corresponded to an allowable biasing voltage of 2700V. To increase the amount of support for the foils, fibers of Kevlar were integrated as a support structure.

As alluded to in the previous paragraph, following the MWPC is a plastic scintillator, attached to 4 light guides, which directed the scintillation light to corresponding photomultiplier

tubes (PMTs) over a meter away, in a much smaller 0.03T field. The scintillator was the 3.5mm thick EJ-204. For the case of a maximally energetic electron, a penetration depth of 3.1mm was calculated, and thus the scintillator was able to capture the full energy spectrum from the decay. The PMT's were also magnetically shielded, using a combination of passive μ -metal coverage, and a solenoid as an active component.




Figure 2.7 Electron detector package diagram. Identical packages were located on either side of the superconducting solenoidal decay trap. Electrons (traveling from right to left) first transit the decay trap windows (which reflect neutrons), through the Multi-Wire Proportional Counters, and end in the plastic scintillator. Reproduced from [18]

For an event to be counted, timing coincidence between the MWPC and scintillator is required, suppressing external background signals from contributing to the measurement. Around both detectors were sealed drift tube counters, 1 with a plastic scintillator muon veto around one “arm” of the spectrometer, removing any events associated with cosmic muons.

To eliminate the possibility of linear drifts of either backgrounds or detector efficiencies, an “octet” structure for β decay runs was used. The 2 possible octets are based on the state of the spin flipper during the loading. The variation of loading states were $-+++---+$, or its inverse/reverse configuration. Which of the two octet structures chosen was randomized by technical means, such as coin flips.

The calibration of the detectors was performed using multiple distinct methods, to properly characterize position-, energy-, and backscattering- dependent effects. For determining the position dependence of the scintillator detector, a novel method was employed, in which ^{135}Xe gas was made, via exposure to neutron capture, and then allowed to diffuse into the decay trap. With a uniform population of β -emitters, the position dependence of the light transport efficiency can be assumed constant across the entire face. By utilizing the position reconstruction of the MWPC's, a position dependent normalization function can be formed for the scintillator.




Figure 2.8 Position dependent energy calibration maps for each of the 4 quadrants in the PMT array (East package). Normalized via neutron activated Xe gas isotropically filling the system.

After position normalization had taken place, energy reconstruction calibration was required. We utilized a sealed source insertion system, which allowed physical placement of

known quantities of radioactive material into the center of the decay trap. 3 sources were typically included on the insertion paddle: ^{137}Ce , ^{113}Sn , and ^{207}Bi . Conversion electrons emitted by each species were of known energies, respectively 130.3 keV, 363.8 keV, and with 2 peaks for Bi; 481.7 keV and 975.7 keV.




Figure 2.9 Example energy distributions during sealed source calibration runs. From [35]




Figure 2.10 Example of position detection signal during calibration run. Reproduced from [35]

During analysis, a blinding factor was also applied to the data, preventing our collaboration from imposing biases, consciously or otherwise. Factors of up to 5% were applied

to modify the measurements and were only removed after all analysis had been completed. The β asymmetry results covered here were published in [35], finding

$$A_0 = -0.12015(34)_{stat}(63)_{syst}$$

$$\lambda \equiv \frac{g_A}{g_V} = -1.2772(20) \quad (2.4)$$




Figure 2.11 Example of energy uncertainty envelopes produced via sealed source calibration runs. Reproduced from [35].

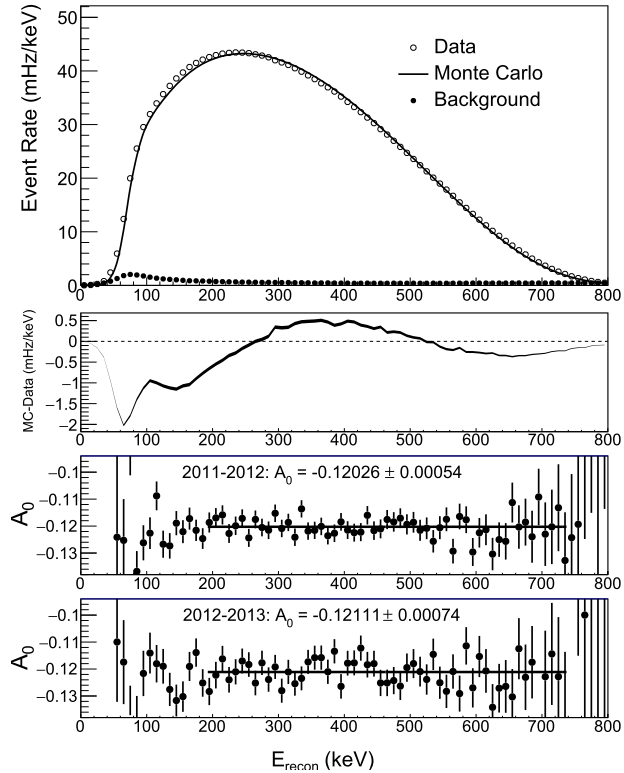


Figure 2.12 Energy reconstruction and fit value to A_0 parameter. Reproduced from [35]

2.5 Backscattering

One crucial systematic associated with determination of the raw asymmetry value is correctly identifying the various forms of backscattering that are possible. By demanding coincidence between the scintillator and MWPC detectors, along with use of timing cuts and position reconstruction afforded by the MWPC's, the majority of events can be reconstructed and assigned between various categories of backscattering types. The primary types of consideration are (See **Figure 2.13**)

- 1) Type 0: No backscattering. Electrons are traverse and are detected successively between the MWPC and scintillator, in the detector correctly associated with their emitted direction. Ideally, all events would be of this form.

- 2) Type 1: Events in which the electron is fully detected by both the MWPC and scintillator of one detector, and is then scattered towards the opposite detector, where it is then also detected by the MWPC and scintillator detector. These are easily identifiable by their characteristic correlation in spatial reconstruction, as well as the restriction on the time of flight required to transit the 4m long guide section between the detectors
- 3) Type 2: Events in which the electron is initially only detected by the MWPC of a given detector, and subsequently detected by the opposite detectors MWPC and scintillator. This form would be misallocated if determined by only twofold scintillator triggering.
- 4) Type 3: The time opposite of type 2, in which electrons are fully detected by the MWPC/scintillator package of the first, correct detector, and then detected only by the MWPC of the 2nd. Type 3 cannot be distinguished from type 2 with scintillator triggering alone, and their symmetry lends to consideration of type 2/3 events being considered jointly in most scenarios.
- 5) Missed events: Events in which the electron scatters backwards without being detected by the correct detector, and is instead detected by the opposite detector package. These may occur on either the decay trap foils, or the fill gas present within the MWPC unit. These events cannot be identified experimentally, and will simply result in an incorrect determination of the asymmetry, relying on simulation reconstruction for systematic correction
- 6) Lost events: Events not detected at all, or only producing a signal in a single MWPC. Since coincidence is required with the scintillator PMT signal, these events will

appear as background, or not at all, and are thus excluded from consideration. These still require correction via Monte Carlo simulation of the experimental sensitivity.

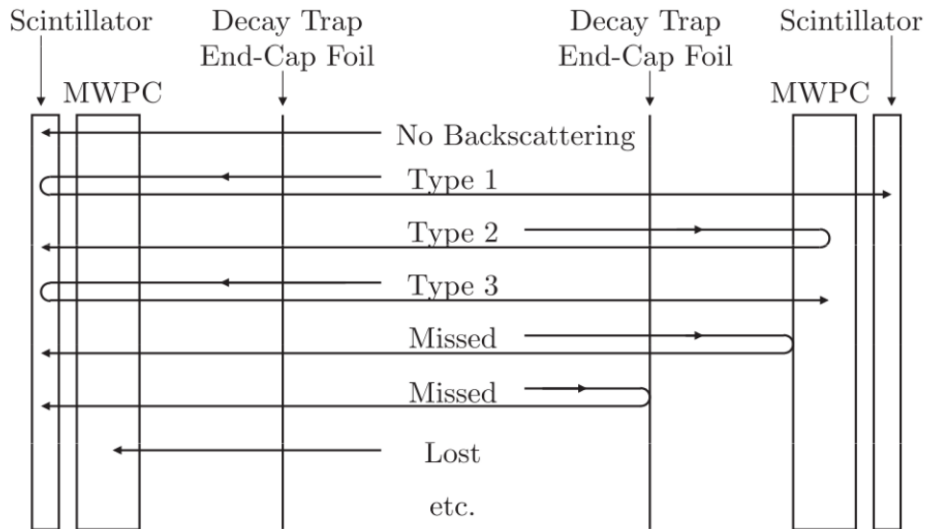


Figure 2.13 Depiction of various backscattering event types. Reproduced from [35]




Figure 2.14 Example of PMT signal comparison with and without coincidence requirement for MWPC detection. [35]

Chapter 3

Polarization Analysis

3.1 Depolarized Fraction

We now turn to the average polarization referenced in eq. (1.12). This polarization term represents the average fraction of neutrons polarized along a given axis, and can be expressed as

$$\langle P \rangle = \frac{N_p - N_d}{N_p + N_d} \quad (3.1)$$

With N_+ representing the number of neutrons polarized along the axis of the magnetic field, and N_- the oppositely polarized. As the polarization required by the experiment is high, our method for determining the polarization is best accomplished instead through measurement of the depolarized fraction, which is defined by

$$\xi = \frac{N_d}{N_d + N_p} = \frac{1 - \langle P \rangle}{2} \quad (3.2)$$

To determine A with sufficient precision for improving current limits, a proportionally stringent range for equilibrium polarizations in the decay trap must be found. One of the primary advantages of the UCNA experiment is the ability to measure this polarization “In Situ”, which is to say with no time delay between measuring polarization after the decay rate observation times, as well as without requiring any alterations or disturbances to the system before initiating the measurement phase. This primary form of run is referred to as the standard depolarization run cycle, or “D” runs. Likewise, in concert with the formation of the super-ratio allowed through alternating the polarized spin state loaded into the decay trap in complementary runs, we reference each run type by their associated spin state. The initially polarized neutrons are in the low energy, high-field seeking state, which is referred to as the “minus” state (As the neutron has

a negative magnetic moment, the up- and down- nomenclature typically adopted for 2-state systems is highly misleading and is abandoned in favor of clarity). Likewise, adding energy is required to produce a spin flip to the opposite state, accordingly referenced as the “plus” state, or the low-field seeking state. This plus- and minus- nomenclature will always be used to designate the state of the spin flipper during the loading process, to maintain consistency, but it can become confusing on first pass when implementing complementary normalizations; I will do my best to clarify when this distinction may be misleading, and to clarify the populations present at each stage of various run types.

Although the primary run types for measurement of the polarization are performed in situ, multiple other run types are developed and implemented ex situ, for measurement of a variety of other parameters. These measurements allow assessment of individual guide components, as well as in evaluating neutron behavior in non-standard loading scenarios, for characterizing and calibrating our system. Only one such measurement, the Polarization, or P-type run, is required for direct calculation of the depolarized fraction, and care was taken to perform these in close time proximity to standard beta decay measurements as well, for the data most accurately reflecting concurrent conditions of the guide system. However, the statistics required for these types of runs are substantially lower than that of their D-type counterparts, and thus few of them were required.

3.2 Spin Flipper

Controlling the spin state of neutrons loaded into the beta spectrometer is crucial to allowing implementation of a super-ratio, drastically increasing the precision to which A can be calculated. While one of the spin states is relatively easy to select for (the low energy, $|-\rangle$ state)

through kinematic polarization, a transition must be induced to populate the higher energy state. This can be accomplished multiple ways, such as through super reflectors [64]. The first, most obvious choice would be to utilize Rabi spin flipping. To do so, one need only apply a small, rotating field, perpendicular to the direction of polarization, for an appropriate time such that $\theta = \omega_1 \Delta t = \pi$; that is, that the angle rotated through during the Larmor precession of the neutron induces a full π radian reorientation.

This can most easily be understood by transforming into the rotating frame of the neutron. In this frame, the neutron is now stationary, and as such it must observe no magnetic field about which it would normally precess. If a perpendicular field is then applied in a single direction, the neutron would begin instead to precess about this new field. So, to create this field in the rotating neutron's frame, one need apply a perpendicular field that also rotates at the resonant frequency of the neutron in the larger, holding field.

This works well for pulsed operation, in which the delay between pulses is sufficiently long as to not be loaded with the population from a separate pulse, and where the population's residency time in the spin flipping region is large compared to the required time to perform the spin flip. If these cannot be accomplished, instead of pulsing the rotating field for some Δt , one may instead have the rotating field operated continuously, over some distance such that $d = v\Delta t$ to produce the correct timing. As may be immediately obvious however, this then requires high monochromaticity of the impinging beam.

To account for the continuous flux of neutrons, as well as the wide spectrum of velocities in UCNA, the technique utilized instead is that of Adiabatic Fast Passage (AFP) spin flipping. To exploit this method, we start with a similar premise to that of the Rabi spin flip; first, we consider a B field with a strong holding field B_0 , and a weaker, rotating field B_1 ,

$$\mathbf{B}_I = B_0(z)\hat{\mathbf{z}}_I + B_1(z)[\cos(\omega t)\hat{\mathbf{x}}_I + \sin(\omega t)\hat{\mathbf{y}}_I] \quad (3.3)$$

Then, transforming to the rotating coordinate system, the magnitude of the z-component of the field is altered by the frequency of the rotational transformation we perform, and the B_1 field remains unchanged in magnitude, but now appears in a constant direction.

$$\mathbf{B}_{II} = \left(B_0(z) + \frac{\omega}{\gamma}\right)\hat{\mathbf{z}}_{II} + B_1(z)\hat{\mathbf{x}}_{II} \quad (3.4)$$




Figure 3.1 Diagram for the initial and rotated coordinate systems

With $\gamma = -1.832 \times 10^8 \text{ T}^{-1}\text{s}^{-1}$ as the gyromagnetic ratio of the neutron. The angle and magnitude of this new B field can be found through standard trigonometry, yielding

$$B_{II}(t) = \sqrt{\left(B_0(t) + \frac{\omega}{\gamma}\right)^2 + B_1^2(t)} \quad (3.5)$$

$$\tan \theta = \frac{B_1(z)}{B_0(z) + \frac{\omega}{\gamma}} \quad (3.6)$$

We would then like to determine the dependence of the rate of change of this angle with respect to the field magnitude and driving frequency. After some application of the derivative chain rule, one finds

$$\dot{\theta} = \gamma \frac{-\gamma B_1 \dot{B}_0 + \dot{B}_1 (\omega + \gamma B_0)}{(\omega + \gamma B_0)^2 + (\gamma B_1)^2} \quad (3.7)$$

Next, it is useful to define the “adiabatic parameter” α as the ratio between the rate of rotations around the effective B_{II} field as seen from the rotating field’s frame vs. the rate of change of the inclination,

$$\alpha = \frac{|\gamma| B_{II}}{|\dot{\theta}|} \quad (3.8)$$

In doing so we can quantify the adiabatic condition: as α becomes large, the system is more adiabatic, and a spin flip more likely. This is in effect stating that the adiabatic condition reflects a large number of precessions about the holding field compared to the rate of change of the holding field’s angle, allowing an equilibrium precession to be established for every infinitesimal change to the holding field.

To determine the efficiency of the spin flipper, we need analyze the quantum mechanical probability of a spin flip occurring. We begin with the Schrödinger equation for a particle in a magnetic field

$$i\hbar \frac{d\psi}{dt} = \mathcal{H}\psi = \boldsymbol{\mu} \cdot \mathbf{B}\psi = \frac{\gamma\hbar}{2} \boldsymbol{\sigma} \cdot \mathbf{B}\psi \quad (3.9)$$

With $\boldsymbol{\sigma}$ the Pauli matrices, and $\psi = a_1(t)|+\rangle + a_2(t)|-\rangle$, the normalized 2 state solution.

Utilizing spherical coordinates for \mathbf{B} , and changing to spinor representation with column vectors, we can simplify Eq. (3.9) further as

$$i \begin{bmatrix} \dot{a}_1 \\ \dot{a}_2 \end{bmatrix} = \frac{\gamma}{2} B \begin{bmatrix} \cos \theta & e^{-i\phi} \sin \theta \\ e^{i\phi} \sin \theta & -\cos \theta \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} \quad (3.10)$$

Next, we transform into the frame of the field. To do so requires rotating the wave function, taking care to include the behavior of spinors under rotation; namely, that they instead transform

with $\theta/2$, to satisfy their antisymmetry under 2π rotations. We will also define $\omega = \gamma B$ for notational convenience. So, we let $\mathbf{a} = \mathbf{R}\mathbf{b}$, with \mathbf{R} the rotation matrix for spinors,

$$\mathbf{R} = \begin{bmatrix} \cos \frac{\theta}{2} & -e^{-i\phi} \sin \frac{\theta}{2} \\ e^{i\phi} \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{bmatrix} \quad (3.11)$$

$$\dot{\mathbf{a}} = \frac{d(\mathbf{R}\mathbf{b})}{dt} = \dot{\mathbf{R}}\mathbf{b} + \mathbf{R}\dot{\mathbf{b}} = -i\omega \begin{bmatrix} \cos \theta & e^{-i\phi} \sin \theta \\ e^{i\phi} \sin \theta & -\cos \theta \end{bmatrix} \mathbf{R}\mathbf{b} \quad (3.12)$$

We then put the \mathbf{b} terms on the right, and apply the \mathbf{R}^{-1} operation to each side, solving for $\dot{\mathbf{b}}$, and complete the application of the rotation matrices through brute force, yielding

$$\begin{bmatrix} \dot{b}_1 \\ \dot{b}_2 \end{bmatrix} = \begin{bmatrix} -i\omega & \dot{\theta} e^{-i\phi} \\ -\dot{\theta} e^{i\phi} & i\omega \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \end{bmatrix} \quad (3.13)$$

This shows that the change in θ is what allows mixing between the 2 states, as one would expect a static field to leave the states unchanged as they precess.

By changing from derivatives in time, to derivatives in θ instead, and allowing $\phi=0$ for simplicity, the equations can be decoupled into two 2nd order differential equations, which produce the transition probability of [65], [66]

$$\epsilon = 1 - \frac{1}{1 + \alpha^2} \sin^2 \left(\frac{\pi}{2} \sqrt{1 + \alpha^2} \right) \geq 1 - \frac{1}{1 + \alpha^2} \quad (3.14)$$

Finally, we include some simplifying assumptions. First, we will take the worst-case scenario as the moment the driving B_1 field is on resonance, creating the lowest α value for the whole of the process. Generalizing then to develop a lower bound, we take the system to behave at this worst case α for the entirety of the process. Additionally, we include a very small gradient in the B_0 field, and with slow particles, allow us to express the rate of change in the B_0 instead as a

function of distance transited in the field, $\dot{B}_0 = (\nabla_z B_0)v$. Together, we arrive at a lower bound on the transition efficiency of our spin flipper,

$$\begin{aligned}
 & \text{On resonance; } \left(\frac{\omega}{\gamma} + B_0\right) \rightarrow 0 \\
 & B_{II} \rightarrow \sqrt{0^2 + B_1^2} = B_1 \\
 & \dot{\theta} \rightarrow \frac{-\gamma^2 B_1 \dot{B}_0 + \dot{B}_1 \cdot 0}{0^2 + \gamma^2 B_1^2} = -\frac{\dot{B}_0}{B_1}
 \end{aligned} \tag{3.15}$$

And, after putting this all together, we obtain our expression for the lower limit to the performance of our AFP spin flipper, dependent on field gradient, velocity spectra, and applied rotational field [63];

$$\epsilon > 1 - \frac{(\nabla_z B_0)^2 v^2}{(\nabla_z B_0)^2 v^2 + \gamma^2 B_1^4} \tag{3.16}$$

Here we see the tradeoffs. We desire a shallow gradient in B_0 to transit the field change more gently, and desire higher B_1 fields to minimize the inefficiency. Yet, the whole method hinges on the very assumption that B_1 be small compared to B_0 , as well as sufficiently uniform in space. With nominal values of 0.1G/cm field gradient, and B_1 strength roughly 1G, one obtains an estimate for the highest speed (8m/s) tail on the order of $1 - \epsilon \cong 10^{-5}$.

Representative efficiencies can be seen in other experiments as a benchmark, with the spin flipper implemented by the PERKEO II spectrometer for cold neutrons as $\epsilon = 99.99(1)\%$ [67]. Performances similar to this are predicted from simulation parameter scans, which will be discussed later in Chapter 4.

With the theory of how to control the spin states of the neutrons established, we turn to the method of achieving it. It happens that producing circularly rotating RF fields is mechanically more difficult than simple transverse fields, and as such, a compromise is reached. First, we note that having instead 2 counter-rotating B fields is functionally equivalent to producing a single oscillating transverse magnetic field,

$$\mathbf{B} = B_t \sin(\omega t) \hat{\mathbf{x}} = \frac{B_t}{2} [\{\sin(\omega t) \hat{\mathbf{x}} + \cos(\omega t) \hat{\mathbf{y}}\} - \{\sin(-\omega t) \hat{\mathbf{x}} + \cos(-\omega t) \hat{\mathbf{y}}\}] \quad (3.17)$$

We see that the requisite field profile for our AFP spin flip is satisfied, albeit at 2 separate frequencies: ω and $-\omega$. However, with B_1 sufficiently small, the departure from resonance of the opposite mode is highly suppressed, and very high efficiencies are still achievable. It is much easier to produce this transverse field than the rotating one, and to do so requires only that a sinusoidal current distribution is produced, at the resonant frequency of the neutrons while in the given field strength of 1T.

To produce this current distribution, UCNA utilizes a birdcage resonator. This consists of an even number of rods, joined at 1 end by a solid, conducting ring, and the opposite with each rod joined to its 2 adjacent neighbors through a capacitor each. As the number of rods increases, the field uniformity throughout the resonator is improved, however it comes at the cost of creating more resonant modes, along with technical difficulties in manufacturing. These rods are also made hollow, such that cooled water may be pumped through them during operation, to avoid overheating from the large power input provided.

The resonator may be driven in 2 different modes; end driven, with leads on opposite sides near the capacitive end, or rung driven, with the leads on opposite sides of a single rod. Each mode produces a different set of resonant frequencies.




Figure 3.2 Picture of water cooled Adiabatic Fast Passage Spin Flipper, from [63]

All conductive pieces of the spin flipper were made of copper or brass, and subsequently electroplated in silver, with a thickness of $\sim 25 \mu\text{m}$, equivalent to many skin depths of electrical current of silver in the RF region of operation. This coating serves to reduce the resistance further for the system, since the resistivity of silver is one of the lowest available $\rho \cong 1.6 * 10^{-8} \Omega \cdot \text{m}$. Capacitors were roughly 470pF, with silver microstrip leads for lowered resistance as well, and within 2% tolerances for uniformity of the field produced.

The resonator was tuned by utilizing an Agilent 4195A network analyzer. Resonances were measured using both direct coupling (physical connection) to the resonator, as well as inductive coupling, by means of holding a loop near the resonator, in multiple positions with respect to the rods, and at different depths within the resonator itself.

The resonant frequencies of the birdcage resonator, of (even-numbered) runs N is given by;

$$\omega_j = \frac{1}{\sqrt{C \left[2M + L - 2M \cos \frac{2\pi j}{N} \right]}} \quad (j = 1, \dots, N), \quad (3.18)$$

Though the strongest resonances are found at the extrema of $j=1, N$. Below can be seen examples of the primary resonance, along with associated, weaker modes.



Figure 3.3 Resonance spectra of spin flipper when coupled to network analyzer inductively. These produce much weaker minima of other modes, as the primary is selected preferentially.




Figure 3.4 Resonance signal output from direct coupling of network analyzer to spin flipper, through end-driven mode. Multiple resonances can be observed as minima in impedance, along with discontinuities in phase angle.

Once the device is assembled to produce resonance in the general region of interest, fine tuning must be performed. The sliding ring allows for slight modifications of the device after installation, and can allow for adjustments required from the change to the inductive behavior of the instrument when located in situ, from RF shielding cages, and nearby mutual inductors.

This resonance can further be tested by the transmission of the neutrons themselves, both as a function of driving frequency, and driving power. Utilizing a crossed polarized analyzer, a given spin state can be made to change from transmitting to nontransmitting, or vice versa, and a corresponding minimum(maximum) will demonstrate the ideal tuning for maximum resonant transition probabilities.

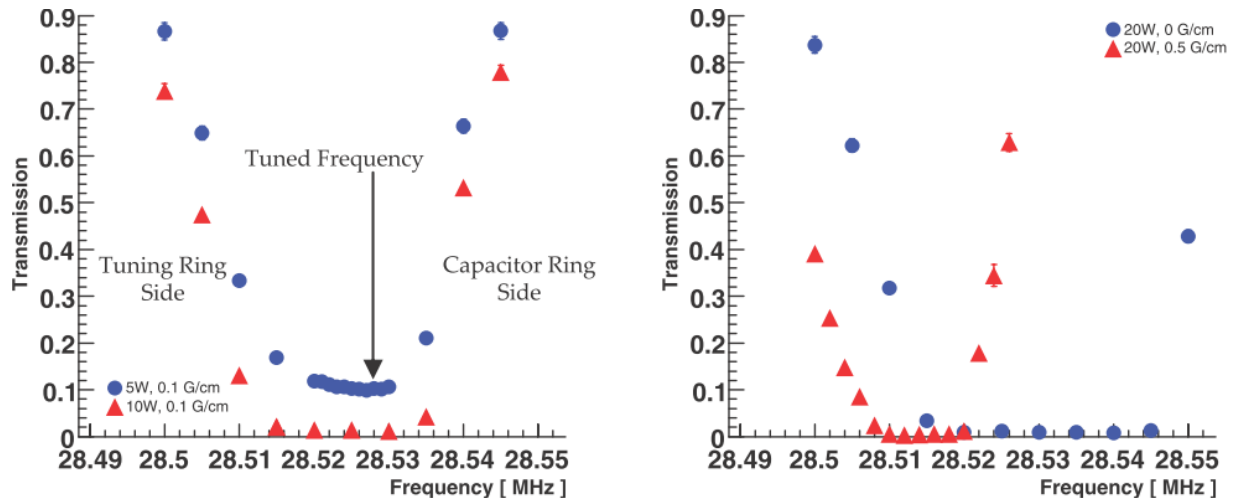


Figure 3.5 Examples of transmission efficiencies during frequency tuning in crossed polarizer analyzer configuration. Reproduced from [63]

Additionally, special care was taken to ensure the uniformity of the holding field gradient in the area of the spin flip. Corresponding to eq (3.16), minimizing the field gradient corresponds to maximizing the efficiency of the flipper. As such, the AFP magnet was designed with 10 additional shim coils, centered around the spin flipper, to tailor the field to field gradients of order 0.1 G/cm, enhancing the efficiency of the adiabatic transport condition greatly (**Figure 3.6**).

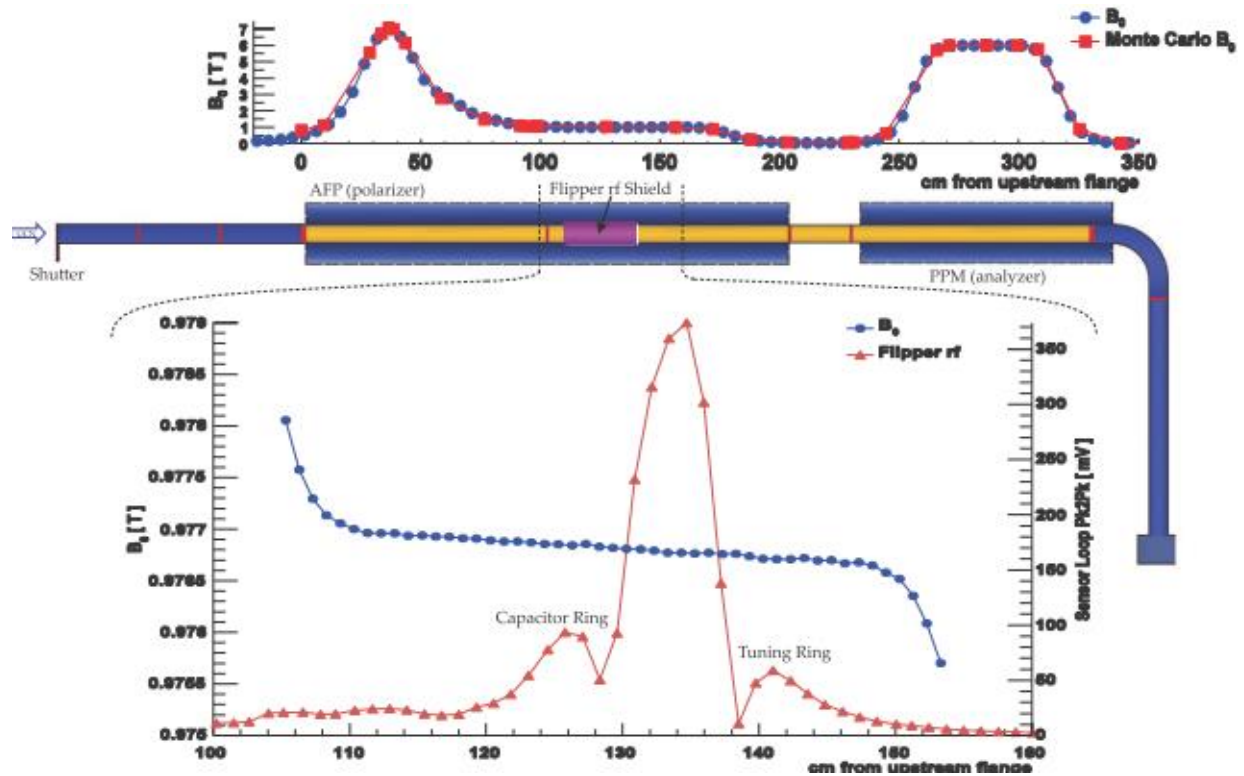


Figure 3.6 (Top) Measured/Simulated values of holding field magnitude along propagation direction. (Bottom) Zoomed in view of holding / oscillating fields at the location of spin flip transition. Reproduced from [63].

3.3 Pre-Shutter Polarimetry

In 2011, a DLC-coated copper plate, functioning as a gate valve at the entrance of the decay trap, was installed. Prior to this shutter, polarimetry proceeded in a somewhat more simplified way. With no shutter available, the only method for determining between the two populations of spin lied in clever utilization of the spin flipper. As such, the typical measurement involved the very direct method of simply changing the switcher state, allowing some time for the guides between the decay trap and switcher to drain, and then flipping the state of the spin flipper to allow the depolarized population inside the decay trap to escape into the switcher.

This primary run type was designated the depolarization trapping (DT) run cycle. This measurement constituted the bulk of statistical data taken. However, a minimum of a 2nd run type was necessary, along with its corresponding systematic, the Reload (R) run type and correction.

This run attempts to account for the fact that there exist neutrons upstream of the spin flipper, but downstream of the peak magnetic field of the AFP, which are able to sample the decay trap as the opposing state when the spin flipper is toggled to the opposite state at the time of unloading.

To address this reloaded population, the reload measurement schema seeks to reproduce this effect. To do so, reload runs proceed much the same as the more standard DT runs, however during the loading phase, the state of the spin flipper is briefly toggled to the opposite state, and returned to its initial state quickly thereafter (of order ~ 5 s). In doing so, the R-type run is assumed to generate this reloaded population twice; once at the initial spin flipper toggle, and again at the usual time of measurement. By inducing a 2nd reload process, one may then find the difference between the 2 measurements and designate this difference as the necessary reload correction to be applied to the DT data sets.

$$\begin{aligned}
 S_{DT} &= \textit{depolarized} + \textit{polarized} \\
 S_R &= \textit{depolarized} + 2 \times \textit{polarized} \\
 S_{\textit{depol}} &= 2S_{DT} - S_R
 \end{aligned}
 \tag{3.19}$$

However, in seeking to reproduce these results, current simulations indicate the method largely underestimates the size of the population introduced via the reload spin flipper cycling, appearing to suggest that those populations reloaded earlier in time behave differently than those at the standard unloading phase.

The leading systematic to this method was, as may be predicted, the very large backgrounds that remained in the system after the guide draining time. This was the direct impetus for construction of the decay trap shutter; by allowing isolation of the decay trap population, the guides could be completely cleaned before proceeding with measurement of the depolarized population.

It should be noted, however, that the shutter installation was not directly an improvement towards the **net** polarization; in fact, it may be the case that polarization was made worse by its simple inclusion into the system, and further that usage of it in determining the polarization leads to additional systematic effects not present in the original run types. Instead, however, its main advantage is solely from its reduction of backgrounds, which did demonstrably improve the **uncertainty** in the measurement of the polarization. With the system sufficiently polarized, it was much more important to know precisely what the polarization was, even at the cost of making it worse.

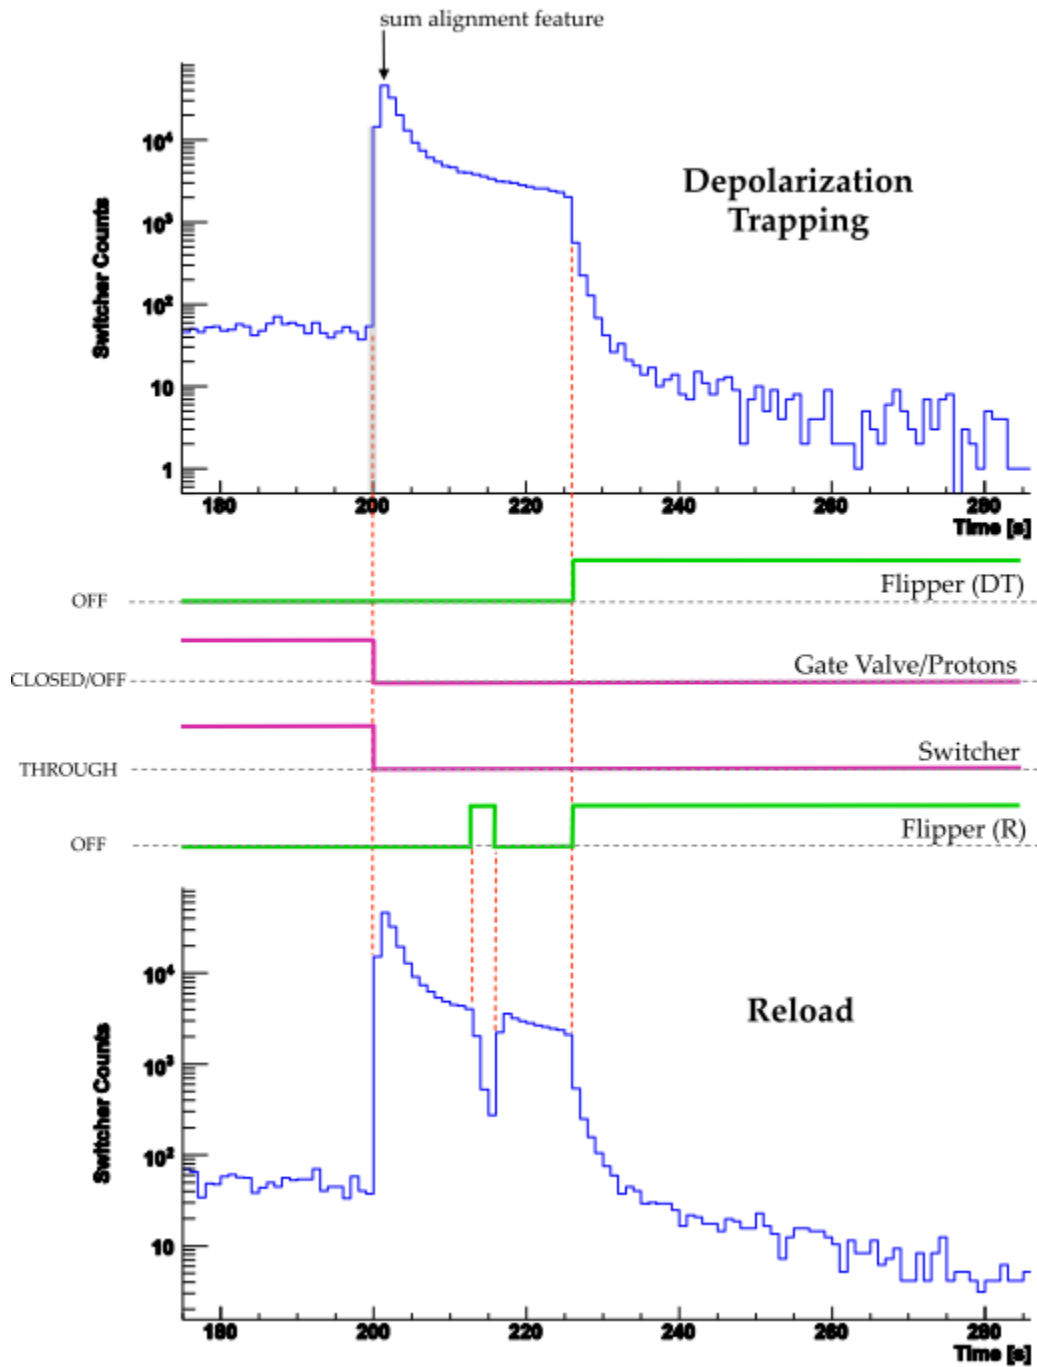


Figure 3.7 Example plots of switcher detector signals during complementary DT and R run types for determining polarization prior to decay trap shutter installation.

3.4 Run Types

3.4.1 Depolarization (D)

The primary method of evaluation is the Depolarization run type. The timing cycle of a D run is as follows:

0s	Loading Phase	
40s	Shutter/Gate Valve/Switcher close, protons off	Right-Spin Guide Clean
65s	Spin Flipper Changes State	Wrong-spin Guide Clean
70s	Shutter Opens	Wrong-Spin Unload
170s	Spin Flipper Changes State	Right-Spin Unload
400s	End of Cycle	

See Figure 3.9 for visualized timing and typical signal profile.

The D runs are typically performed in-situ, at the end of a standard beta decay octet. During the octet, the system is simply loading UCN into the decay trap, and develops an equilibrium population. The start of each D measurement is then chosen such that $t=40s$ corresponds to the beginning of the system changes, allowing the 40s following $t=0$ to constitute the measurement of the steady state loading rate in the decay trap. Upon initiation, multiple components states are altered in concert; the decay trap shutter is closed, thus trapping and separating the population located inside the decay trap. Simultaneously, the gate valve separating the UCN source from the rest of the guide system is closed, and the proton beam is diverted from the experiment; these both serve to reduce the backgrounds during measurement, as well as prevent further loading of UCN into the system. Finally, the switcher mechanism is flipped to the closed, or “down” state. In doing so, the system upstream of the switcher is no longer connected such as to allow direct flow of neutrons downstream into the decay trap, and instead the switcher is moved to connect the downstream components with a (He-3?) UCN detector. The detector is located ~1m lower than the guide system, accelerating any incoming neutrons into the detector,

increasing detection efficiency by allowing higher penetration depths of the foil separating the detector gas from the vacuum of the guides.

Next, the system remains in this configuration for 25 seconds. This means that the components located between the switcher and the decay trap shutter (the population trapped in the guides) is now connected to a detector at one end. This serves as a highly efficient absorber, especially due to the height drop to the detector, and as such, removes the population that has the same polarization as the loaded state. This then is referred to as the “Right-Spin Guide Cleaning” phase. Following this, the spin flipper state is reversed, permitting the opposite spin state to escape from between the high-field region and the shutter, into the decay trap. This is correspondingly named the “Wrong-Spin Guide Cleaning” phase. As this population is small to begin with, and has also been decaying during the previous 25s of cleaning, only 5s of cleaning is required for this stage. At this point, the guide system has all UCN removed, leaving only those neutrons downstream of the shutter, i.e. in the decay trap.

After the guide cleaning phase is over, the shutter is opened, allowing neutrons in the decay trap to flood back upstream. Now, however, the spin flipper remains in the OPPOSITE state from that in which loading occurred; so, although both spin states populate upstream, reaching the high-field region of the AFP magnet, only the wrong-spin states will be permitted through the magnetic field, into the switcher monitor. This phase is referred to as the “Wrong-Spin Unload” phase, while the right-spin neutrons are “Spin Flipper Trapped”. After being allowed to proceed for 100s, the final step is to change the spin flipper back to its original state, thus permitting the right-spin population that was heretofore trapped downstream of the magnetic field to drain into the switcher monitor.

It is worth exploring the 2 possibilities in some detail, as this practice of spin flipper trapping is common across many run types, and is an essential component for the effectiveness of the depolarization measurement scheme. Suppose first that the system is loaded with the spin flipper turned off; this means that the high-field seeking states (a.k.a. “off-loaded”, “low-energy”, or “down” states) are the only neutrons permitted to transit the peak AFP magnetic field of 7T, and remain in this polarization state all the way to the decay trap. After guide cleaning, the spin flipper state would then be energized, changing to “on”. Now, as the high-field seeking neutron state that was loaded passes through the spin flipper region, they are excited to the low-field seeking state (“on-loaded”, “high-energy”, “up” states). This state then continues upstream, encountering the 7T peak field; but as it is low-field seeking, it is now slowed to a halt by the kinematic barrier, and forced to turn around. As it passes the spin flipper once again, it is returned to its original state as it returns towards the decay trap.

Likewise, assume instead the spin flipper to be on during the loading phase. The decay trap will then be populated by the high energy, low-field seeking state. After the guide cleaning phase, the spin flipper would then be changed to its off state. As such, the neutrons remain in their low field seeking state, even as transiting the spin flipper region, as it causes no transitions while de-energized. Thus, the same scenario of the low-field seeking state encountering the 7T field occurs again, and it is likewise trapped.

In the first scenario, of spin flipper off loading, the spin flipper trapping phase functions by converting the right-spin population into the wrong-spin population, preventing it from traveling further upstream. However, in the second scenario, the spin flipper trapping is accomplished by instead allowing the right-spin population to remain as such, naturally being reflected by the magnetic field. It is this asymmetry that lends to difficulty in nomenclature, and

is the impetus for usage of both the right- and wrong-spin designations, as well as the low- and high-field seeking terminologies in tandem.

Table 3.1 Various count rates during standard D-type run cycles. Includes Gate Valve loading counts (near the source), SCS decay trap counts for correlating transport efficiency, counts in the SCS monitor during the wrong spin draining portion, SCS monitor background rates, and the wrong-spin counts reaching the switcher. Includes both spin states, and both run years.

	2011 off	2011 on	2012 off	2012 on
Gate Valve Loading Counts	107181	105522	112229	123348
SCS Loading Counts	245425	165547	205854	161123
SCS draining counts	138680	70305	100341	59696
SCS background rates (per second)	136.9	126.9	206.4	222.4
Switcher counts	1216	1796	734	1248

In principle, the depolarized fraction can be determined by these D runs alone. The goal then would be to utilize the separated detectors' signals to determine measured populations at a given time, and reconstruct for each the $|+\rangle$ and $|-\rangle$ states their equilibrium population at the start of the run cycle. For the different loading configurations then, we would have

$$\begin{aligned}
 \text{Off - loaded} : & \begin{cases} N_p = \epsilon_p^- \cdot N_{0p}^- \cdot e^{-\frac{t_s}{\tau_s}} e^{-\frac{t_f}{\tau_f}} \\ N_d = \epsilon_d^+ \cdot N_{0d}^+ \cdot e^{-\frac{t_s}{\tau_s}} \end{cases} \\
 \text{On - loaded} : & \begin{cases} N_p = \epsilon_p^+ \cdot N_{0p}^+ \cdot e^{-\frac{t_s}{\tau_s}} e^{-\frac{t_f}{\tau_f}} \\ N_d = \epsilon_d^- \cdot N_{0d}^- \cdot e^{-\frac{t_s}{\tau_s}} \end{cases}
 \end{aligned} \tag{3.20}$$

Where ϵ represents a transport efficiency from the decay trap to the switcher detector, the N_0 terms are initial equilibrium populations inside the decay trap, the t 's are the amount of time a

population remains trapped behind either the shutter or spin flipper respectively, and the τ 's their associated lifetimes during trapping.

In each case, the polarized population must wait for 30s while shutter trapped, with characteristic lifetimes τ_s , and then 100s while flipper trapped with characteristic lifetimes τ_f . The depolarized sample, however, must only endure through the shutter trapping, and is represented by the spin state opposite of that which was loaded. Note that in the off-loaded case, the polarized signal is the $|-\rangle$ spin state, while in the on loaded case the polarized signal is instead the $|+\rangle$. These characteristic lifetimes can be extracted directly from the count rate inside the decay trap, for the polarized population of each set; since the polarizations are very high, signals inside the SCS will be dominated by the primary population. The depolarized population can have its lifetime extracted from the corresponding D-run of opposite loading state. In practice, however, separate runs are used for determining these lifetimes, and will be discussed in a later section.

Finally, we have the ϵ terms, which represent the transport efficiencies of each population, and are in principle time dependent, as well as dependent on the populations production mechanism; that is to say, a given state produced via a depolarization mechanism may behave differently than when produced via direct loading. These transport efficiencies can be approximated, but there is in fact a much better way for dealing with these.

With each population's calculation defined, we may combine them with eq. (3.2) to determine the depolarized fraction in each case:

$$\begin{aligned}\xi^- &= \frac{N_d^+}{N_p^- + N_d^+} \cong \frac{N_d^+}{N_p^-} = \left(\frac{\epsilon_d^+}{\epsilon_p^-}\right) \cdot \left(\frac{N_{0d}^+}{N_{0p}^-}\right) \cdot e^{-t_s\left(\frac{1}{\tau_s^+} - \frac{1}{\tau_s^-}\right)} \cdot e^{-\frac{t_f}{\tau_f^-}} \\ \xi^+ &= \frac{N_d^-}{N_p^+ + N_d^-} \cong \frac{N_d^-}{N_p^+} = \left(\frac{\epsilon_d^-}{\epsilon_p^+}\right) \cdot \left(\frac{N_{0d}^-}{N_{0p}^+}\right) \cdot e^{-t_s\left(\frac{1}{\tau_s^-} - \frac{1}{\tau_s^+}\right)} \cdot e^{-\frac{t_f}{\tau_f^+}}\end{aligned}\tag{3.21}$$

Where we eliminated the depolarized populations in the bottom terms, due to their negligible size compared to the polarized populations.

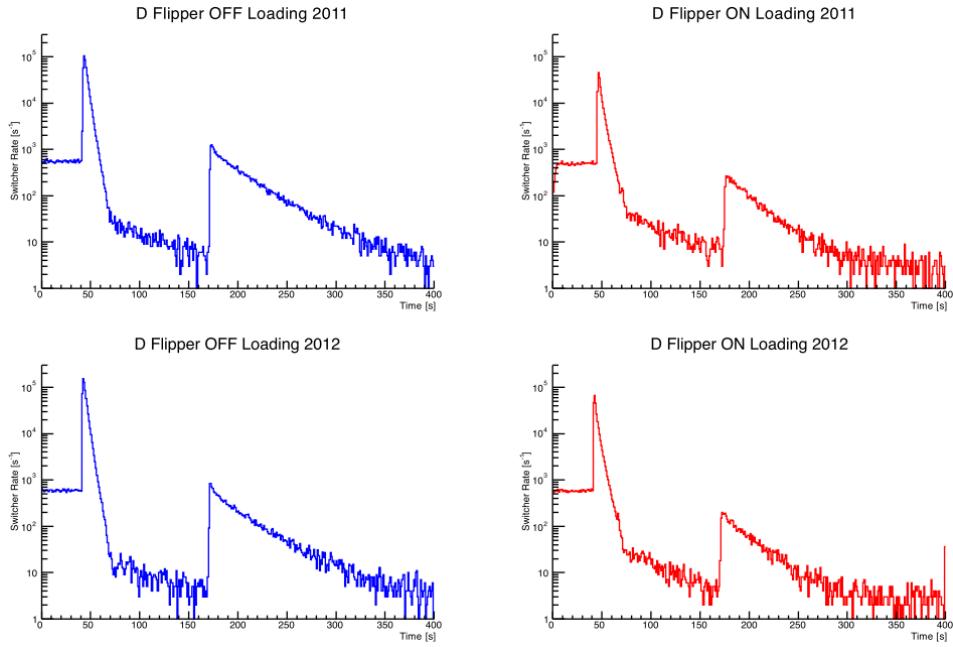


Figure 3.8 Switcher detector signals during standard Depolarization run types, performed at the end of each β octet segment. These are analogous to the DT run types of the pre-shutter polarimetry.

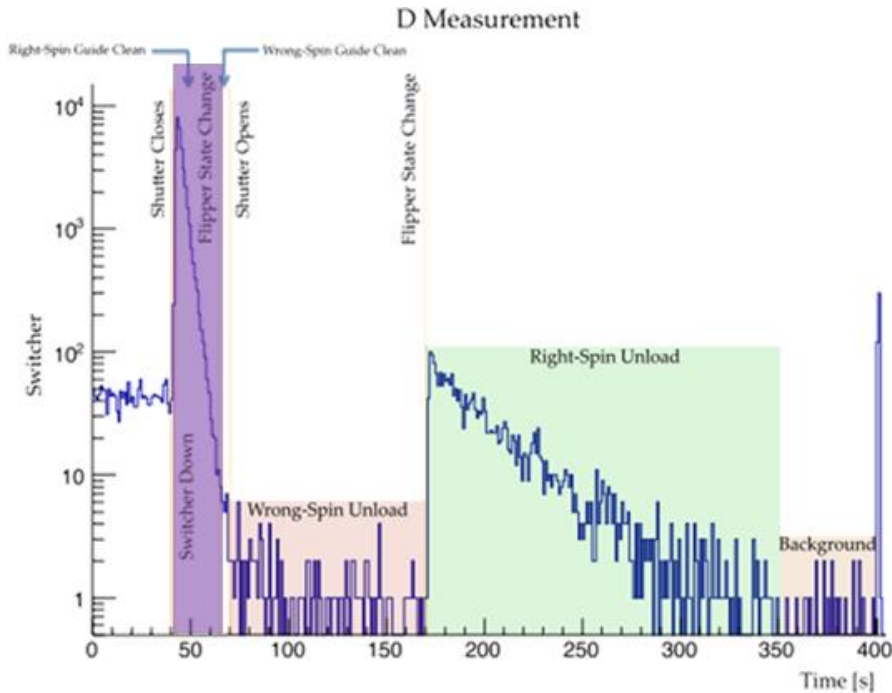


Figure 3.9: An example of a D-type run signal in the switcher detector. The 5 phases are illustrated, though everything preceding the Wrong-Spin Unload is not used in calculations, as it represents a background reduction process.

3.4.2 Shutter and Flipper Trapping

As the shutter trapped time is only of 30 seconds, extended runs better characterize the decay behavior during this period. The D+P method circumvents the requirement for the lifetime of each population while inside the decay trap, and behind the shutter, since the method directly correlates initial populations with resultant switcher detector counts, but it nevertheless serves as a helpful characterization of the decay trap components, and as a secondary method of calculating the depolarized fraction for self-consistency checks.

Two run types used for determining the lifetimes required by the first analysis method are the Shutter Lifetime Drain (SLD) and the Spin Flipper Trapping (SFT) measurements. These runs are among the simplest, and consist only of loading the trap to equilibrium, and either closing the shutter (SLD) or changing the state of the flipper (SFT). By monitoring the SCS-

monitor signal, the decay rate can be extracted for one spin population. We assume the counts to be all from the majority population, which is valid for the high polarization rates within the experiment. This then precludes use with any other run types in which the spin flipper efficiency is altered by any significant amount, as this would then contradict the assumption that only the majority spin state contributes to the signal. Indeed, all determinations of lifetimes require flipper efficiency to be set to the values identical to those used in standard D-type runs, for properly representative decay rates.

Table 3.2 Fits to lifetime for populations in both shutter and flipper trapped scenarios, for each run year and load state.

	2011 off	2011 on	2012 off	2012 on
Shutter trapped lifetime	59.6 ± 0.26	64.8 ± 0.46	44.2 ± 0.26	45.3 ± 0.46
Flipper trapped lifetime	50.9 ± 0.42	36.3 ± 0.39	45.2 ± 0.43	34.8 ± 0.47

3.4.3 Polarization (P)

This run type is an answer to the problem of unknown draining efficiencies for each population inside the decay trap during a D-type run. In examining this method of normalizing out transport efficiencies, we will find additional cancellations that improve statistics dramatically from our previous approach of including each term directly.

First, we will define the timing for a P run as

0s	Loading Phase	
150s	Shutter/Gate Valve/Switcher close, protons off	Right-Spin Guide Clean
180s	Shutter Opens	Right-Spin Unload
400s	End of Cycle	

As these are performed ex-situ (as complementary runs to the standard beta decay measurements), we require a sufficiently long loading time to develop an equilibrium population. Measurements and simulations indicate that 110s, corresponding to roughly 3-4 lifetimes, is sufficient to develop this steady state, and a corresponding 40s is added for measurement of the plateaued loading rate.

Table 3.3 Various counts for P type runs, for each load state and run year. These include the normalizing Gate Valve rate at the beginning of the system(near the source), the SCS decay trap counts for correlating transport efficiency, the SCS monitor counts during the draining period, the backgrounds for the SCS detectors, and the number of counts reaching the switcher detector.

	2011 off	2011 on	2012 off	2012 on
Gate Valve Loading Counts	26079	43041	46253	52301
SCS Loading Counts	31965	37185	49167	41659
SCS draining counts	17064	14124	22735	14615
SCS background rates (per second)	16.6	23.0	23.6	27.1
Switcher counts	49667	33874	60859	30182

The P-type run proceeds much the same as did the D-type runs, however the notable difference is the absence of any change to the state of the spin flipper. We can then write down the populations' evolution over time, as before, however this time we measure only the polarized, or "right-spin" population;

$$\begin{aligned}
 \text{Off - loaded: } N_p^- &= \epsilon_p^- \cdot N_{0p}^- \cdot e^{-\frac{t_s}{\tau_s^-}} \\
 \text{On - loaded: } N_p^+ &= \epsilon_p^+ \cdot N_{0p}^+ \cdot e^{-\frac{t_s}{\tau_s^+}}
 \end{aligned}
 \tag{3.22}$$

Combining these with eq. (3.2), we find that the transport efficiency terms cancel out, as do the flipper trapped lifetimes altogether. Qualitatively, P-type runs answer the question "how

many polarized neutrons were there at the end of the shutter trapped period?”, whereas D-types answered how many depolarized neutrons there were. However, we still must grapple with the shutter trapped lifetimes, as well as the fact we have used an approximation to find the depolarized fractions. To do so, we exploit a key insight; although it’s true that the number of polarized and depolarized neutrons reaching the switcher detector is approximately just the polarized portion, it is also true that the SCS monitor measures the TOTAL population. Not only that, but the SCS is also capable of measuring the neutrons during the equilibrium state before the alteration of the system (i.e., changing of the switcher state), and can measure the shutter trapped lifetime.




Figure 3.10 Switcher detector signals during the Polarization run type. These ex-situ runs are performed independent of standard β -octet run sequences, but due to the higher particle counts during these runs, fewer are required to be completed than for the D-type runs.

With this final fact, we can now construct terms that answer a broader, and more apropos question to our desired measurement; given a certain number of neutrons detected by the SCS monitor, how many neutrons of that population are expected to reach the switcher under given conditions. We may then replace the population decay during the shutter trapped period, to simply be the number actually measured when allowed to release as the right spin population after the shutter holding time.

$$\begin{aligned}\xi_{D+P}^- &= \left(\frac{P_{scs}^+}{P_{swt}^+}\right) \cdot \left(\frac{D_{swt}^-}{D_{scs}^-}\right) \\ \xi_{D+P}^+ &= \left(\frac{P_{scs}^-}{P_{swt}^-}\right) \cdot \left(\frac{D_{swt}^+}{D_{scs}^+}\right)\end{aligned}\tag{3.23}$$

Where all populations are found by counting from the opening of the shutter, such that for D-runs one is tabulating wrong-spin neutrons detected, and in P-runs, the right-spin neutrons instead. In this way, we transform a detected population in the switcher, to the corresponding expected rate in the SCS when a majority state, as in the P runs. For example; for the off-loaded case, the depolarized population is instead the high energy $|+\rangle$ state, and as such requires scaling according to the oppositely loaded flipper-on P run.

This also permits the complete removal of uncertainties associated with the effective lifetimes, excludes the necessity of measuring the much later right-spin population at the end of D-runs, and incorporates the transport efficiencies of each population, such as to cancel them to first order.

$$\sigma_{\xi^-} = \xi^- \sqrt{\frac{1}{P_{scs}^+} + \frac{1}{P_{swt}^+} + \frac{1}{D_{swt}^-} + \frac{1}{D_{scs}^-}} \quad (3.24)$$

$$\sigma_{\xi^+} = \xi^+ \sqrt{\frac{1}{P_{scs}^-} + \frac{1}{P_{swt}^-} + \frac{1}{D_{swt}^+} + \frac{1}{D_{scs}^+}}$$

Where we've used $\sigma_N = \sqrt{N}$ for each of the detector signals. This then demonstrates the multiple advantages in utilizing this different calculational method; reducing the uncertainties to depend only on detector counts, avoiding extra terms, and circumventing the need for fitting procedures. Additionally, since the signals are significantly smaller for the switcher unloaded population, this one measurement is the limiting factor in minimizing the uncertainty.

Following this, we may produce our 0th order calculation for the polarization for each year:

Table 3.4 0th order calculations of the depolarized fractions ξ for each of the 2 run years, and loading states, determined from experimental data, with no systematic corrections included.

	2011	2012
ξ_{off}	0.0062 ± 0.0004	0.0045 ± 0.0005
ξ_{on}	0.0099 ± 0.0003	0.0070 ± 0.0003

3.4.4 Depolarization Calibration (DC)

Although simulations depend on setting spin flipper efficiencies directly, and theory provides a framework for extrapolating from measured B fields, in-situ experimental controls only allowed 2 inputs for the spin flipper; the input power, and the driving frequency. Although input frequency necessarily results in being driven at this given frequency, correlating directly to theory, input power is highly contingent as to exact geometric, electronic, and physical

parameters of the system, such that the only sufficiently accurate method of determining the results is direct measurement. As such, for consistency it is ideal to correlate this input power to spin flipper efficiency directly. Though, this presents an obvious paradox; the very nature of the problem faced is to determine the polarization, of which a large part is determining this spin flipper efficiency. How then can one approach the problem using the same populations and techniques?

In short, one cannot; at least not perfectly. However, a clever implementation of a specific run type allows correlating each of spin flipper efficiency, input power, and loading rates for a wide range of values. This “Depolarization Calibration”-type run functions by alternating the state of the spin flipper during loading, for short bursts. While polarizations are high, and bursts are sufficiently short (order ~ 0.1 s), the ratio of flipper off vs. flipper on *durations* quickly approaches the effective efficiency in the system; that is, if the flipper is switched to the on state 5% of the time during loading, the system will be loaded with almost exactly 5% inefficiency, as the remaining fractional percentage is washed out.

As such, we may proceed as we do in normal D-type runs, while artificially inflating the depolarized population. In doing so, the low statistics of the depolarized population are improved dramatically, and highly accurate depolarized fractions can be determined. By setting this fractional time lower and lower, a curve may be constructed for modeling the behavior; as it approaches 0, there will remain a non-zero depolarized population. For sufficiently small polarizations, this evolution will behave linearly, and the y-intercept of the fitted line represents the depolarized population at maximum efficiency.

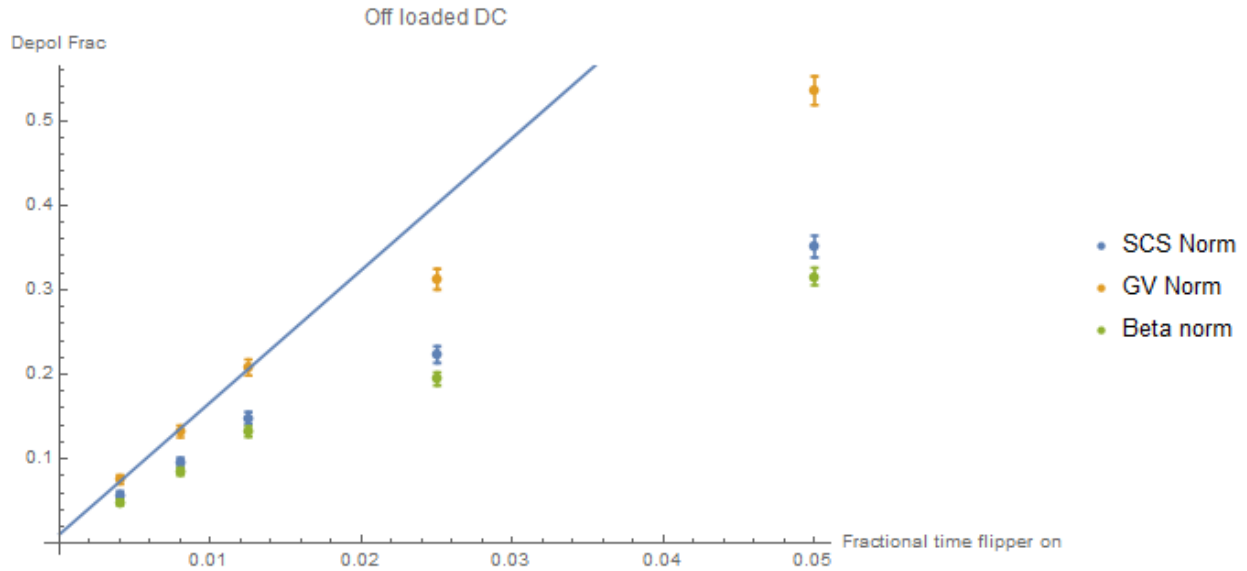


Figure 3.11 Depolarized Fraction measured during multiple Depolarization Calibration runs. The line represents the deviation from linear evolution of lower efficiencies. One also observes the normalization scheme by the gate valve becomes in worse and worse agreement, due to its inability to account for the changing loading rate efficiency in the trap.

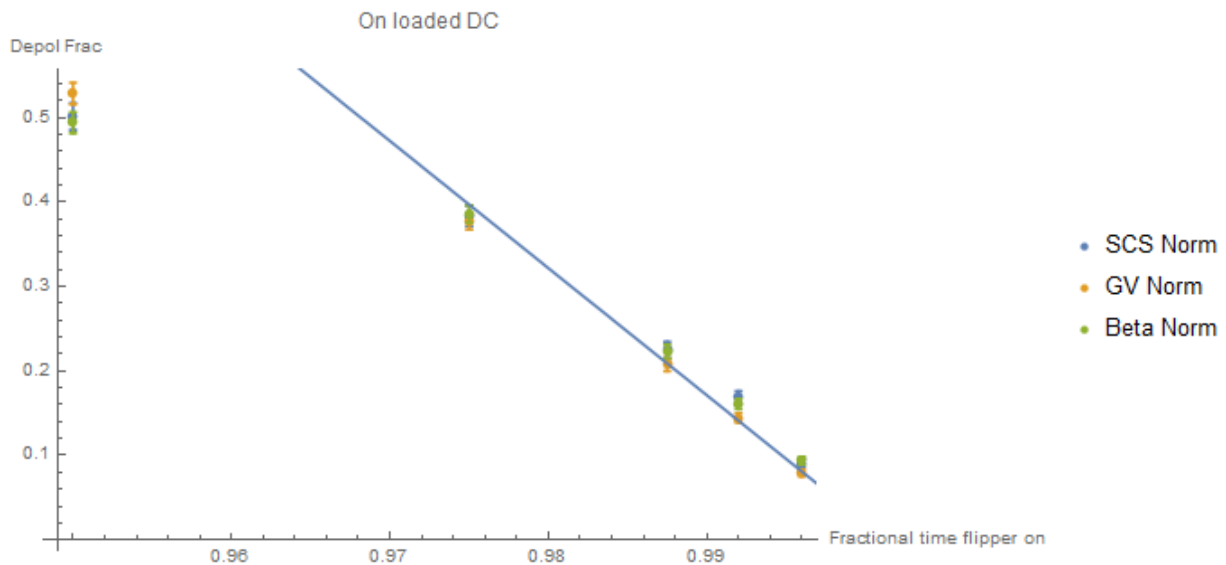


Figure 3.12 On loaded case of depolarized fraction dependence during Depolarization Calibration run types. Exhibits similar behavior to that of off loaded (See Above)

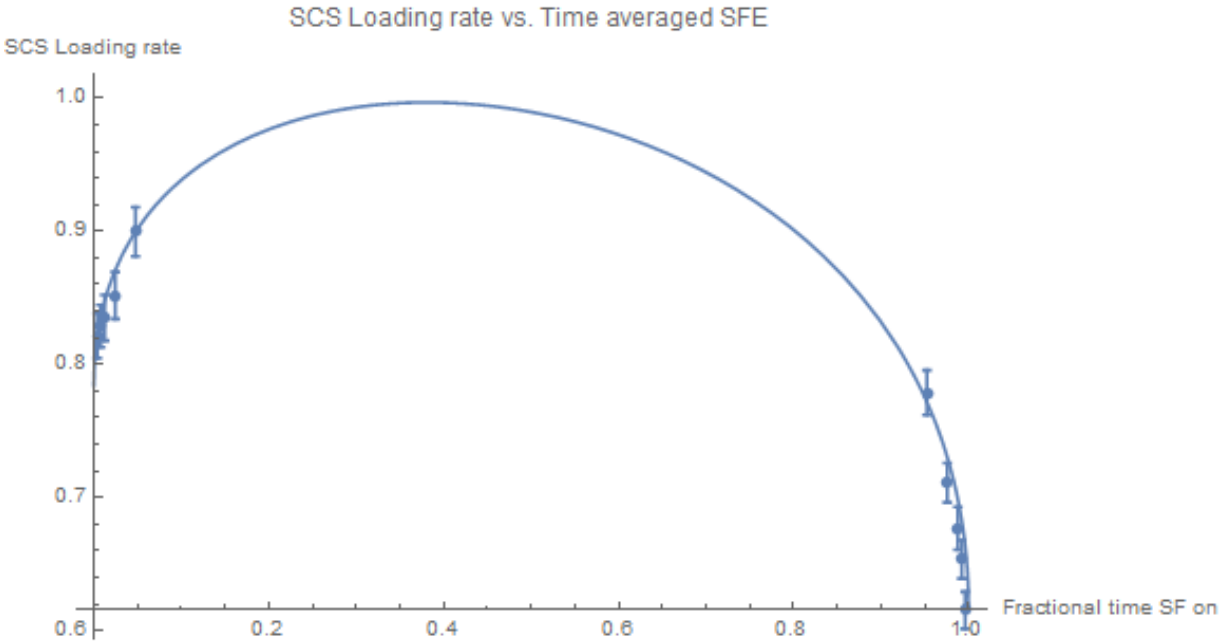


Figure 3.13 Measured loading rate seen in the SCS monitor during depolarization calibration runs. Both the on and off case have been combined into 1 graph, on either side. An empirical curve of the form $\sqrt{x} + \sqrt{1-x}$ has been overlaid on the data, simply as a method for comparing to computation. It is now known that this is likely more accurately modeled as a Beta function, but sufficient data was not available to confirm or test this hypothesis.

However, a critical difference is important; while the spin flipper has some maximum efficiency during the on-loading phase, it has *perfect* spin flip efficiency while in the off-loading phase; 0 input power **must** correspond to 0 spin flip efficiency. If this were the only inclusion of the term, we would then have a simple and direct means of separating those depolarization sources that originated from the spin flipper efficiency itself, and that of the inherent depolarization via interaction with the guide system. However, things are of course more complicated, and further attention to the procedure of the D run reveals that the spin flipper efficiency *does* participate in the off-loaded state, during the spin flipper trapping portion.

This is then our first encounter with the largest issue for determining the depolarized fraction in UCNA; the spin flipper efficiency is convolved with other mechanisms of

depolarization, and the spin flipper efficiency affects both the off- and on-loaded states, but in *different ways*.

3.4.5 Spin Flipper Efficiency (SFE)

To deconvolve the true efficiency of the spin flipper requires some method of varying this efficiency. As seen from the DC run type, artificial inefficiency can be injected into the system by varying the time averaged loading state of the spin flipper, but other than approaching the upper limit of efficiency while fully powered during loading, the method contains no predictions from theory as to the limits of this efficiency. Upon examining eq (3.16), it becomes clear there is only 1 variable that can easily be controlled, for use as a predictive mechanism for the efficiency.

The holding field B_0 cannot be easily changed; first, because of the large and slow process of ramping both the holding coil along with the shim coils, and second, because doing so would alter the resonance point's field, thus detuning the spin flipper itself, as well as causing a difference in velocity spectra of the 2 spin states upon reaching the decay trap, along with a new tailoring procedure for producing the appropriate shim fields. Adjusting the field gradient is similarly difficult and would fundamentally alter the system from that which we hope to measure. Instead, then, the remaining control is allocated to the rotating field's magnitude, or in terms of what is experimentally determined: the input power to the spin flipper.

It is with this in mind that the spin flipper efficiency (SFE) run type was developed. It is performed similarly to the DC run type, but instead of varying the spin flipper's running time, the input power during loading is varied, and a normal D run may proceed.

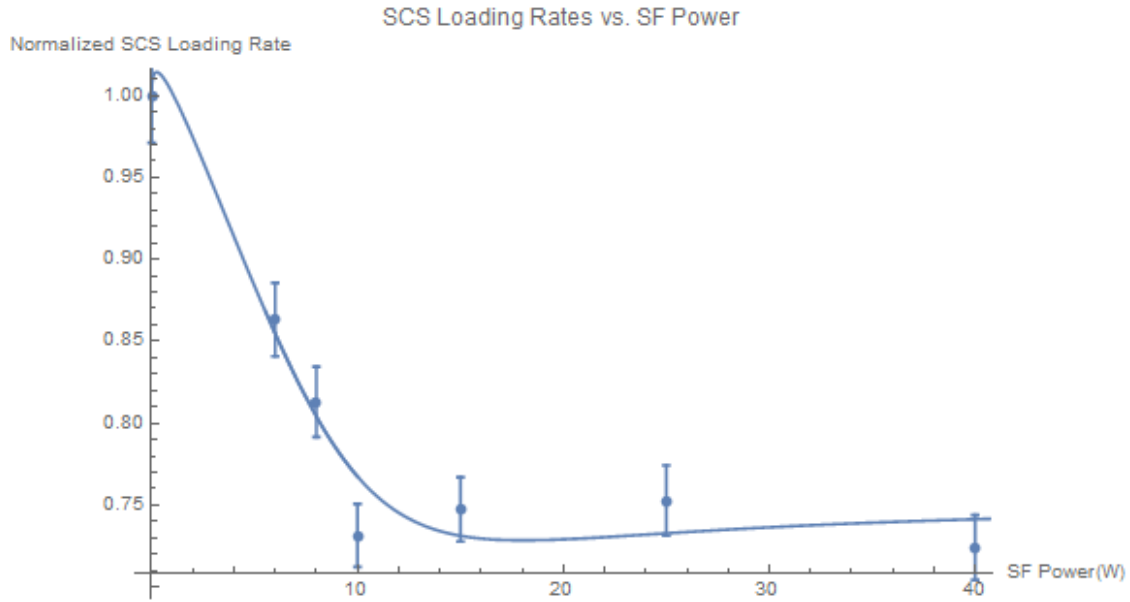


Figure 3.14 Loading rates measured in the decay trap during the loading phase, while varying Spin Flipper power. 40W corresponds to typical input power during on-loading, and 0W corresponds to the simple off-loading case. Note; Due to the $1/B^4$ dependence of the efficiency, most of the interesting behavior actually lies between 0-1W, as can be seen by the fitted curves maximum being produced slightly past the y-axis. Data is fit to CDF of Beta function.

3.4.6 Depolarization Evolution (DE)

Intrinsic in the assumptions for both the D and D+P analysis methods is that the populations of each species behave consistently in time, and independently from their origin. However, this is known to be false, albeit a reasonably small effect. Therefore, the Depolarization Evolution run type was developed to probe any deviation from the simple exponential behavior measured when each spin state is trapped behind the shutter, to compare majority and minority loaded cases.

An example of measurement results can be seen in **Figure 3.15**. Here, we would expect in the simplest case for the data points to fall on the exponential decay curve, determined separately from Shutter Trapped lifetimes, wherein the depolarized state was loaded as the polarized(majority) population. Deviations from the curve indicate a difference in behavior for

those neutrons produced via depolarization, rather than when loaded as the polarized state. Additionally, simple closed form solution points have been included (see eq (4.4)), that describe the population for which all depolarization is fed by the majority population, and none from spin flipper inefficiency. We note that the largest discrepancies between the 2 models appear at times less than the effective lifetimes of the shortest lived population while in the trap. However, probing this time scale was unavailable experimentally, as it coincided with the required time to clean the upstream guide before measurement could proceed.




Figure 3.15 Graph of wrong spin counts in Switcher detector for Depolarization Evolution runs, normalized to 1 for the standard D run holding time at 30s.

3.5 Data Quality Checks

3.5.1 Runs 21k-22k: Anomalously High Shutter Trapped Decay Rates

One set of runs (from roughly 21k-22k), exhibited anomalous behavior in terms of lower lifetimes inside the decay trap for both spin states. At first glance, these errors went overlooked, but by considering a variety of metrics, we are able to establish those runs which are suspect, the effect on calculations predicted by these differences, and the physical mechanism for the

discrepancies. This error was also noted by A.T. Holley in an internal document, but was not recognized at the time as an issue with the system, especially since multiple other characterizations used for evaluating the system appeared nominal.




Figure 3.16 Decay rates fit to shutter trapped populations. Excessively high values between 21k-22k indicated a possible issue with the decay trap.

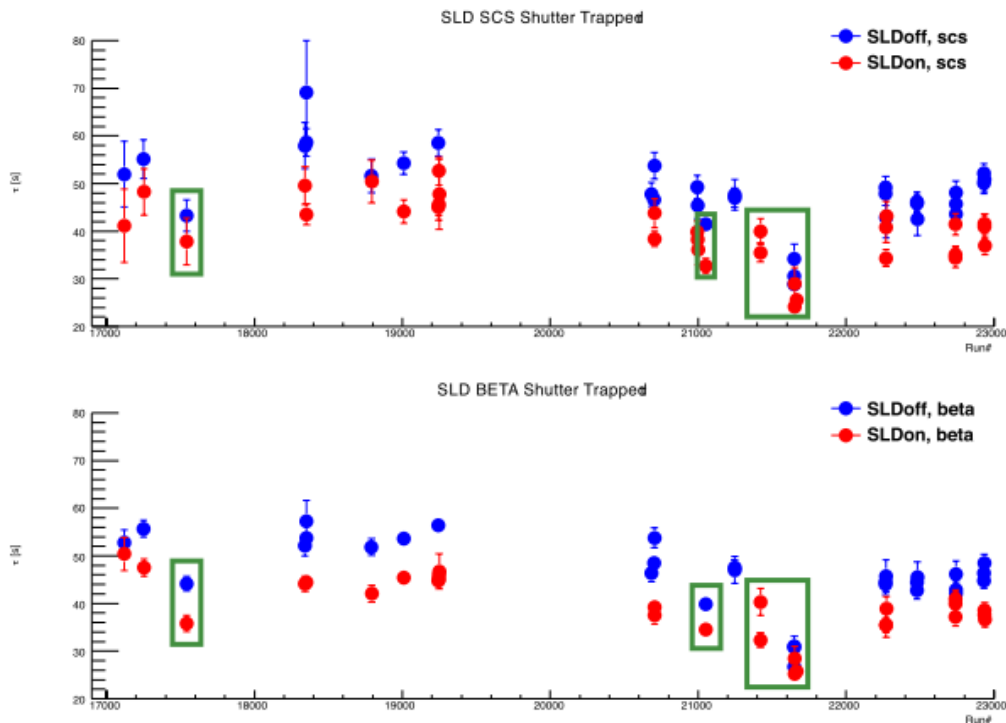


Figure 3.17 Initial observations of anomalous shutter trapped lifetimes. Fit lifetimes suggested experimental issues, but inclusion of spurious runs in analysis masked the degree of the effect at the time.

Our first observation is the establishment that the ratio of the steady state loading rates as measured by the iron foil monitor against the gate valve monitor remained constant over the entirety of running, including no discrepancy between annual run cycles (which switched at run #20k). **Figure 3.18** demonstrates that 1) there was little to no change in detector efficiencies for either the iron foil or gate valve monitors over the course of each run cycle, nor from year to year, and that 2) the behavior of the guide systems from the gate valve ratio up to the iron foil monitor (after the spin flipper, before the shutter) was largely unchanged from year to year. We will omit the possibility that all detectors varied in efficiency by the same fraction, as it is enormously unlikely, due to them being at different locations, powered and measured by different electronics, and utilized different methods of detection (^3He and B counters).




Figure 3.18 Ratio of Iron-foil monitor to Gate-Valve monitor signals during loading.

Next, we observe a change in the ratio of the iron foil monitor to the SCS monitor (located inside the decay trap) from the 2011-2012 cycle (17k-20k) compared to the 2012-2013 cycle(20k-23k). This suggests a gain difference was present between the 2 years for the SCS monitor. However, no peculiar behavior is exhibited during the region of interest between runs 21-22k, suggesting that the difference in lifetime could not be attributed to faulty gain (nor would one expect varying gain fractions to cause such an interpretation).




Figure 3.19 Ratio of Iron-foil monitor to SCS monitor signals during loading.

Next, we address anomalous background counts. For the SCS monitor, very low backgrounds were detected throughout all running, with no apparent difference in any set of runs (within statistical uncertainties). However, inflated values seem to occur for the beta monitors during runs centered first around #19,500, and next for runs near #22,500. These, however, also do not appear during the region of interest, and as only 1 of the detectors is affected (beta monitors), suggests a detector pedestal issue, instead of a physical mechanism inherent to the system, that would affect both spin states.

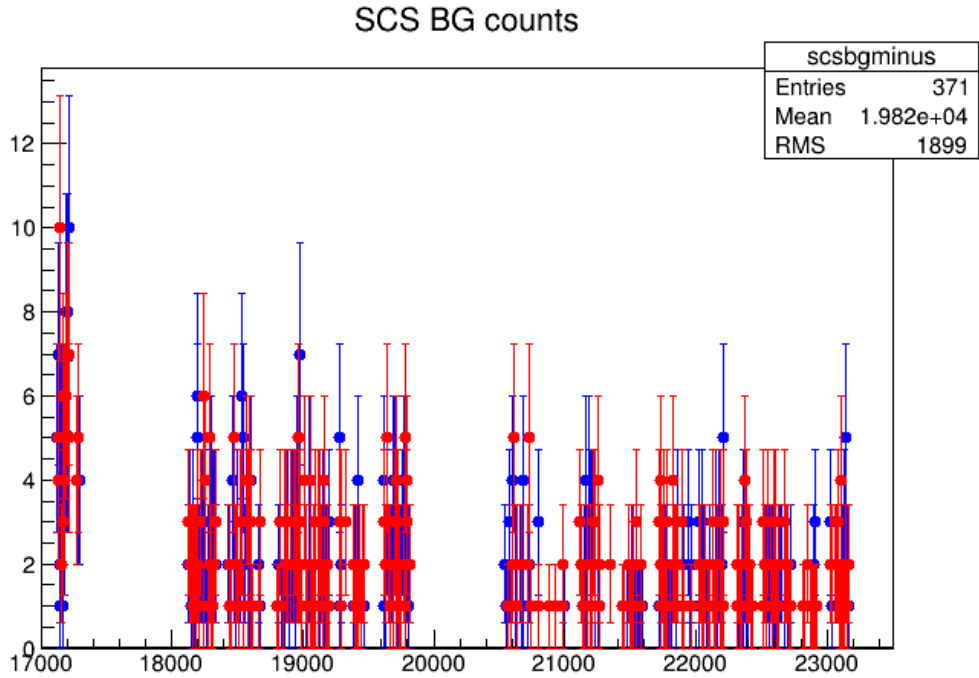


Figure 3.20 Background counts in SCS monitor vs run #.

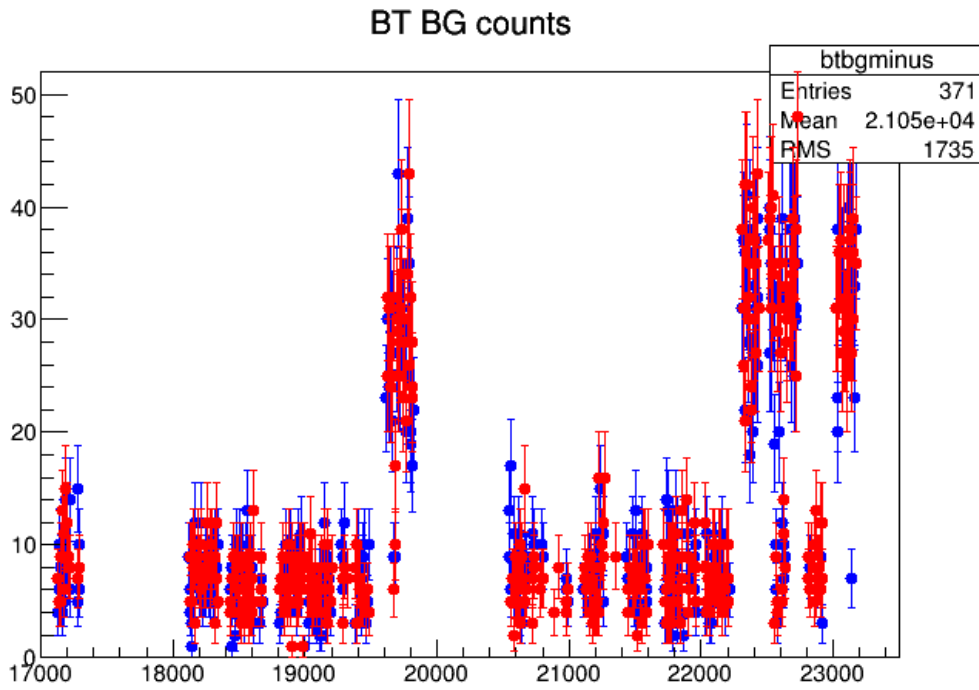


Figure 3.21 Background counts in β monitors vs run #.

As only 2 shutter lifetime drain runs were performed during the area of interest, evaluations on the behavior of the populations during D and P runs was required to establish consistency of the behavior. By observing the number of neutrons drained after shutter holding times in either run type, it was found that there existed a deficit of expected counts for both spin states as well. This then leads to the conclusion that the system downstream of the shutter must have had some effective absorber inside the system. End-cap foils having structural issues is in general possible, though this would result in dysfunctional behavior of the MWPCs present in the beta detectors, which was not observed, therefore eliminating them as the culprit. This then leaves only a few possibilities; 1) An issue with the shutter's movement leading to high losses on only the upstream side, or 2) An issue with some portion of the guide system, i.e. the rectangular guide or the decay trap guide. Finally, as no repairs were performed on the shutter, yet the behavior ceased, we are left with only option 2). Upon inspection of the run-log, it was indeed found that sealed source calibrations occurred both near runs 21k, and 22k. Therefore, we are led to believe that improper reinstallation of the source holder plug in the system behaved as a hole in the decay trap, decreasing lifetimes when subjected to extended sampling of the decay trap.

We must then ask which data, if any, would be reliable in this situation. It so happens that the actual asymmetry of the decay is not affected by changes in the lifetime of either spin state; or rather, that the measurement of the depolarized fraction in each case wholly includes this effect, and if measured properly, can be included in calculations during the region. However, with varying lifetimes, it is then required to evaluate the system separately during this run period; All calculations and normalizations must be performed within the region of interest, mutually exclusive to those outside of it. With this the case, performing separate calculations of the poorly performing time range, it was found that the very sparse number of P runs, and lowered statistics

from the increased losses, led to a negligible increase in statistics for overall evaluation of the asymmetry. Therefore, although the systematics were in principle able to be determined, the data set during this region was forced to be excluded from final calculations.

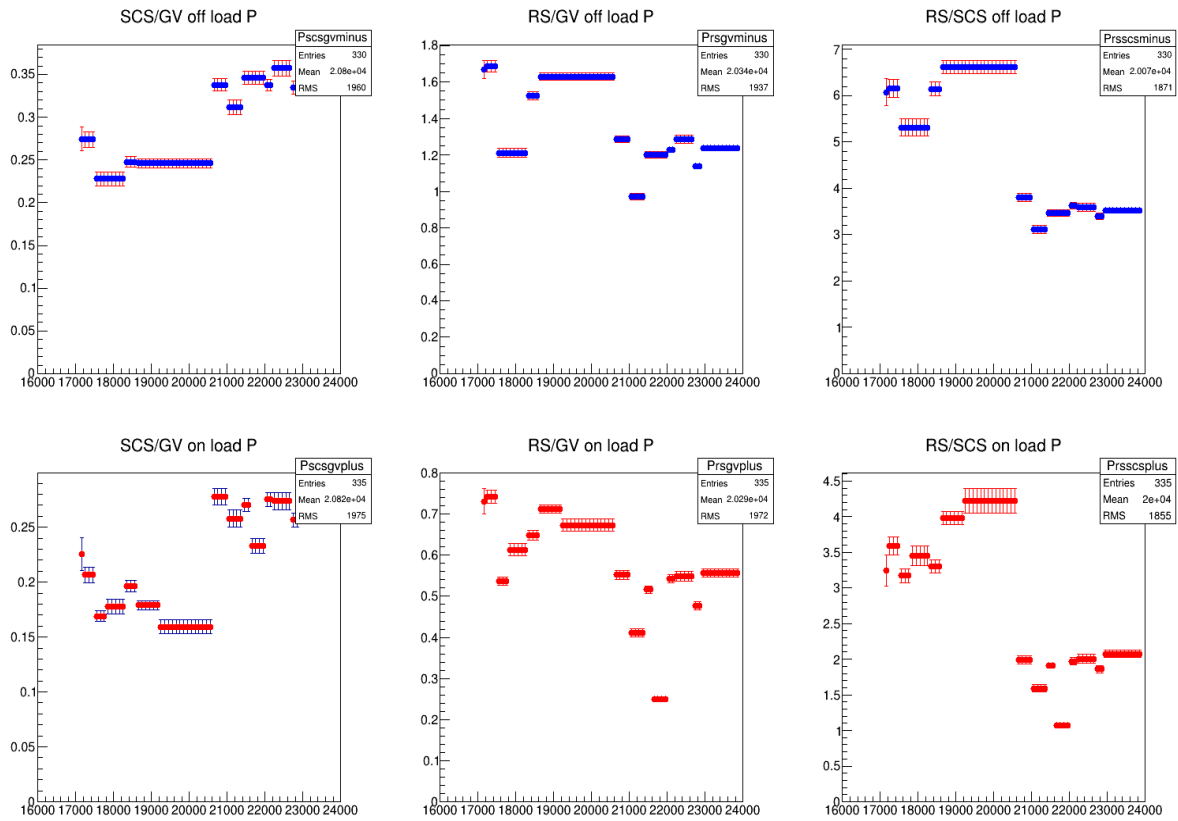


Figure 3.22 Signal ratios for P type runs for (left to right) 1) SCS monitor vs gate-valve monitor, 2) right-spin counts normalized by gate-valve monitor, and 3) right-spin counts normalized by scs monitor.

Chapter 4

Computational Simulations

4.1 Computational Resources

In addition to theoretical analysis and experimental characterizations, a critical element in confirming the reproducibility of the data, and serving as a check on systematics, is our use of computational transport simulations. We implement an in-house C/C++ based Monte Carlo code, first developed by Adam Holley, specifically for observing the evolution of UCN populations within material traps over time. By tracking individual particles from creation to loss, we are able to implement highly scalable models amenable to use in cluster computing.

Randomization is integral to nearly every aspect of the evolution and interaction of each neutron in the system. As such, whenever a probability is required, a specific subroutine is called to generate a random number, or “roll” the odds, producing a value between 0 and 1. This randomization is seeded initially, and again every ~1000 particles. To prevent repeating data appearing as independent runs, each seed is determined using a combination of particle number, run number, processor number, and a modulus of the Unix Epoch Time. Although only a pseudo-random generator provided by the C package, by restricting our values to between 0 and 1, and coupling the seeding to every variation that may be used to delineate a given particle, with a running count, serves to function exceedingly well as a truly random generator.

Particles are modeled semi-classically; each neutron is assumed to have a parabolic trajectory inside the system. A quartic solver is implemented to determine the intersections of the neutrons’ predicted trajectory with that of the guide walls (typically cylindrical). The neutron is then evolved to the intersection, and a check is performed to determine if a decay occurs during

the transit time. If so, the particle is lost, and a new particle is generated at the start of the system. If no decay occurs, however, various material interaction checks are performed. Losses can occur through straightforward loss per bounce (lpb) parameterizations, or from more complicated material interactions.

At the intersection, we then determine whether the particle reflects or transmits through the barrier, using the relation

$$LPB = 2 \left(\frac{W}{V_F} \right) \sqrt{\frac{E_{\perp}}{V_F - E_{\perp}}}, V = V_F + iW \quad (4.1)$$

For particles reflecting from a finite potential barrier as seen earlier, however with the introduction of a complex component W , representing a loss parameter. Each guide component is constructed with assigned material potential values. If penetration does occur, the particle is then evolved in a bulk interaction state. This state develops much the same as the normal (vacuum) trajectory solutions but includes an additional loss term associated with the material itself. This term can represent losses from upscattering and absorption, but in practice serves to modify the effective lifetime. If there are no boundaries defined outside of the material, the neutron is considered lost. However, this is not the case for some critical components, such as the various gate valves and shutters, or thin foils used to separate the vacuum of the system from the D_2 source.

If reflection at an intersection is evaluated to occur, the system then rolls the material associated property of specularity. A low specularity refers to a very ‘rough’ material, resulting in isotropic scattering from the surface. A high specularity, associated with highly polished and clean components, instead produces a reflected angle equal to the incident angle with respect to the normal vector, as per the classical law of reflection.

The system utilizes CERN's ROOT package to handle particle tracking. In doing so, we are able to consider our output in terms of 'events'; each time an interaction of some form occurs, an 'event' is recorded, including the various properties of the neutron (position, velocity, spin, etc.), along with a numerical event code, which can be used to examine trends of individual interactions. For example, one may cut on a specific event code to be given only those interactions resulting in a bulk material absorption, and all events returned will consist of the final location, time, velocity, etc. of each neutron lost in this way. With the ability to select individual interaction forms, problematic areas or trends can be quickly diagnosed, aiding in either debugging in the case of errors, or determining critical areas of improvement/sensitivity for physical mechanisms, such as disproportionately high losses in one guide section.




Figure 4.1 Depolarized Fraction evolution inside decay trap over time, in the old "Depolarization Trapping" and "Reload" run types. As the Reload type injects more intentionally depolarized UCN, the increase in the depolarized fraction can be observed. However, this procedure relied on the assumption of re-establishing an equilibrium, producing double the depolarized population, which may not be the case.




Figure 4.2 Depolarized fraction evolution over time, for the spin flipper on loading state. See **Figure 4.1**




Figure 4.3 Plot of UCN density inside decay trap vs. height. As one expects, the population is slightly denser at lower regions, from the gravitational potential energy being similar in scale to the kinetic energies.

4.2 Experimental Reproduction

First requirements were to reproduce given signals as closely as possible, by finding those sets of parameters which would produce close agreement to measured observables. With upwards of 200 guide components in the geometry file, and half a dozen parameters each, it was important to categorize those regions of interest which behaved in bulk similarly to their individual components (which is to say, changing of a parameter for any of multiple pieces giving effectively the same result, so as to be consolidated into a general Region of Interest).

Table 4.1 Table of various independent parameters to be tuned, and ways to measure their effects, in reproducing the experimental behavior of the guide system. Note; \pm denotes those values which must be found separately for the on and off loaded spin states.

Regions of Interest	Parameters	Observables
Decay trap	Loss per Bounce $\sim 1 - 5 \cdot 10^{-4}$	Shutter Trapped $\tau(\pm)$
Rectangular Guide	Material Potential $\sim 70 - 350$ neV	Flipper Trapped $\tau(\pm)$
Shutter	Bulk Potential	Open Trap $\tau(\pm)$
Low Field Region	Specularity $\sim 0.8 - 0.99$	Guide Cleaning $\tau(\pm)$
Spin Flip Region	Depolarization per Bounce $\sim 10^{-6} - 10^{-5}$	Depolarized Fraction (ξ_{\pm})
Peak Field Region	Spin Flipper Efficiency $\sim 0.99 - 1$	Loading Rate Ratio
		Loading Rate SFE Dependence (\pm)

Each of the parameters listed in **Table 4.1** represent highly conservative bounds for possibilities of parameter explorations, representing at least a 2- σ confidence of the materials falling within the posited ranges. However, even extrema cases outside these bounds were tested as consistency checks, and while they behaved as one would predict, they inevitably produce meaningless results. Naturally, these same materials of interest have been measured repeatedly, for decades prior to this experiment, and indeed directed our selections of materials to begin with. Many estimations and evaluations of the parameters of the UCNA components specifically,

were detailed recently in [3], and an exploration into the topic of creating and quantifying the UCNA DLC guide component coatings can be found from [1] (See also Appendix E).

In theory, each parameter, with each component, can be varied independently, and are all, to some extent, coupled to one another, such that effects may compound in non-additive ways. Let us consider a quick example for illustrative purposes. For a single component, say the Decay Trap, and 2 parameters of interest (here we will choose loss per bounce, and material potential), it may be possible to at the same time, increase the rate of loss per bounce, decreasing all lifetimes, though some more significantly (namely, those which sample the decay trap a greater portion of the time). However, we may also simultaneously increase the material potential, such that we reduce losses associated with higher energy neutrons penetrating the walls or being upscattered. In doing so, we can observe right away that we have now selected preferentially in favor of a higher energy spectrum population; if we tune the lifetime to be the same (for at least 1 case), increased the losses to every neutron with a flat %, but decreased losses from higher energy neutrons, we have conditioned our population with a bias towards higher energy. What comes to be even more of a problem, however, is that it is likely in so doing, we may be able to match ONE, or a couple of the aforementioned observable lifetimes, but as each observable is effected differently, we will have a cascade of effects, that thus require further and further tuning to undo changes we made, while still implementing new conditions on the decay trap parameters.

This, then, is the fundamental problem with the reconstruction of the experiment in simulated space: the highly interdependent web of parameters, observables, and finally, statistics. If we were to try to test these correlations, say with 3 values for each parameter, to give the ability to fit the most basic of a linear curve, to understand evolutions and dependences, then we see that we would have a choice for each of the 30 parameters; region 1 may have each of 5

parameters, so we may examine changing each of them individually, or varying both the 1st and 2nd, and so on, such that they behave as individual elements.

Additionally, we must also consider we intend to vary the spin flipper efficiency, as this is one of the most significant contributions to determining the depolarized fraction. And finally, we must consider which types of runs must be completed, to extract the desired observables. For these, the Guide Cleaning τ can be obtained simultaneously with the Shutter Trapped τ , and the Loading Rate Ratio is able to be measured in all run types, so requires no extra contribution. However, the depolarized fraction requires both D and P type runs, for a total of 4, and each other lifetime not mentioned will require 1 run type per spin state, totaling 12 runs to fully diagnose a given system. If we allow the number of values the spin flipper will take to be the same as that of all the other parameters (for demonstration purposes), where we let N be the number of different values we would like to test for each parameter, we obtain an estimate of;

$$\begin{aligned}
 \widehat{12} \text{ Runs} * \widehat{N} \text{ SFE Values} * \overbrace{\left(\widehat{N}^5 \right)^6}^{6 \text{ Regions}} &= \mathbf{12 * N^{31} Runs} \\
 12 * 3^{31} &= \mathbf{7.4 \cdot 10^{15} Runs}
 \end{aligned}
 \tag{4.2}$$

Which, amusingly, would require almost exactly 1 run per minute, every minute, since the beginning of the universe to complete. This leads to the simple conclusion that exhaustively probing parameter space is not a viable means for determining values.

However, it does raise some concerns; this order of magnitude estimate was not wholly unreasonable: we reduced the whole system to only 6 critical regions, which do indeed experience different physics, and only 5 parameters, all of which ARE expected to, and observed to, have effects on population behavior. Finally, 3 values each is a very conservative list to test; even with direct constraints on allowable values from experimental measurements, small

variations in these values causing large deviations would want to be explored. Yet we still arrive at an incomprehensibly long time to check.

As such, our method was not one of rigorous exhaustion, but of judicial choices between 1- and 2-D scans. The first, and most obvious method, is to simply ignore correlations; explore along 1 axis of a parameter at a time, find the best fit, and repeat for a different parameter. This is indeed useful, and of course the primary method of attack. For example, agreement with the flipper-on, shutter trapped lifetime was extremely sensitive and selective to the material potential of the shutter. In this regard, the one parameter scan immediately singles out a strongly constrained parameter for all other possible scans (see **Figure 4.6**). Multiple iterations of this procedure are required for each variable, until either total convergence, or bounds of stability are found. However, as demonstrated in **Figure 4.4**, there are scenarios in which varying only 1 parameter, while remaining within a solution region, may never find other, potentially better parameter spaces of agreement.




Figure 4.4 Simple cartoon depicting 2 potential solution regions (circles), being inaccessible by exploration of 1-D scans from anywhere within one solution region (rectangles). It is these cases in which correlation studies are required.

Furthermore, one must ask, what IS the solution region, in these cases? On first approach, one may think it obvious; construct a global χ^2 metric and minimize. However, not all parameters are equal, and to blindly consolidate all information into a single χ^2 just increases the difficulty unnecessarily.

For example; while tuning the depolarization per bounce can have an enormous effect on the depolarized fraction (as one would expect), it has next to no effect on the lifetimes, since even the highest estimates for depolarization rates are still orders of magnitude lower than those for losses. As such, to include the depolarized fraction on equal footing with that of lifetimes, when concerned with exploring the depolarization space, one buries away key insights that lead to more appropriate levels of separating the variables.

However, this is not to say that one must ignore these effects entirely; varying the loss per bounce in a region can indeed alter the depolarized fraction significantly as well, depending on the relative sampling rate of each spin state separately to that region.




Figure 4.5 Example of dependence of depolarized fraction versus depolarization per bounce. While the DPB will not significantly affect the lifetimes, it can control the produced polarization levels significantly. It is therefore a 2nd order tuning parameter, as it has effects on fewer observables, and should be tuned near last.

So, the general approach was first, to estimate parameters for all guide elements, associated with those that had been measured, or where measurements were not available, on nominal values reported for similar materials in the literature. From here, we first attempted to reach agreement with the various lifetimes, as well as the loading ratio of off to on loaded states. To do so required identifying which components, and which parameters of each component, contributed most significantly to which lifetimes. In this way, a sort of linear algebra problem began to form, in which each tunable parameter would affect each observable by varying magnitudes, yielding a form of coupling coefficients. With this, we were then able to tune until good agreement was had.

However, as mentioned, 1-D scans are sometimes not enough, and introducing correlated variations is required. In one instance, for example, an area of good agreement was found, but in

only 5 of the 6 lifetimes. No further deviations could improve agreement. In this instance, we imposed a radically different value for one of the parameters (in this case, the material potential of the rectangular guide), and with a qualitative understanding of the behaviors governed by each parameter type and effects expected to be produced on each run type, adjust material potentials and loss-per-bounces for multiple surrounding materials, to try to partially compensate for the change. Afterwards, we repeated our linear grid-walk approach, and found a separate area of agreement, that was now able to facilitate all measured lifetimes to much improved consistency.

One set of high-agreement geometry parameter sets is provided in **Appendix C**. This, along with the associated connections file in **Appendix D**, describe the physical parameters used in many of the most recent simulations. As discussed, it is not necessary that there exist only 1 set of parameters capable of accurately reproducing all observable behavior within the system, and further optimization is certainly available for exploration.

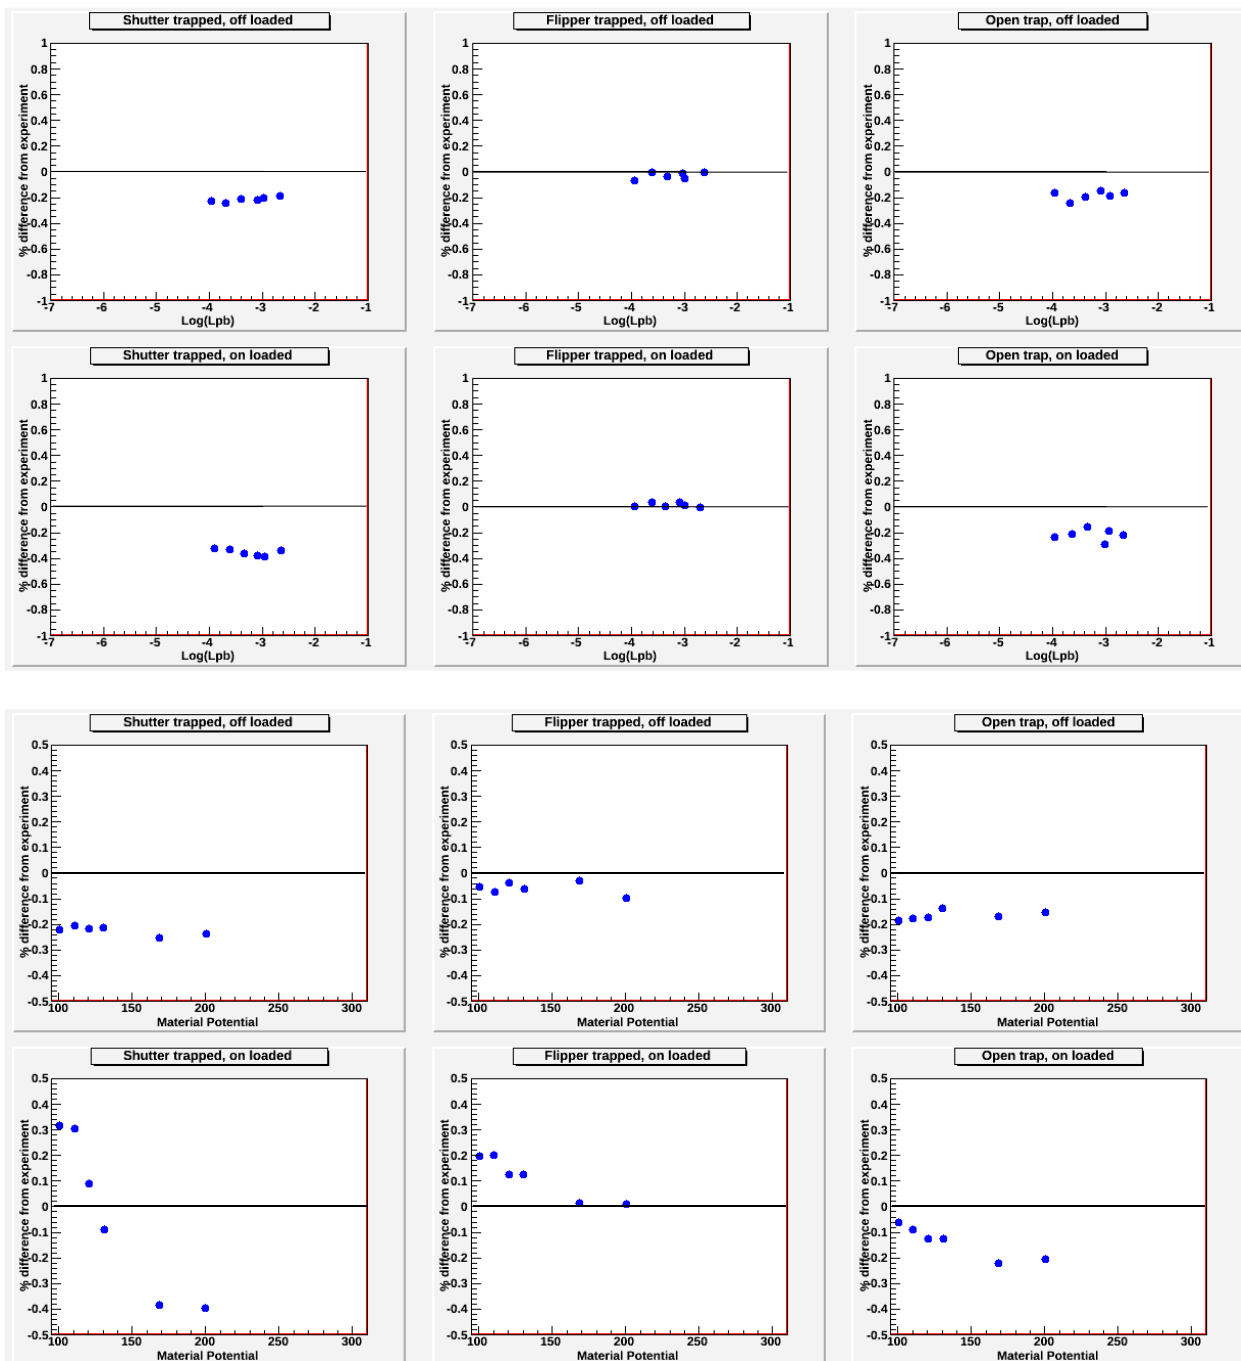


Figure 4.6 Comparison of changing parameters of the UCNA decay trap shutter. (Top) Varying the material potential demonstrates a stronger dependence on the flipper-on loaded state, compared to flipper off. (Bottom) Varying the loss per bounce of the shutter has little to no effect, on either population. This is attributable to the lower field at the shutter position, resulting in an acceleration of the on-loaded states. This also demonstrates the signature of selectively-differential component-parameter combinations, identification of which are essential to narrowing parameter scan scopes.

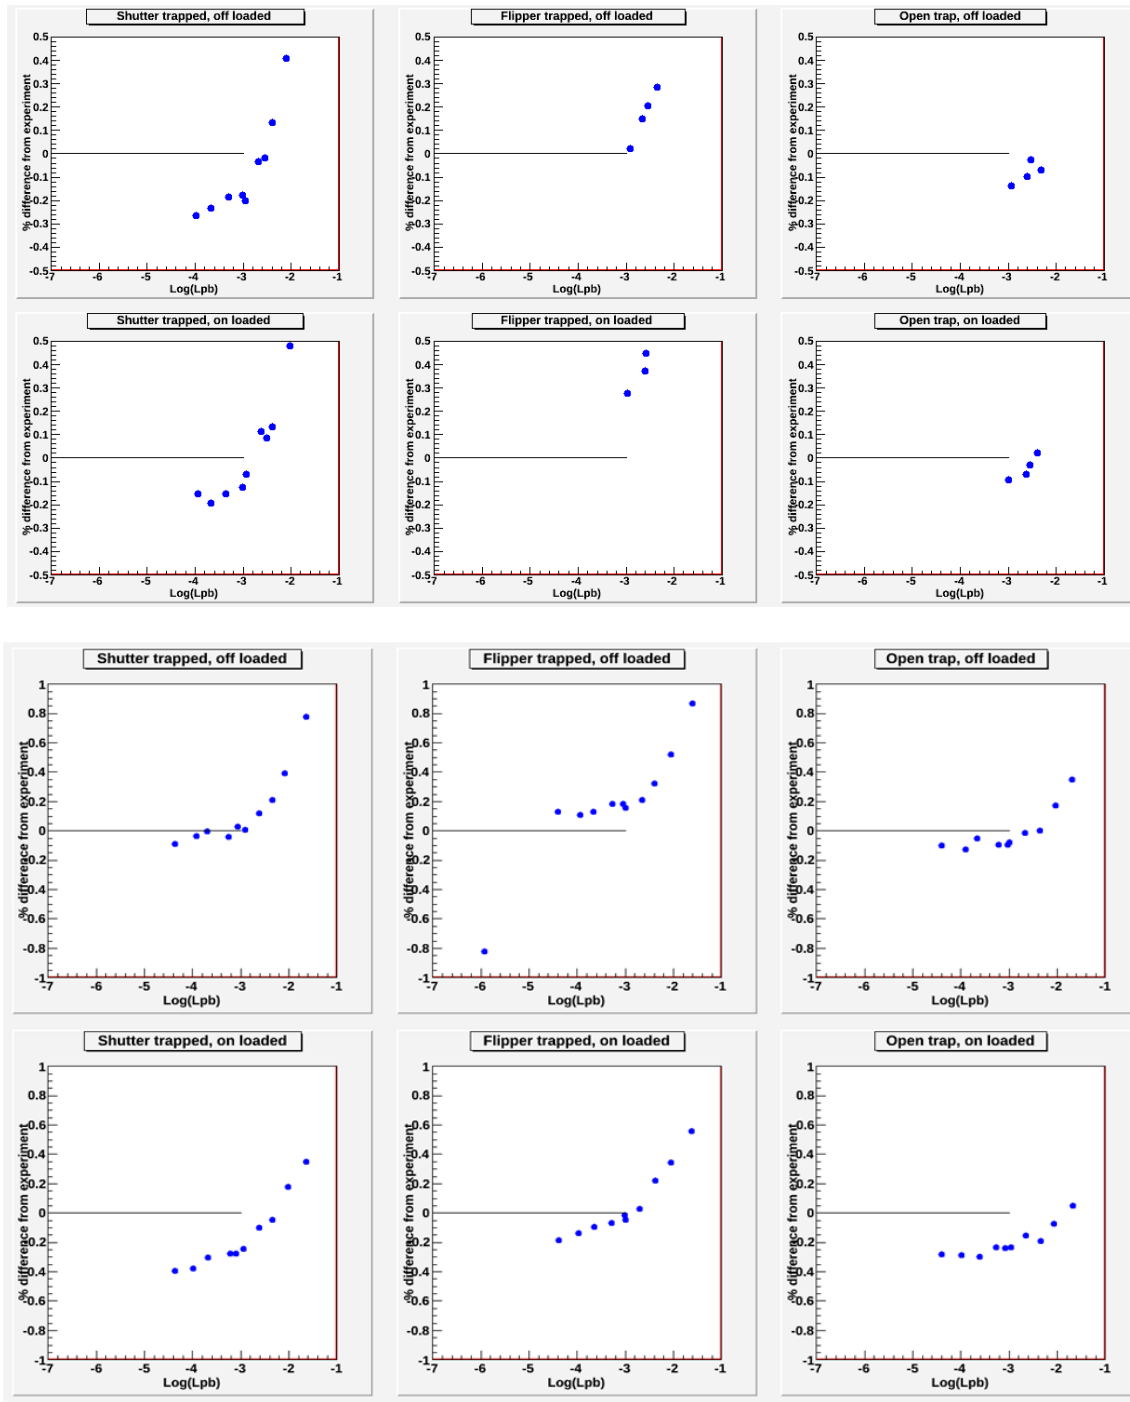


Figure 4.7 Example of behavior at an island of stability. Top and Bottom representing varied alternative parameter choices, such as rectangular guide material potential. Varying the Loss per Bounce (of the decay trap in this case) was able to bring a few lifetimes in near agreement, but some were fundamentally separated (Ex, top-middle vs. bottom-right). The regions of these discrepancies can hint towards the regions that may need adjusting, and subsequent compensation.

4.3 Systematic Corrections

4.3.1 Feeding

One inherent flaw in the method of reconstructing the depolarized fraction within the trap is the measurement process itself affects the population, over an extended period of time. One major systematic contribution is the effect caused by closing the population within the trap behind a shutter; the very mechanism that allows the experiment to significantly reduce the previously leading uncertainty of background during measurement.

For a qualitative understanding, we will first construct the differential rate equations for each population inside the decay trap. We would then have

$$\begin{aligned}\frac{dN_p}{dt} &= +R\epsilon_{SF}\epsilon_p - \frac{N_p}{\tau_p} - BLN_p \\ \frac{dN_d}{dt} &= +R(1 - \epsilon_{SF})\epsilon_d - \frac{N_d}{\tau_d} + BLN_p\end{aligned}\tag{4.3}$$

With R the loading rate from the upstream beam, ϵ_{SF} the spin flipper efficiency, ϵ_p/ϵ_d the loading efficiency of the polarized and depolarized populations respectively, τ_p/τ_d the lifetimes within the trap for each population, B the rate of wall bounces per second, and L the losses per bounce. These equations represent being created by the beam, each decaying exponentially, and finally that the much larger polarized population will also feed the smaller depolarized population via wall depolarization events. Although a reverse term of the depolarized population feeding the polarized population is technically present, it is negligibly small for all scenarios considered.

These equations can then be solved simply, as they are not mutually coupled; the polarized state is assumed to have only negligible contributions from the depolarized, which then means it can simply be expressed as an exponential, and its depolarization rate can be absorbed into its effective lifetime. The depolarized population is slightly more complicated, but the

general form of the solution can be readily guessed, and the requirement for the coefficients can be derived. The solution is that of a double exponential, and can be given by the form

$$N_d = \frac{N_0^d}{1 + \frac{1}{\frac{1}{\tau_d} - \frac{1}{\tau_p}}} \left(e^{\frac{-t}{\tau_d}} + \frac{BLN_0^p}{\frac{1}{\tau_d} - \frac{1}{\tau_p}} e^{\frac{-t}{\tau_p}} \right) \equiv c_1 e^{\frac{-t}{\tau_d}} + c_2 e^{\frac{-t}{\tau_p}} \quad (4.4)$$

There are then 2 major contributions that create a significant deviation from the expected evolution of the populations. First, when the measurement is begun, the source R is turned off, leaving only the remaining 2 terms for each population. This then means that while the polarized population decreases exponentially, with the depolarization rate being included in the effective lifetime, that the depolarized population maintains a positive growth term, from that depolarizing rate of the first population. In other words, the depolarized population need not obey simple exponential decay in early times. Additionally, for depolarized lifetimes much shorter than polarized lifetimes, at later times populations will instead exhibit decay corresponding to the lifetime of the larger population.

Furthermore, these issues are compounded significantly by the closing of the shutter. Upon closing the shutter, lifetimes of both populations are increased, by magnitudes of order 50%. In doing so, the polarized population can be easily adjusted by simply measuring this decay rate under these circumstances through the detectors within the trap. However, when any given population is under conditions such as to be easily detected via the monitors within the decay trap, they will have been primarily generated by the loading beam, and thus not represent crucial characteristics of the depolarized population's behavior when fed by a primary population. This is to say; if the depolarized population is significantly created by the depolarization of the majority (as opposed to from inefficiency in the spin flipper, which will appear instead as a

loading rate dependent effect), then the closing of the shutter will leave the primary growth term of depolarizations unchanged, and will *slow* the decay of the population, which was already in a steady state equilibrium. With the growth term maintained, and the loss term decreased, for a given equilibrium, it may then be possible to observe that closing the shutter, and turning off the beam, will not only cause an inflation of the depolarized **fraction in the trap**, but also counterintuitively **increases the net population** of depolarized neutrons, and accordingly will drastically deviate from the assumed behavior in the initial calculations.

By observing this uncharacteristic growth term, we can conclude that both spin states will necessarily **produce** more depolarized neutrons through the measurement process, leading to **overestimates** of the depolarized fraction in both cases. If the difference between the polarized and depolarized populations lifetimes is significant enough, then the systematics develop as a slower decay rate than would be predicted, leading erroneously to the conclusion that the initial population must have been much higher, if it were to be decaying at the previously measured decay rate of the population when in the majority state. Additionally, if the difference in the depolarized population's lifetime **with** vs. **without** the shutter are significantly different, a short-term net growth can be produced.

A critical observation from this analysis is the development of a symmetry between the 2 spin states. While spin flipper efficiency behaves differently for each population, and each spin state has respective lifetimes, and velocity spectra for each different area of the experiment, this effect is the only one that depends not just on the spin state, but on whether or not that state is the **majority** population within the trap. Similarly, although we may artificially construct scenarios that increase the depolarized population via the behavior of the spin flipper during loading, in doing so we only dilute this effect, and indeed we are unable to modify the rate of depolarization

via this feeding term. Finally, through (mostly) coincidence, the time scales within which the net population growth may be observed, roughly overlap with the times required to complete the draining of the upstream guide sections for background removal. This means then that without significant alteration to the experiment, to hasten the cleaning process, that direct observation of the net growth will be impossible. Likewise, the decay rates within the trap are not easily adjustable either, and of course have an upper limit from the inherent instability of the neutron.

With the short-term behavior unable to be approached, we are then left with the possibility of observing the long term behavior, appearing as a modification of the depolarized populations decay rate. This then was the motivation behind the creation of the depolarized evolution run-type. However, this method also required multiple data points, each consisting of multiple lifetimes worth of delay behind the shutter before measurement, utterly crippling the statistics able to be generated in our limited availability of UCN production time. Runs of this kind were performed, but at the time it was not known the significance they held, and the data produced were so small as to be within the uncertainties of the background estimates. We were then left only with reliance on our ability to reproduce these effects in simulations, and interpret their magnitude as related to our measured population's behavior.




Figure 4.8 Qualitative plot, showing the potential increase in **net** depolarized population after sealing the decay trap.




Figure 4.9 Simulated depolarized population in decay trap over time. While being the majority population, or depolarization being generated primarily from inefficiency in state selection at time of loading, population behavior remains similar during the shutter trapped holding time. However, when fed primarily through depolarization interactions within the trap, the populations exhibit non exponential decay, causing significant overestimations of the truly present counts.

4.3.2 Leaking

The other main source of depolarization is via interaction with the spin flipper. For the on-loaded case, a correction is counterintuitively **not** required. As the goal of the analysis is to provide a determination of the equilibrium state of the neutron polarization during β decay runs, it is critical to realize that although a lower polarization is expected during the spin flipper on loading than when off, it still represents the actual depolarized fraction within the trap during the run as well.

However, the spin flipper does require a correction for the off-loaded system. Although polarization is maximum during spin flipper off loading, the measurement schema requires a period in which the population is trapped by the spin flipper. During this trapping, a large population of UCN sample the spin flipper multiple times, and only a small fraction (those which were actually depolarized) will be permitted through. With the imperfect efficiency from the spin flipper, then, some fraction of the much larger population will then tunnel through when missing a spin flip and will be detected during the draining of the wrong-spin portion. This will lead to an uncharacteristically **high** measurement of the depolarized fraction, as missed spin flips will appear as depolarized neutrons. This term has been coined as spin flipper “leaking” and is an asymmetric systematic correction between the 2 loading states.

4.3.3 Conditioning

Once the system is changed in any manner, in order to measure the equilibrium population, that same population is immediately driven from its equilibrium state. Additionally, although the decay trap is tuned to match its B field to that of the resonance point for the spin flip in the AFP, the magnetic field between these 2 points is known to decrease in magnitude for a stretch. As such, we then must account for the evolution of the populations in sampling the

changed system (such as closing the shutter, changing a spin flipper state, etc.), but must also include emergent discrepancies in the corrections between the 2 states due to the different kinetic energies that develop from interaction with the different B fields.

Likewise, we must consider that the spectrum, and similarly the most heavily sampled guide regions, are different for each spin state, and what's more, different between the same spin state, depending on how and where they were produced; those wrong spin states who are produced by an imperfect spin flipper will simply behave as they do when in the majority loaded state, however all populations produced by depolarization will transit different regions, in some cases rather significantly.

4.4 Results

To produce our final estimate of the polarizations, we implemented a method in which the systematic corrections of the experimental data were determined by convolving the statistical uncertainty's projection into the Monte Carlo simulation, and the statistical uncertainty associated with the Monte Carlo simulated measurements themselves. These statistical uncertainties within the simulation dictate the limits of the systematic correction; as the true depolarized fraction at equilibrium is known in simulations, uncertainty in this value becomes nearly 0, and the composite of all the individual corrections can be formed from simple mapping of simulated detector signals to singular depolarized fractions. We are of course able to disentangle these corrections, to find estimates for individual contributions, which is especially useful in determining methods for improving subsequent experiments, by focusing on the leading contributors. We will primarily focus on the magnitude of the total correction here.

We are able to analyze the χ^2 of each simulated value versus experiment, as a metric for evaluating the believability of a given parameter set, based on its production of measured observables consistent with those found in experiment. We use here the standard definition,

$$\chi^2 \equiv \sum_{i=1}^N \left(\frac{(O_i - E_i)^2}{\sigma_{O_i}^2 + \sigma_{E_i}^2} \right) \quad (4.5)$$

With O the observed values from experiment, and E the expected values from the Monte Carlo simulations. We then found our minimum χ^2 of all values produced in our parameter scan, and subtracted this constant from all values, allowing us to search for an increment in the χ^2 . With multiple values of χ^2 , we then turn to the related (but confusingly, not identical to) χ^2 distribution. The probability density function of the χ^2 -distribution is given by

$$\chi_k^2(x) = \frac{1}{2^{k/2} \Gamma(k/2)} x^{k/2-1} e^{-x/2} \quad (4.6)$$

$$\Gamma(z) = \int_0^{\infty} x^{z-1} e^{-x} dx \rightarrow (z-1)! \text{ for } z \in (1,2,3 \dots) \quad (4.7)$$

With k degrees of freedom, Γ the ‘‘Gamma Function’’, and x the dependent variable. We must likewise define the cumulative distribution function, CDF, of a probability distribution function, PDF, of a given statistical distribution, f ;

$$CDF(PDF(f), \alpha) = \int_0^{\alpha} PDF(f, x) dx \quad (4.8)$$

Where α is the upper limit of the integral, and thus indicates what portion of the function is expected to lie below α . Similarly, the value of $1 - CDF$ is often used, and referred to as ‘‘the P-value’’. By utilizing the cumulative distribution function of this χ^2 distribution, we are able to set our confidence interval by choosing our increments in χ^2 that we include, and their respective values for the original depolarized fraction.

A standard p-value accepted as a general metric of significance is the 2σ result, of 5%. Various tables exist, listing values of χ^2 required, for multiple degrees of freedom, to achieve multiple p-values. However, modern computers allow us to find these with ease, by utilizing “Quantile” functions; effectively, the inverse of the CDF.

Since our scan has 2 independently tunable parameters, namely the depolarization per bounce, along with the spin flipper efficiency, $k = 2$ for our distribution. Likewise, we will want to implement a $CDF = 95\% \leftrightarrow p = 5\%$. It is then easy to solve that we require $\alpha = 5.99$.

With our χ^2 increment of 5.99 now known, we then choose to include all parameter set combinations that produced a χ^2 value at *most* this much above the minimum of the set, which, we recall, has been subtracted off from all elements already. This then provides our set of parameter space which encompasses that 95% confidence interval.

However, we know that there are errors associated with the measured values of the depolarized fraction, that still require correction. With the parameter space which gives agreement in the measurement regime, we then scale each of these members to the accurate depolarized fraction, that existed at the close of the shutter, independent of measurement scheme. It is important to note, that simply picking the minimum and maximum depolarized fractions earlier, BEFORE the scaling, would not necessarily produce the correct range; although this is typically the case, it is also possible that the further deviated from the central allowable area of the distribution, could require larger magnitude corrections, such that the corrected depolarized fraction lay within a smaller envelope of error than may other values. Our uncertainty instead assumes the worst-case scenario, of accept all possible agreements with measurements within the 95% confidence interval, and includes all of these parameter pairs for inclusion in the corrected values, which determine the overall uncertainty.

The simulated values were determined through the same combination of runs and calculation as the experiment; by construction of the ratios of the decay trap monitors to the switcher detector signals for normalization of loading rates and transport efficiencies, along with inclusion of D and P type runs, to normalize the 0th order evolution.

Table 4.2 Table of corrected polarization values, within 68% (1 σ) confidence interval

68% confidence interval	2011 off	2011 on	2012 off	2012 on
P	0.9970	0.9939	0.9979	0.9952
$\Delta P/P$	0.0030	0.0025	0.0015	0.0020

Table 4.3 Table of corrected polarization values and respective uncertainties for each year and loading state, within 95% (2 σ) confidence interval.

95% confidence interval	2011 off	2011 on	2012 off	2012 on
P	0.99774	0.99110	0.99841	0.99380
$\Delta P/P$	0.00426	0.00526	0.00332	0.00426




Figure 4.10 Plot of χ^2 for both load states simultaneously. Contours represent χ^2 of 10, and 100. Region of agreement is widened, as the requirement for a single parameter set to satisfy both spin conditions simultaneously is more rigorous.


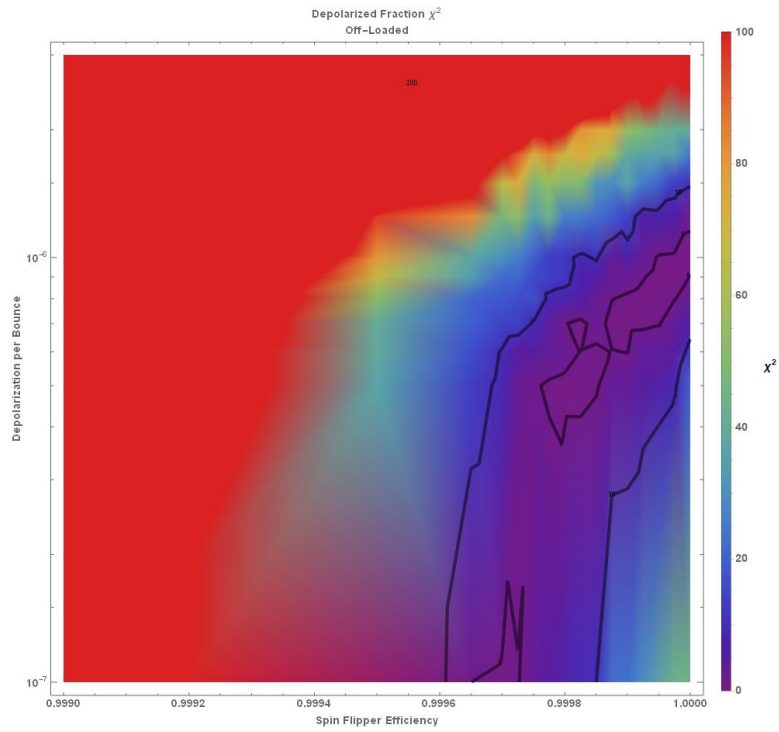


Figure 4.11 Plots of χ^2 values for simulations versus experimental observations, for a grid of spin flipper efficiency and depolarization per bounce, for each of the 2 spin states. Contours represent χ^2 values of 10, and 100. These demonstrate the very narrow available parameter space to reproduce those values determined from experiment, and constrain the uncertainty for the requisite systematic corrections.

Chapter 5

Gravity

5.1 Gravitational Eigenstates

Neutrons' lack of charge is experimentally highly valuable, as it serves to remove a leading systematic in many types of interactions; the electromagnetic influence. Along with the relatively long lifetime, the relatively light mass, and the inability to be electrically polarized (see nEDM), neutrons can be invaluable tools in probing any perturbative forces.

One such interaction being probed is gravity; more specifically, modifications to gravity utilizing 5th force interactions. Various models, such as the Yukawa potential, or axion coupling, predict shifts in energy eigenvalues associated with spatial distribution. One particularly perfect scenario for exploring these shifts are in the eigenstates of the gravitational field itself.

To observe this effect, we first write down the Schrödinger's equation for the simplest gravitational interaction: that of a constant acceleration, in this case from earth's gravitational field at earth's surface;

$$E\psi = -\frac{\hbar^2}{2m}\nabla^2\psi + mgz\psi \quad (5.1)$$

This potential is for the region of space above a floor, which we set to $z=0$ for convenience, and below which is assumed an infinite potential. UCN are then uniquely positioned for these studies, as their ability to be totally internally reflected via the effective Fermi potential corresponds to the floor potential required to produce eigenstates.

By moving all terms to one side, using the substitution $z = \frac{E}{mg} + lx$, and defining a characteristic length

$$l = \sqrt[3]{\frac{\hbar^2}{2m^2g}} = 5.87\mu\text{m} \quad (5.2)$$

this equation can be recognized as Airy's equation, and as such, the solution may simply be written down, $\psi = aAi(x) + bBi(x)$, with a and b representing normalization coefficients. We then impose boundary conditions to construct a wave function which vanishes at $z=0$ and $z \rightarrow \infty$. The $z \rightarrow \infty$ term simply results in setting $b=0$, leaving only Airy functions of the first kind. The $z=0$ condition then requires that $-\frac{E}{mgl}$ correspond to any of the zeroes of the Airy function, yielding a characteristic energy scale of $E_0 = mgl = 0.602 \text{ peV}$, such that the energy eigenstates are of the form $E_n = E_0 * \text{AiryZero}(n)$. Therefore, we find the solution

$$\psi_n(z) = c \cdot Ai\left(\frac{z}{l} - \frac{E_n}{mgl}\right) \quad (5.3)$$

With c the normalization constant such that $\int_0^\infty \psi^* \psi dz = 1$. It so happens in this case that the leading coefficient has a simple functional form, which can be obtained through integration by parts. As ψ is a real number, $\psi^* \psi = \psi^2$, and so

$$\begin{aligned} & \int_0^\infty Ai^2\left(\frac{z}{l} - \frac{E_n}{mgl}\right) dz \\ &= z \cdot Ai^2\left(\frac{z}{l} - \frac{E_n}{mgl}\right) \Big|_0^\infty \\ & - \int_0^\infty 2z \left(\frac{1}{l}\right) Ai\left(\frac{z}{l} - \frac{E_n}{mgl}\right) Ai'\left(\frac{z}{l} - \frac{E_n}{mgl}\right) dz \\ &= z Ai^2\left(\frac{z}{l} - \frac{E_n}{mgl}\right) \Big|_0^\infty - \left(\frac{1}{l}\right) Ai'^2\left(\frac{z}{l} - \frac{E_n}{mgl}\right) \Big|_0^\infty \\ &= \left(\frac{1}{l}\right) Ai'^2\left(\frac{-E_n}{mgl}\right) \end{aligned} \quad (5.4)$$

Where the last step comes from the Airy function being evaluated at an Airy zero.

For scenarios in which there is a ceiling, the boundary conditions will change, requiring introduction of Airy functions of the 2nd kind; however, these typically modify the potentials only slightly. Interestingly, each eigenstate is the same Airy function, simply translated such that one of the Airy zeroes coincides with the floor's position, with magnitudes scaled only for normalization purposes. Additionally, these Airy functions form a complete, orthogonal basis, as required.




Figure 5.1 Reproduced from [68]. The first 5 eigenstates in earth's gravitational well for a neutron. Note that each is the same Airy function, only scaled for normalization, and translated to subsequent Airy Zeroes

The next issue then is to prepare an ensemble in one of these eigenstates. The simplest approach is to include a roughened neutron scatterer as a ceiling for the neutrons to sample during transit. By selecting this separation between floor and ceiling to correspond to the characteristic height of the ground state ($\sim 20\mu\text{m}$), all higher order eigenstates will sample the

roughened layer more often, and be scattered out of the UCN regime, and the experiment, leaving only the $n=1$ ground state solution.




Figure 5.2 Sketch from [68], illustrating an apparatus for selecting $n=1$ eigenstates. Higher order states sample the absorber more completely and are removed from the ensemble rapidly.

5.2 Ramsay's separated oscillatory fields

Once eigenstates are prepared, measurements for neutrons transiting the state selector can be made, as well as while varying the height of the absorber; these measurements were performed extensively by Jenke et. al. in 2011, for the qBounce experiment[69]. However, even more interesting information can be obtained by producing and observing superpositions of these eigenstates.

As discussed in the previous section, the eigenstates produced are indeed orthogonal to one another. Therefore, putting an ensemble into an eigenstate implies they would statically remain in that eigenstate indefinitely. To produce transitions between the states, then, a new coupling must be introduced. To do so, we employ use of Rabi's method, without yet specifying the form of the interaction to take place.

Following the notation of Abele et al [70], we simplify our gravitational problem to that of a 2-state problem, for insight into the transition process. Rabi's method requires only that one couple the eigenstates with a sinusoidal interaction form, which resonates at or near the

frequency corresponding to the energy difference between the 2 states. In the GRANIT experiment, this is done via the magnetic interaction with the neutrons magnetic dipole moment, with the magnetic field varying in the vertical direction[71]. By coupling the 2 states in this way, we obtain the traditional Rabi transition probability,

$$P(t) = \left(\frac{\Omega_R}{\Omega'_R}\right) \sin^2\left(\frac{\Omega'_R}{2} t\right) \quad (5.5)$$

$$\Omega'_R = \sqrt{\Omega_R^2 + (\omega_{pq} - \omega)^2} = \sqrt{\Omega_R^2 + \delta^2}$$

With δ the detuning from resonance.




Figure 5.3 Examples of transition probability widths for Rabi resonances (drift distance $\rightarrow 0$)

However, [70] demonstrate that in lieu of an oscillating magnetic field, one is instead able to take advantage of the floor potential created by the Fermi potential of the material. We will reproduce this short derivation here for clarity. First, we write the Schrödinger's equation,

with the gravitational potential from before, however with a new, oscillating Heaviside function, to represent the oscillation of the floor,

$$-\frac{\hbar^2}{2m} \frac{\partial^2 \psi}{\partial z^2} + mgz\psi + V_0 \Theta(-z + a \sin \omega t)\psi = i\hbar \frac{\partial \psi}{\partial t} \quad (5.6)$$

While the V_0 term is finite (typically of order ~ 150 neV), the vertical energies associated with these states are 5 orders of magnitude smaller (~ 1 peV), and as such the floor potential can be treated as an infinite one. Therefore, as before, we require that the wavefunction vanish at the floor, such that $\psi(z = a \sin \omega t) = 0$. It is also beneficial to change reference frames, to that of the oscillating floor, as will be seen shortly. For now, we take $z' = z - a \sin \omega t$, allowing us to separate the Hamiltonian into both perturbed and unperturbed portions $H = H_0 + W(z', t)$, defining

$$H_0 = -\frac{\hbar^2}{2m} \frac{\partial^2}{\partial z'^2} + mgz' + V_0 \Theta(-z') \quad (5.7)$$

$$W(z', t) = a \left[mg \sin \omega t + i\hbar \omega \cos \omega t \frac{\partial}{\partial z'} \right]$$

Where H_0 is the Hamiltonian with the oscillator at rest, and $W(z', t)$ represents the potential supplied by the oscillating floor. Additionally, we may utilize the fact that the Airy functions produced in this way form a complete basis, such that we may express any solution as a linear combination of these eigenstates,

$$\psi(z, t) = \sum_n C_n(t) e^{-\frac{iE_n t}{\hbar}} \psi_n(z) \quad (5.8)$$

Combining equations (5.8) and (5.6) yields

$$i\hbar \frac{dC_n(t)}{dt} = \sum_n \langle \psi_n | W | \psi_k \rangle \cdot C_k(t) \cdot e^{i\omega_{nk}t} \quad (5.9)$$

$$\langle \psi_n | W | \psi_k \rangle = a \left[mg \delta_{n,k} \sin \omega t + i \hbar \omega \cos \omega t \int_0^\infty \psi_n(z) \frac{d}{dz} \psi_k(z) dz \right]$$

While the sine term is present only when $n=k$, the opposite is true for the cosine term. As such, the integral portion of the 2nd term is sometimes referred to as the “overlap integral”,

$$Q_{n,k} = \int_0^\infty \psi_n(z) \frac{d}{dz} \psi_k(z) dz \quad (5.10)$$

The entries of this matrix represent the likelihood of a transition between 2 eigenstates while undergoing excitation and can be calculated with numerical integration.

Table 5.1 $Q_{n,k}$ values for first 5 eigenstates, from equation (5.10), in m^{-1} . These represent the degree of coupling between eigenstates, induced by an oscillating floor potential, sometimes referred to as the “overlap integrals”

	k=1	k=2	k=3	k=4	k=5
n=1	0	97139	-53411	38210	-30321
n=2	-97139	0	118650	-62984	44080
n=3	53411	-118650	0	134249	-70136
n=4	-38210	62984	-134249	0	146859
n=5	30321	-44080	70136	-146859	0




Figure 5.4 Proposed experimental geometry, reproduced from [70], for a Separated Oscillatory Fields measurement of gravitational eigenstates.




Figure 5.5 Transition probability for monochromatic neutron beam vs. driving frequency. Near perfect transition rates can be achieved, and the high frequency oscillations allow more precise determination of resonance widths.




Figure 5.6 Ramsey transition probability (as above), but with $v^2 dv$ velocity spectrum. Note that maximum transition probability has decreased from unity, and the resonance width is widened, allowing for less accurate determination than with a monochromatic source.

5.3 5th force modifications

With the ability to observe modifications in transition energies, it is then useful to examine likely forms these interactions could take, to begin to predict sensitivity of the apparatus. One common form is the Yukawa interaction, which is expressed as

$$V(r) = -G \frac{m_i \cdot m_j}{r} \left(1 - \alpha e^{-\frac{r}{\lambda}} \right) \quad (5.11)$$

Where α represents the coupling strength, and λ the fundamental length scale of the proposed force. By imposing a physical geometry of an infinite plane below the neutron at $z=0$, and integrating over the parallel components of the wavefunction, such as to separate the z -component(height) from the x - y of the plane's potential, we can obtain a simpler form of

$$\Delta V(z) = 2\pi m_n \rho \alpha \lambda^2 G e^{-|z|/\lambda} \quad (5.12)$$

With ρ the density of the underlying attractor material, assumed to continue infinitely down from the $z=0$ plane. This Yukawa interaction is a form taken from Abelian gauge fields [72], [73], or from extra spatial dimensions [74]–[76].

A similar form is a spin dependent axion coupling, whose form is given in [77] as;

$$V(\vec{r}) = \hbar g_p g_s \frac{(\vec{\sigma} \cdot \vec{n})}{8\pi m c} \left(\frac{1}{\lambda r} + \frac{1}{r^2} \right) e^{-\frac{r}{\lambda}} \quad (5.13)$$

With σ the neutron spin polarization, and \vec{n} the surface normal. Both forms follow similar exponentially damped behavior, with $\frac{1}{r}$ terms thus being enhanced at smaller distance scales.

However, the presence of spin dependence would require the simultaneous use of a spin flipper, possibly similar to those mentioned in previous chapters, for proper control of sensitivity.

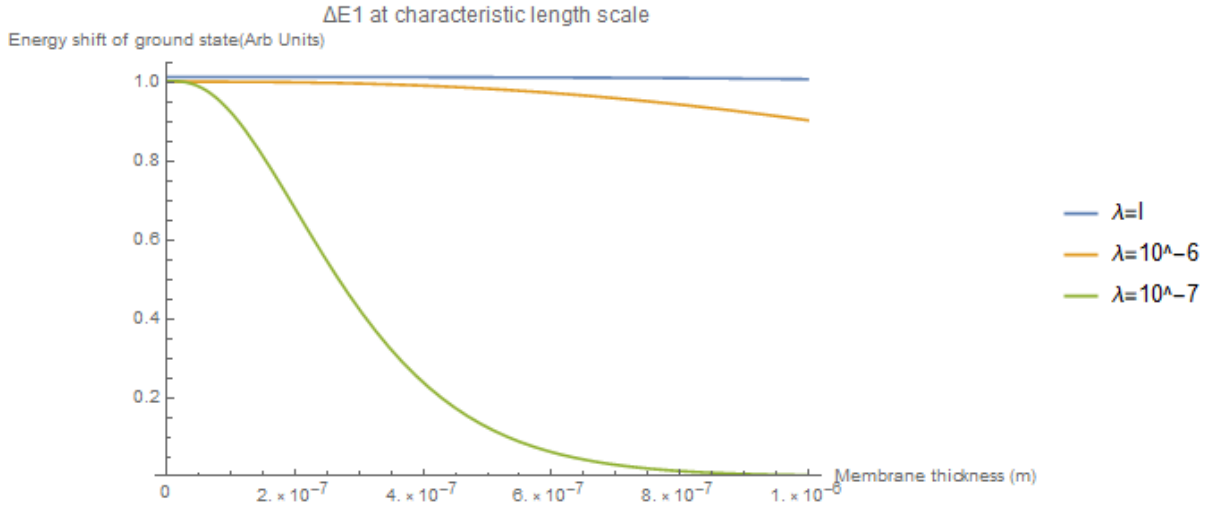


Figure 5.7 Example curves for energy shifts of 1st eigenstate of neutron in gravitational potential, versus thickness of membrane dividing neutrons from attractor material, for varying values of the interaction range λ in a Yukawa-like potential. Values normalized for 0 thickness (optimum sensitivity). “1” corresponds to fundamental length scale of neutrons in earth’s gravitational field, as seen in eq. 5.2

5.4 Microfluidic arrays

To test for interactions at intermediate length scales with sufficient sensitivity for detection, it is desirable to perform a relative measurement shift through introduction of varying conditions, as opposed to an absolute measurement of theoretically predicted values for measurements. To do so, some experiments exploit neutron interferometry, by having coherent neutrons traverse separate paths, one of which may be filled with 1 or multiple gases, of varying densities and pressures, and observing interference fringes develop in response to varying of the density.

Another method for probing these micron-length-scale forces is by inducing resonance frequency shifts through modifying the potential experienced by the neutron. In this case, we would like to adjust the floor potential seen by the neutron. While it is possible to do so by brute force, simply exchanging the entirety of the floor with one of equal flatness and higher/lower density, doing so would be potentially costly. A desirable attractor material must have high density, be machinable to sufficient flatness across large surface areas, exhibit nonmagnetic behavior, and be sufficiently hard as to avoid distortion by resonances with the driving oscillatory field.

As the energy range for these gravitational eigenstates are of order \sim peV, fermi potentials are largely unimportant for the floor material; it need only be that the material is not highly neutron absorbing.

To avoid the catastrophic reconstruction associated with dismantling of an entire experimental apparatus, and the subsequent recalibrations required that would mar the interpretation of a 'control', we developed a microfluidic array through which attractor may be flowed. By utilizing the array, experimental modification can be performed in situ, allowing

continuous running, removing large errors induced by significant changes to the geometry of the system by instead flowing a fluid underneath an exposed top layer.

Due to the length scale being probed, our film need only be sufficiently thin as compared to the order of the range of the interaction; in this case, no thicker than a few microns. Mechanical constraints limit the lower bound to of order $\sim 500\text{nm}$, but this corresponds only to minimal sensitivity changes for the length scale of interest (see **Figure 5.7**).

The most straightforward approach may first seem to be; simply machine a thin enough, flat enough material, and perhaps utilize a bottom layer for support. However, this is beset by many pragmatic issues. First, as the films must be on the order of $\sim 1\mu\text{m}$, no material will have significant resistance to deformation to support itself against the sagging induced by its own weight. Next, the requirement for the film to be locally smooth to $\sim 1\mu\text{m}$, and globally flat to within the same margin over lengths of $\sim 10\text{cm}$, is too stringent for machining technology to achieve. Lastly, placing a thin film on top of a support structure will necessarily create a potential cliff for the neutrons to transit, at the border formed by the edge of the film with the support structure. Even with an inset designed, matching the exact thickness, along with the area and shape, are too stringent of conditions for mechanical production.

For this purpose, silicon was chosen as the substrate upon which to design this array. Silicon was largely chosen for the already prevalent technology and machinery involved with manufacturing desired patterns. Additionally, although silicon is not easily machinable, it has a crucial property to be exploited; the ability to be chemically etched anisotropically. Amusingly, the primary value of silicon in most applications (ie, as a semiconductor) affords no advantage in this application and is of no direct importance.

Multiple methods exist for shaping silicon. Possibly the most direct available would be the process of Deep Reactive Ion Etching (DRIE), in which alternating layers of protective coatings, followed by directed high energy ion exposure are used. In doing so, all surfaces normal to the incident beam are etched more quickly than sidewalls, which are continuously re-coated with the protective gas deposition. However, this method has some issues; sidewalls produced are not smooth, but instead ‘scalped’, as the alternation between beam and protective coating wear into the walls at varying amounts as depth increases. More importantly, however, is the depth control. DRIE works exceptionally well for etching *through* a surface but relies on timing exposure durations properly. With the dynamic conditions associated with etching and re-depositing a substance, run routines can produce varying results from batch to batch. As such, the etch process to form a thin film would instead be a slow and tedious process of guess and check, with thickness variations occurring over the span of the film, and a reliability and reproducibility issue inherent in its production.

Instead, we utilized the process of KOH wet etching. The separate crystal planes within the diamond-lattice shape of the silicon react at different rates, such that certain orientations can be dissolved, while the overall structure remains intact[78]–[80]. These crystallographic planes also impose limitations on the shapes that are able to be etched, and dictate which axes may be used.

Etching a silicon wafer alone would simply create holes in a wafer; this means there is still the requirement of a film. Luckily, chemical vapor deposition (CVD) allows the growth of many different chemical layers on silicon. For our purposes, we utilized a silicon nitride deposition, which has a high Young’s modulus (200-300 GPa), and does not react to KOH. This layer on both sides of the wafer then acts as the etch stop for pattern designs, as well as the thin

film separating the fluidic channels from the neutrons. This film thickness can be controlled by length of exposure during film growth, and typically is of order ~600nm, though no upper bound on the thickness has yet been determined.

The first pattern attempted was a series of 17 elongated parallel sidewalls of thickness ~5mm, spanning the ~75mmx75mm square area in a winding channel configuration(see **Figure 5.8**).The depth is set by the thickness of the wafer itself, and was generally of order ~500 μ m, over a diameter of 15cm. After coating in a Si_xN_y layer, the double side polished wafers were spin coated with a chemical known as photoresist. After being allowed to harden through baking, the wafers were then loaded into a photolithography machine. By inserting a aluminum coated quartz mask between the wafer and the exposure light, the desired pattern can then be selectively exposed to a UV light source, reacting with the photoresist coating. After the exposure, immediate washing in a developing solution occurred. Depending on the type of photoresist, either negative or positive exposure types were used, in each case determining whether the exposed material was washed away by the developing fluid, or hardened.




Figure 5.8 Diagram of early design for winding array(grey), on quartz block(red), with silicon sidewalls(green).

After having the desired pattern outlined by the hardened photoresist, wafers were then exposed to developing solution. After sufficient time in this environment, removal of all top layers occurred. This resulted in a removal of silicon nitride in the developed area, exposing the underlying silicon, while those areas which remained coated in photoresist only had excess thickness of photoresist removed, leaving the underlying nitride layer intact. Finally, a wash in a removal solution was performed, cleaning the remaining photoresist from the wafer, leaving exposed the bare silicon in only the pattern desired, on 1 of the sides of the wafer, while the opposite side remained completely covered in the film of nitride.

After development was complete, the etching process began. Silicon wafers can come in a variety of orientations, but the most widely used and available type is the $\langle 100 \rangle$ plane; the simplest configuration, which in plane carries the remaining 2 preferential orientations at 90 degrees to one another. This means that etched shapes will share a propensity towards creating squares and rectangles in these directions. In our first attempts, nitride layers became fully

exposed, but the elongated sidewalls (silicon fingers) were completely dissolved as well, leaving behind only the nitride coating that was meant to protect them. Additionally, the expected right angles at each corner were instead notched to create an asymmetric octagon (see **Figure 5.9**). This feature led to the discovery that the manufacturers had not followed conventional cleavage plane nomenclature, and that our wafers were in fact 45° misaligned to the desired orientation, resulting in the dissolution of the silicon fingers which were oriented as to be maximally dissolved. Despite this setback, these experiments proved the proof of concept for the selective etching procedure, and confirmed the feasibility of such an atypically large suspended film being produced and maintained, even when lacking the underlying support structures of the winding array pattern.




Figure 5.9 [100] array with array walls etched out. The remaining silicon nitride strands coating the walls can still be seen vertically, and the yellow tinted film is still in-tact.

Because of concerns over nonuniform cohesion forces acting on the membrane during the drying process, a displacement system was implemented such that the KOH was slowly removed while replaced with water, so as to maintain the full immersion of the wafer, and subsequently the water removed and replaced with isopropanol, due to its fast rate of evaporation, water solubility, lower density, and specifically low surface tension, to avoid stressing the membrane. However, further tests showed this procedure to be unnecessary, and that simple removal of the wafers from any solution did not cause membrane rupture either during removal or drying. Afterwards, isopropanol was then used only as a simple rinse, without requiring the wafers to remain immersed throughout the process.

After correcting the orientation issue, sidewalls were able to be produced. However, multiple issues then arose. Most noticeably, the wafer appeared wrinkled in many areas between the channels. This was caused by multiple factors, one of which was the internal stress of the nitride deposition layer. Typical silicon nitride is formed at the stoichiometric ratio of Si_3N_4 , but because of the mismatch between the geometry of these layers compared to the diamond lattice structure of the underlying substrate, an internal compressive stress is present throughout the film, causing buckling, or in cases in which the film was too thick, full rupture. However, this ratio of 3:4 can be modified by altering the ratio of gases flowed during the deposition process, along with other key parameters such as heating and cooling rates, and in doing so, internal film stresses can be modified. In this way, we had our nitride layer specified to an undetermined ratio of Si_xN_y , resulting in a low net *tensile* stress[81]. This pulling force then helped smooth the membrane when fully exposed, much like pulling the sheets tight over a bed to remove wrinkles.




Figure 5.10 Properly etched fluidic array, with compressive stress in silicon nitride film. Waves/ribs can be seen on the film between the side walls. Additionally, the [100] orientation over-etching can be seen most visibly as the small colored rectangles at the end of each silicon finger.

The 2nd issue that became prominent was the structure of the silicon fingers after etching. While the thickness of the walls remained the same at the very top of the array (the side touching the fully intact membrane), they were thinner on the opposite face (the one which etching occurred on). This angle of 54° corresponds to the $\langle 111 \rangle$ crystallographic plane, at which the etch rate reaches a near halt[80]. Due to this fixed angle, then, there exists a minimum width of sidewalls required, depending on the thickness of the wafer, beyond which the fingers would be detached from their membrane, providing no structural support (and causing further strain on the membrane from their own weight/bowing force), and eliminating the feature of a meandering channel, instead being a fully connected layer underneath each wall. Further, this then would

limit the maximum useful surface area of the exposed fluid layer, reducing sensitivity of the experiment.

It was then we turned to a new orientation. Hölke & Henderson report the possibility of high contrast ratio, high depth, vertical sidewall channels being possible to produce through wet silicon etching, similar to the procedure we'd already developed. Instead, however, was the requirement of using a different orientation wafer, $\langle 110 \rangle$ [82]. While $\langle 100 \rangle$ wafers are by far the most common, and $\langle 111 \rangle$ wafers have their place in practical and research settings, $\langle 110 \rangle$ orientation wafers are by far the least used, and thus least commercially available, of standard silicon wafers. This placed restrictions on available substrates, but luckily properties like dopant concentration, and thickness, were manageable for this application.




Figure 5.11 SEM of $[110]$ silicon wafer etching pattern producing deep grooves with high contrast ratio, reproduced from [82].

With the $\langle 110 \rangle$ orientation, the original pattern was reproduced, but with vertical sidewalls. We were then able to move to much finer detailed patterns, with $\sim 10\times$ the number of

silicon fingers, and improve the filling factor from roughly 1:1, to instead ~5:1. This served to increase our sensitivity through a larger portion of the film being exposed, and to reinforce the film more thoroughly with the more numerous, yet smaller sidewalls.

However, the array was also affected by the crystallographic planes' orientation in the wafer. Whereas in $\langle 100 \rangle$ wafers, the planes are situated at 90° to one another, $\langle 110 \rangle$ wafers are arranged 71.9° apart [83]. This then demanded 1 preferential direction for the sidewall orientation, and precluded utilizing horizontal features that may have been helpful. However, due to the high ratio of length to width of the array walls, the pattern was produced to within very similar effect, with the difference that the very ends of each sidewall protrusion were angled at the 71.9° instead of 90° as in the $\langle 100 \rangle$ case.

The other effect of the differing relative orientations of the crystal planes, besides the strange behavior of both convex and concave corners during etching, is for the cleavage process. If one attempts to break crystalline silicon, the result will be many shards of varying sizes, which will also exhibit the underlying symmetry of the crystal. However, if one is to create a preferential cleavage plane, by introducing a small nick along the crystal plane with the use of a diamond-tipped scribe, then the wafer will cleanly break along this axis, into 2 whole pieces. For the $\langle 100 \rangle$ orientation, this then meant that shaping the wafer to the rectangular quartz bonding substrate was as simple as a few careful markings with the scribe, as the $\langle 100 \rangle$ planes are also at 90° to one another. However, as the $\langle 110 \rangle$ cleavage planes are separated by 71.9° instead, this then precluded all 4 sides being cleaved to fit the rectangular shape required for bonding. This then requires an additional step, in which one must etch the rectangular perimeter by some other means (such as via the same wet etching process, DRIE, etc.).

With the array manufacturing process reasonably well established, we then need turn to the issues of flatness, and a mechanism for fluid flow through the channels. Both were approached with the same solution; a supportive 1” thick block of quartz, machined to the required flatness, with 2 fluidic input/output channels bored at a 90° angle from 1 of the sides, upward towards the top face. By having a sufficiently thick block of quartz, we were able to greatly diminish the flexible nature of any thinner substrate, as in the case of the silicon wafer. When measuring the wafer with a Dektak profilometer, significant bowing was observed before etching, and extreme bowing was found after the etching procedure. This remains consistent with the production of a tensile membrane layer on each side, with 1 side having significant portions of the surface removed prior to the etching procedure; thus, with 1 side covered with more tensile material, the wafer tended to bend upwards towards this more fully covered surface.

It is possible to directly bond quartz to silicon, but doing so involves high temperatures and pressures[84], which would likely destroy the delicate membrane. However, this is a permanent effect, and we required a reversible one. Due to the difficulty, time, and cost associated with production of the quartz substrate prisms, each of which would need to reversibly, chemically bond with a removable compound. Also, the procedure for bonding

However, quartz’s chemical composition lends itself to being chemically active with silicon-based polymers as well, yet resilient towards many forms of chemical reactions, and is hard enough to prevent excessive scratching of the surface, which would distort the planar flatness required. This common silicon component allowed bonding via an activated layer of polydimethylsiloxane (PDMS), a widely used synthetic polymer.

In all scenarios, the fused silica wafer was used primarily for reducing expenses associated with the bonding procedure, as well as due to the reusability of the silica wafers. By

desiccating the fused silica wafers in the presence of 100 μ L of 1,1,2,2-fluorooctotrichlorosilane for 2+ hours, a passivation layer was deposited, making the PDMS adhesion very weak, and allowing easy removal of the fused silica wafers after curing was completed. This passivation was easily testable, as the coated surfaces demonstrated a drastic change to being highly hydrophobic, allowing simple contact angle/wetting observations of DI water to confirm proper surface preparation.

PDMS was prepared immediately prior to use, by mixing the elastomer in a 10:1 ratio with a curing agent, and then placing the mixture into a desiccator for ~15 minutes to remove bubbles formed from the mixing. Although complete curing could take up to 24 hours, significant setting occurred within only ~2 hours, and this could be sped up significantly through the addition of heat, i.e., through baking on a hotplate.

Initial tests with the PDMS involved pouring an excess of PDMS onto the quartz substrate, and then applying pressure to a fused silica wafer placed on top. The goal was for the polymer to wick from the center of the block outwards due to capillary action[85], however the plastic did not have sufficient mobility to completely cover the surface. Additionally, the surface produced was highly nonuniform, exhibited strong surface tension effects on edges, and thickness was difficult to control. N-hexane is a known solvent for PDMS[86], and attempts were made to utilize a mixture of this with the prepared PDMS, such as to significantly reduce the surface tension of the PDMS during deposition, and then evaporate out during the curing process[87]. Although this did have the desired effect, the evaporation process resulted in many bubbles throughout the mixture during curing, and the surface produced by this method varied significantly in thickness across the face.




Figure 5.12 Improper bonding of fused silica wafer to quartz substrate via PDMS. The diffraction rings of varying thicknesses can be observed in the center, and the incomplete coverage is observable at the corners. Input ports were required to be blocked, to prevent flow of PDMS into input/output channels, to prevent clogging.

It was believed at the time that proper bonding to PDMS was only possible if it were deposited on a given surface, and allowed to cure while in contact. For the surface covered with the fused silica wafer, the hopes were that the PDMS surface being in contact with the silica until the moment of bonding for the silicon wafer would prevent deactivation. While somewhat true, bonds obtained via direct deposition were quite weak, with bonded wafers being removeable without causing any shattering of the silicon wafer. However, multiple sources reported the ability to activate passivated PDMS via exposure to an oxygen plasma[88][89]. With reactivation a possibility, a new method for developing uniform thickness coatings was made available; spin coating.

PDMS was prepared as above, but deposition occurred directly onto a fused silica wafer. Because of the shape and mass, it was infeasible to perform spin coating of the quartz block.

Care was taken to pour the PDMS solution smoothly onto the wafer, to maintain laminar flow, and prevent introduction of bubbles through pouring. The wafer was then ramped to 1000rpm, and spun for 75s. Directly after spinning, the wafer was placed onto an 80°C hotplate for ~2 hours, and allowed to fully harden. If any bubbles or dust were introduced during the process (as evidenced by highly obvious surface deformation around the contaminant), the wafer could be cleaned prior to curing, and the process repeated. To remove uncured PDMS, the wafer was spun to 3500 rpm, and then N-hexane was applied from the center outwards. After this, a series of rinses with acetone, methanol, and finally isopropanol was performed for final cleaning.

For creating the final bonding step, both the quartz block and the array were descummed in an O₂ plasma asher at 300W for 10 minutes. This high energy oxygen plasma helped remove any organic contaminants immediately preceding the bonding stage. Next, the same quartz block and the fused silica wafer were placed face up in the same plasma asher, and this time exposed at only 60W for 60s. [89] speculate as to the mechanism of this activation, but it is confirmed that attempting to use normal operating power (300W) for any amount of time did not produce the same surface activation required for bonding. Immediately after exposure, both components were placed firmly together, and allowed to seal (as observable via diffraction rings at the bonding interfaces). These were then placed onto a hotplate (wafer side down, to conduct heat more efficiently to the PDMS layer), and allowed to bake at 110°C for 1 hour. The same procedure was used to bond the array to the quartz block, after removing the fused silica wafer by simple peeling. For the heating procedure, however, due to the fragility of the array's membrane, an oven was used in lieu of a hotplate, and baking time was increased to between 2 and 3 hours.

Although multiple quartz blocks were available, it was still highly preferable to reuse each as many times as possible. This then required removal of any deposited PDMS from

previous attempts, as well as recleaning and preparation of the surface. To remove the excess PDMS from the block, as well as any portions of silicon wafer still bonded to the surface, a variety of methods were attempted. N-hexane is known to absorb into cured PDMS, and caused wrinkling/swelling, and so was attempted as a means of disrupting the uniform bonding of the polymer. However, this affected only areas in which poor bonding was achieved, and had no effect on well bonded systems. The only method found to work was to submerge the block into 97% H₂SO₄, and heat to ~160°C. At this temperature, H₂SO₄ behaves as an azeotrope, and can be concentrated no further. Because of this, a steady stream of vapor is emitted by the container, and thus required covering.

Etch rates varied significantly depending on exposed area for the surface. A completely exposed block face, coated in PDMS, was able to be cleaned in a matter of hours, which was also able to be hastened significantly with physical agitation. However, for those cases in which the silicon wafer remained bonded, the exposed layer consisted of the edges of the bonding, both from the outer edge, as well as an inner circumference defined by the (now ruptured) membrane. In these cases, multiple days of cleaning were required, and could be even longer if no physical agitation occurred at all. Finally, physical agitation in-situ was made difficult by the requirement of chemical inactivity with the acid. For removal and placement, tools and stands either made of, or coated with, polytetrafluoroethylene (PTFE, Teflon) were used, rendering them unaffected by the sulfuric acid. Additionally, the softness of this material prevented any significant abrasions of the quartz surfaces during agitation.

Once the PDMS was removed, threaded PTFE rods were used to extract the quartz block, and transfer to a separate bath of H₂SO₄, also held at 160°C. This was to prevent resolidification of the PDMS (which became a foamy top layer to the acid bath) onto the block, and the

temperature was controlled to prevent cracking from thermal shock. After the bath was allowed to cool, the block was transferred to a beaker of deionized (DI) water, and rinsed thoroughly to remove any acid that may have stayed in the fluidic channel portion in particular. Finally, the system was dried with dry N₂ gas, extracted from the boiloff of a liquid nitrogen dewar.

Before all processes, an RCA-1 cleaning process was used, in which a 5:1:1 mixture of water: NH₄OH: H₂O₂ was prepared and heated to 70°C, and either the wafer or the quartz block allowed to remain submerged for between 15-30 minutes. This step not only is standard industry practice, known for removing certain metal and other contaminants (???), but had the added benefit of activating the surfaces, changing them to a highly hydrophilic state. Any material that did not exhibit the characteristically shallow wetting angle for test drops of DI water were recleaned in H₂SO₄ and put through the RCA-1 procedure again.

Profilometry measurements indicate that the wafer's deformation is drastically reduced by bonding to the quartz substrate. Over a span of 5cm, a total variation in height of less than 0.5µm was achieved (see **Figure 5.15**). By increasing the pressure exerted by the tip of the profilometer, we are also able to observe the array pattern bowing under the force, giving us a maximum expected deflection of <1µm (see **Figure 5.14**).




Figure 5.13 Completed microfluidic array with highest contrast ratio. In addition to have significantly more and smaller supporting silicon fingers, the relative area of channels to walls was increased by a factor of 10, permitted by use of the vertical sidewall etching.




Figure 5.14 Contact profilometry of wafer after bonding to quartz block substrate via activated PDMS. The large spike was simply physical debris. Oscillations show the variation across each channel only constituting total variations of $\sim 1\mu\text{m}$, over the span of 22mm.

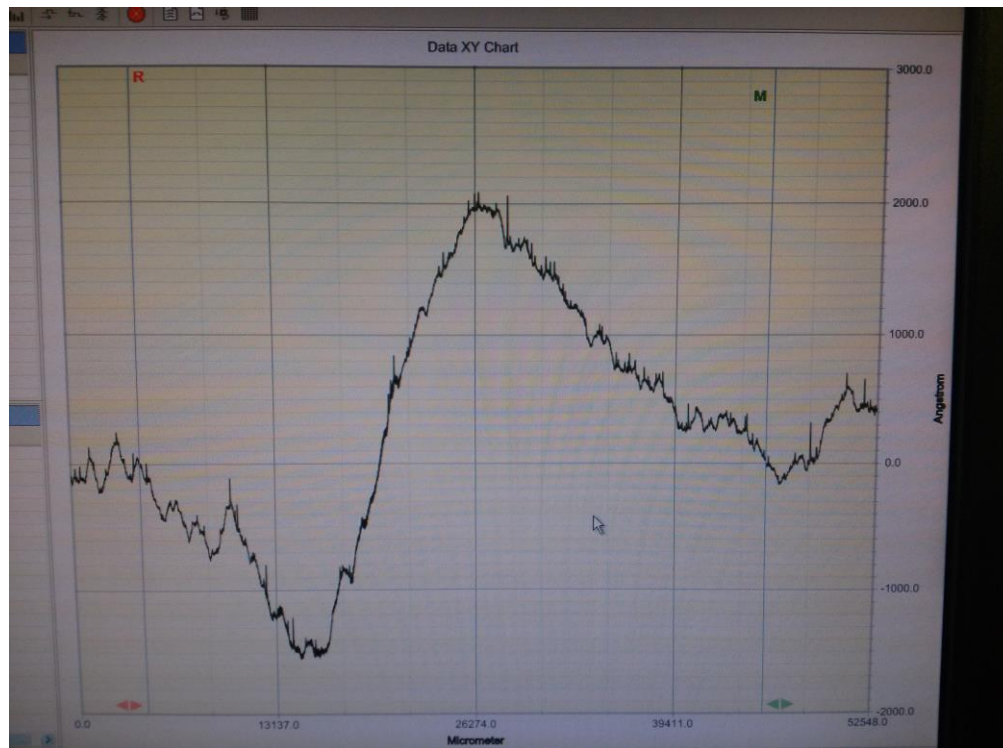


Figure 5.15 Contact profilometry of bonded wafer across flat portion (excluding central array). This demonstrates that the total variation in height over 5cm distance is less than $0.5\mu\text{m}$.




Figure 5.16 Mercury being flowed through completed and bonded array from top to bottom. Mercury was introduced through simple elevation of source input, and can be seen winding halfway across one of the channels from left to right.

5.5 Detector Package

To perform initial tests with the array, a detector assembly was constructed to allow measurement of neutron populations after either prolonged exposure to, or transit across, the array surface. Because of the exceptionally low count rates associated with detection of neutrons prepared into a gravitational eigenstate, reduction of backgrounds was the foremost consideration in developing the setup. Our setup consisted of a thin strip of ZnS, coated in a film of ^{10}B , placed in front of 3 separate PMT's. The ^{10}B served as a neutron absorber, creating ^7Li and ^4He as reaction products. Boron was chosen for multiple advantageous properties; first, that the neutron capture cross section for boron remains quite high in the low energy limit, as opposed to many other nuclear species, whose cross sections fall off with decreasing neutron energy. Secondly,

boron's decay products are stable, allowing for direct observation. These charged products are then able to cause scintillation by interacting with the ZnS strip, and this emitted light is then detected by the PMT array. Finally, boron is quite safe to handle, and simple to deposit, so working with boron is less dangerous and time intensive than other products may be. To remove spurious counts, a triple coincidence between PMT's was required for event triggering.

Likewise, to further reduce backgrounds, the ZnS strip can be made extremely narrow and wide, to match the spatial distribution of the incoming slit of neutrons, prepared in the ground state.

To measure the system, a radioactive sealed source of Gd was used as an alpha-decay source. This was chosen to emulate the alpha particle to be detected from the neutron absorption event.

PMT's were biased to -1300V as an initial test voltage, and individually were swept up in voltage, to observe the reactions of both the separate PMT signals, along with the coincidence requirement. The PMT signals were then directed through an amplifier, then to an ADC, and this output was routed, through a coincidence gate, to an event counter. Amplifiers were set to shaping times of 2 μ s, coarse gains of 50, fine gains of 0.5, with positive input designation.

After sorting out termination and cabling issues, the signal behavior as a function of bias voltage was explored. Individual detectors demonstrated approximately exponential signal growth with increasing voltages (See **Figure 5.17**). However, enforcing triple coincidence between the detectors cancelled electronic noise quite effectively, and produced a wide range of effective voltages at which to operate with consistent results (See **Figure 5.18** and **Figure 5.19**).




Figure 5.17 Example of single detectors count rate increasing with voltage.




Figure 5.18 Count rates of triple coincidence enforced on detectors, in presence of Gd source.




Figure 5.19 Plot of normalized counts rates of 3 detectors, and the signal resulting from their coincidence, versus lower threshold turns on an analog to digital converter. While optimal ranges were available between ~5-15 turns, the triple coincidence signal remains quite consistent even at the lower values, implying better independence from this factor's exact value.

Once reliability was established with the detectors and the imposition of triple coincidence to their signal while in the presence of the Gd source, we proceeded to evaluate background levels (See **Figure 5.20**). This was performed in the approximate area in which we expected the apparatus to be assembled; in the safe work zone, approximately 5m from the biological shielding blocks. These results were confounded by the issue that during experimentation, the accelerator was being ramped down for the year. While some notes were kept tracking the beam current being delivered over the span of each measurement, these values varied, and the reported values were not always accurate. As such, **Figure 5.20** is ordered, from left to right, in order of execution, as beam strength gradually dropped to 0, preventing the final datum from being collected. However, keeping in mind that backgrounds should be falling over time, some useful and interesting information can still be gained from analyzing the figure.

In all cases, we would expect the count rates in which the ZnS strip was present to produce a higher signal, as it simply served as a scintillator directly in front of the detectors, and should emit more from incident particles than would just air. While this is observed in most cases, a notable exception is that in which rubberized boron shielding was used. We believe this was simply an effect of the decreasing beam magnitude over time, which is consistent with the ordering of the 2 measurement periods. Especially interesting, then, is how much higher the count rates were for both cases of the rubberized boron shielding, than with no shielding at all. We suspect this has to do with a secondary shower like effect from the boron, and is likewise seen at similar levels for the polyethylene. However, as the polyethylene shielding run came later, it may be the case that for normalized beam signal, higher count rates would have been observed.

Near the end, lead blocks were used to make a wall between and above the detector package and the direction of the source. Confusingly, even with decreased signal, the geometry with the ZnS inserted produced a *higher* count rate than with no shielding at all, while the absence of the strip produced a lower count rate. It is possible this lower rate would correspond to similar levels as with no shielding, however the increased with the ZnS present suggests a similar effect to that of the other forms of shielding, albeit at possibly lower levels. Finally, the rubberized boron was draped over the lead wall, and a significantly reduced rate was found, however the beam turned off completely during reassembly to test the final configuration without the ZnS present. While it cannot be asserted that this combination was necessarily better at shielding than any other, due to the uncertainty of the beam at the time, it is still consistent with the possibility.

Overall, it appears that nearly all forms of shielding (excepting the possibility of the rubberized boron draped over a lead wall) actually *increased* background. However, we were able to observe in the best-case scenario, a background rate of 0.03Hz. However, due to low beam levels at the time, it is assumed that this will match closely to the situation of the original measurements, with no shielding present, producing backgrounds between 0.03-0.06Hz. This is a sufficiently low rate with which to perform the experiment, with particle rates of ~ 0.1 Hz, though of course further background reduction, or increased particle rates, are always desirable.




Figure 5.20 Background counts versus varying types of shielding, with and without presence of ZnS strip facing detectors. Note: data proceeds in time from left to right, as experiment was conducted during beam shut off.

Chapter 6

Conclusion

6.1 Summary

In this thesis, we have tried to demonstrate multiple continued uses of ultracold neutrons, and the control of their polarization, in investigating interactions both currently within, and potentially beyond, the Standard Model. The utilization of SD_2 as a UCN source has been developed over multiple decades, and the constant improvements in production and reliability have afforded the opportunities to expand the tests able to be performed with the particles.

Because of their basic structure, lack of electric charge, relatively light mass, relatively long lifetimes, and participation in all of the 4 fundamental forces, neutrons provide a window into those forces which remain less well understood, such as the weak and strong forces, can be more easily controlled by the use of the electromagnetic force, and can both be manipulated by, and help probe, gravity at quantum mechanical and classical levels. Examination of free neutron decays have helped to improve determinations of the axial vector to vector coupling constant of the weak force, more accurately determine the still bifurcated results obtained across the world for the current neutron lifetime with unique systematics, and can be selected into individual eigenstates of gravity, to be excited by simple vibrating mirrors, and test for new modifications to the classical Newtonian force from a quantum regime.

Experiments in development for decades, as in UCNA, have concluded, while giving rise to additional experiments, in whose construction and development we assisted, such as measuring the lifetime and electric dipole moment, utilizing the same UCN source. Possible new approaches to separately existing experiments, such as the qBounce experiment and

collaboration at ILL, were created and refined, to improve reliability, as well as facilitate new technologies for controlled interactions and length scales.

It is hoped by the author that these investigations into methods for evaluating depolarization, as well as quantifying and characterizing the primary contributing factors to its development, its evolution, and its measurement, will be helpful and generalizable to further experiments seeking to obtain highly controlled levels of polarization, and constraints on the uncertainty therein. Additionally, the microfluidic arrays herein described have been developed in as a unique solution to address specific systematic experimental concerns, but promise to offer advantages in myriad other applications, following similar techniques as outlined here.

6.2 Future Efforts

After working extensively on various experimental setups, there still remain multiple experiments which I'd liked to be seen to completion, that nevertheless time did not permit. While of course the scope of possible experiments to be performed with the existing apparatuses is vast, and being continued and expanded upon, there are a few in particular that I believe deserve special attention, and require relatively little additional time and effort to reach potentially promising conclusions.

The first, most brute force method is already done; the upgrade of the UCN source to densities at 5x previously available levels immediately grant extensions to achievable statistics, a particularly limiting factor for the depolarization uncertainty currently. So, simply being able to conduct the experiments for an additional cycle, would grant extended reach for determination of A.

Next would be a simple implementation of an absorbing layer to the UCNA decay trap shutter component. By adding even a thin layer of a neutron absorber on the upstream side,

draining times required for emptying the upstream guide components of key backgrounds could be reduced significantly, potentially doubling or greater the decay rate, and likewise reducing the storage time required from the current 30s, to easily achievable levels of roughly 15s. This modification would instantly increase the overall statistics of every depolarization run type, by virtue of the remaining population having simply had less time to decay. More subtly, and perhaps more interestingly for exploration of systematic effects, this reduction of time would then have the cleaning time less than the lifetimes of the populations stored within the decay trap. This then would imply the ability to explore in detail the characteristic non-exponential feeding behavior of the minority states, and allow much finer constraints for determining the magnitude of the depolarization rates from evolution, in turn narrowing the uncertainty window on the complementary correction associated with the spin flipper efficiency.




Figure 6.1 Relative drain curves over time for the existing shutter (left), and a shutter coated in a neutron absorber, such as TPX(right). The greatly reduced time required to clean the system could lead to quicker turnaround for opening the shutter, reducing the degree of evolution that would be permitted.

A few particular depolarization run types are also left inconclusive as a result of sparse statistics, and could benefit significantly from slightly longer beam time allocation, even smaller

allocations of which would be required in light of the source upgrade mentioned above. More time invested in the run types of for depolarization evolution could help resolve the question of both feeding and velocity conditioning over time stored in the trap, however would be extremely effective if measured after the aforementioned shutter upgrade, so as to probe the short time behavior. Additionally, more data points should be collected for the spin flipper power run types, in the region below 5W, concentrated most heavily in the 1W range for better determination of the proper fit parameters, to help extract the implied upper limits to the efficiency. Likewise, the complementary Depolarization Calibration runs would benefit greatly from a finer grain measurement nearer to their maxima than the current 0.1/25 ratio. It may additionally be helpful to incorporate a spectrum scan through the AFP, such as the one performed in 2013, into the analysis as well. This would then help constrain uncertainties regarding flipper on/off transport effects, and conditioning, providing another metric by which to evaluate the robustness of the model.

As for depolarization runs, one final metric of investigation that appears highly promising, is the evaluation of simple loading rates. Whereas all of these run types involve similar timing protocols to those used in standard depolarization measurements, requiring low counts and long run times to determine the wrong spin population at the end of loading and holding phases, loading rates instead provide very high rates, requiring very little time to perform individual measurements. Additionally, if correct, simulations indicate a surprisingly high sensitivity to spin flipper efficiencies, or fractional spin flipper operation times, at the extrema of off and maximally on. This could serve as a very quick gauge of diminishing returns from driving power for the spin flipper, fitting the behavior for intermediate efficiencies could lead to better understanding of the changes in holding lifetimes, and thus sampling rates and population

densities, and being able to map the isomorphic transformation from spin flipper power to effective efficiency would serve to strongly establish the consistency desired from the use of the depolarization calibration runs, while allowing more accurate projections from the maximally powered state, to the limiting case, as an alternative method for extrapolating the spin flipper efficiency.

Additionally, while the parameter space exploration suited this project, it represents a much larger, fundamental problem of high-dimensionality correlated systems. Because of the impossibly vast expanse of the parameter space that could be explored, further investigation into quantifying and modeling the specific coefficients of interdependence, for each of the metrics (lifetimes of varying runtypes, loading rates, etc.). While a solution in general may be impossible, it is nevertheless certain that determining efficient methods for eliminating large areas of phase space, will remain a paramount concern for all computational physics. However, with sufficient simplifications, it is still reasonable to next explore functional modeling of the 2nd order correlation functions, and attempt similar branching parameter space explorations. Upon final reflection, I feel the problem quite lends itself to an investigation from the perspective of “machine learning” algorithms, for possibilities in accelerating these simulations.

Separately, are the uses of the microfluidic arrays that were developed. It goes without saying that their implementation in the intended apparatus is the ideal scenario, however generalizations of the array to alternative uses seem available, and may provide unique advantages. The ability to control membrane thickness, with large area exposure to underlying liquids, able to be flowed in situ, may serve as a helpful tool in evaluating any form of micron to millimeter length scale forces, by providing uniform, controlled surface interactions, while varying parameters only at the length scale desired. While density is the most obvious, there may

very well be other qualities which are desirable to vary, such as magnetic polarizability, surface tension, etc. And, depending on further exploration, it may be possible to flow, and subsequently freeze fluids within the channels, to test structural dependences for hypothetical longer range forces.

REFERENCES

- [1] M. Makela, “Polarized Ultracold Neutrons: their transport in diamond guides and potential to search for physics beyond the standard model,” 2005.
- [2] A. T. Holley, “Ultracold Neutron Polarimetry in a Measurement of the β Asymmetry,” NCSU, 2012.
- [3] B. Vorndick, “Guide Fabrication for UCN Transport,” 2015.
- [4] J. CHADWICK, “Possible Existence of a Neutron,” *Nature*, vol. 129, p. 312, Feb. 1932.
- [5] W. Heisenberg, “Über den Bau der Atomkerne. I,” *Zeitschrift für Phys.*, vol. 77, no. 1, pp. 1–11, Jan. 1932.
- [6] M. Tanabashi *et al.*, “Review of Particle Physics,” *Phys. Rev. D*, vol. 98, no. 3, p. 30001, 2018.
- [7] F. L. Wilson, “Fermi’s Theory of Beta Decay,” *Am. J. Phys.*, vol. 36, no. 12, pp. 1150–1160, 1968.
- [8] I. S. Towner and J. C. Hardy, “Comparative tests of isospin-symmetry-breaking corrections to superallowed $0^+ \rightarrow 0^+$ nuclear β decay,” *Rev. C*, vol. 82, p. 65501, 2010.
- [9] C. Quigg, “The Electroweak Theory,” *Electroweak Theory*, pp. 1–65, Apr. 2002.
- [10] “File:Standard Model of Elementary Particles.svg - Wikimedia Commons.” [Online]. Available: https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg. [Accessed: 14-Nov-2018].
- [11] T. D. Lee, “New insights to old problems,” *Chinese Phys.*, vol. 15, no. 6, pp. 1125–1133, 2006.
- [12] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, “Experimental test of parity conservation in beta decay [5],” *Phys. Rev.*, vol. 105, no. 4, pp. 1413–1415, 1957.
- [13] M. Goldhaber, L. Grodzins, and A. W. Sunyar, “Helicity of neutrinos [7],” *Phys. Rev.*, vol. 109, no. 3, pp. 1015–1017, 1958.
- [14] S. Gardner and B. Plaster, “Framework for maximum likelihood analysis of neutron β decay observables to resolve the limits of the $V - A$ law,” *Phys. Rev. C*, vol. 87, p. 65504, 2013.
- [15] M. González-Alonso, O. Naviliat-Cuncic, and N. Severijns, “New physics searches in nuclear and neutron β decay,” *Prog. Part. Nucl. Phys.*, vol. 104, pp. 165–223, 2019.
- [16] T. M. Ito and G. Prézeau, “Neutrino Mass Constraints on Decay,” 2005.
- [17] J. D. Jackson, S. B. Treiman, and H. W. Wyld, “Possible Tests of Time Reversal Invariance in Beta Decay,” *Phys. Rev.*, vol. 106, no. 3, pp. 517–521, 1957.
- [18] B. Plaster *et al.*, “Measurement of the neutron β -asymmetry parameter A_{β} with ultracold neutrons,” *Phys. Rev. C*, vol. 86, no. 5, p. 055501, Nov. 2012.
- [19] G. J. Mathews, T. Kajino, and T. Shima, “Big bang nucleosynthesis with a new neutron

- lifetime.”
- [20] Y. I. Izotov and T. X. Thuan, “THE PRIMORDIAL ABUNDANCE OF 4 He : EVIDENCE FOR NON-STANDARD BIG BANG NUCLEOSYNTHESIS,” *Astrophys. J. Lett.*, vol. 710, pp. 67–71, 2010.
 - [21] S. Burles, K. M. Nollett, J. W. Truran, and M. S. Turner, “Sharpening the Predictions of Big-Bang Nucleosynthesis,” 1999.
 - [22] 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,” 1996.
 - [23] R. W. Pattie *et al.*, “Measurement of the neutron lifetime using a magneto-gravitational trap and in situ detection.,” *Science*, vol. 360, no. 6389, pp. 627–632, May 2018.
 - [24] V. I. Morozov, “Neutron lifetime determination by ultracold neutron storage,” *Nucl. Inst. Methods Phys. Res. A*, vol. 284, no. 1, pp. 108–110, Nov. 1989.
 - [25] A. K. Fomin and A. P. Serebrov, “Detailed analysis and Monte Carlo simulation of the neutron lifetime experiment,” *JETP Lett.*, vol. 92, no. 1, pp. 40–45, 2010.
 - [26] A. Serebrov *et al.*, “Measurement of the neutron lifetime using a gravitational trap and a low-temperature Fomblin coating,” *Phys. Lett. B*, vol. 605, pp. 72–78, 2005.
 - [27] W. Paul, F. Anton, L. Paul, S. Paul, and W. Mampe, “Measurement of the neutron lifetime in a magnetic storage ring*,” 1989.
 - [28] J. S. Nico *et al.*, “Measurement of the neutron lifetime by counting trapped protons in a cold neutron beam,” *Phys. Rev. C*, vol. 71, p. 55502, 2005.
 - [29] A. T. Yue *et al.*, “Improved Determination of the Neutron Lifetime.”
 - [30] W. J. Marciano and A. Sirlin, “Improved Calculation of Electroweak Radiative Corrections and the Value of V_{ud} ,” *Phys. Rev. Lett.*, vol. 96, no. 3, p. 032002, Jan. 2006.
 - [31] S. Ando *et al.*, “Neutron beta-decay in effective field theory,” *Phys. Lett. Sect. B Nucl. Elem. Part. High-Energy Phys.*, vol. 595, no. 1–4, pp. 250–259, Aug. 2004.
 - [32] J. C. Hardy and I. S. Towner, “Superaligned $0^+ \rightarrow 0^+$ nuclear β decays: A new survey with precision tests of the conserved vector current hypothesis and the standard model,” *Phys. Rev. C - Nucl. Phys.*, vol. 79, no. 5, pp. 1–29, 2009.
 - [33] V. Tishchenko *et al.*, “Detailed report of the MuLan measurement of the positive muon lifetime and determination of the Fermi constant,” *Phys. Rev. D - Part. Fields, Gravit. Cosmol.*, vol. 87, no. 5, pp. 1–30, 2013.
 - [34] H. Abele, “The neutron. Its properties and basic interactions,” *Prog. Part. Nucl. Phys.*, vol. 60, no. 1, pp. 1–81, 2008.
 - [35] M. A.-P. Brown *et al.*, “New result for the neutron β -asymmetry parameter A_0 from UCNA,” *Phys. Rev. C*, vol. 97, no. 3, p. 035505, Mar. 2018.
 - [36] E. DI Amai and E. Fermi, “On the Absorption and the Diffusion of Slow Neutrons,” 1936.
 - [37] R. Golub, D. Richardson, and S. K. Lamoreaux, *Ultra-cold neutrons*. New York, NY: Adam Hilger, 1991.
 - [38] J. Benlliure, “Spallation reactions in applied and fundamental research,” *Lect. Notes*

- Phys.*, vol. 700, pp. 191–238, 2006.
- [39] A. Steyerl *et al.*, “A new source of cold and ultracold neutrons,” *Phys. Lett. A*, vol. 116, no. 7, pp. 347–352, 1986.
- [40] A. I. Kilvington, R. Golub, W. Mampe, and P. Ageron, “Scattering of ultra-cold neutrons (UCN) by superfluid helium at temperatures around 1 K,” *Phys. Lett. A*, vol. 125, no. 8, pp. 416–420, Nov. 1987.
- [41] R. Golub, C. Jewell, P. Ageron, W. Mampe, B. Heckel, and I. Kilvington, “Operation of a superthermal ultra-cold neutron source and the storage of ultra-cold neutrons in superfluid Helium⁴,” *Zeitschrift für Phys. B Condens. Matter*, vol. 51, no. 3, pp. 187–193, 1983.
- [42] R. Golub and J. M. Pendlebury, “The interaction of Ultra-Cold Neutrons (UCN) with liquid helium and a superthermal UCN source,” *Phys. Lett. A*, vol. 62, no. 5, pp. 337–339, 1977.
- [43] H. Yoshiki *et al.*, “Observation of ultracold-neutron production by 9- cold neutrons in superfluid helium,” *Phys. Rev. Lett.*, vol. 68, no. 9, pp. 1323–1326, 1992.
- [44] Y. Masuda *et al.*, “Spallation Ultracold-Neutron Production in Superfluid Helium,” 2002.
- [45] O. Zimmer *et al.*, “Superfluid-Helium Converter for Accumulation and Extraction of Ultracold Neutrons,” 2007.
- [46] P. Schmidt-Wellenburg, K. H. Andersen, and O. Zimmer, “Ultra cold neutron production by multiphonon processes in superfluid helium under pressure,” *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.*, vol. 611, no. 2–3, pp. 259–262, Dec. 2009.
- [47] O. Zimmer *et al.*, “Ultracold neutrons extracted from a superfluid-helium converter coated with fluorinated grease,” *Eur. Phys. J. C*, vol. 67, pp. 589–599, 2010.
- [48] E. Korobkina, R. Golub, B. W. Wehring, and A. R. Young, “Production of UCN by downscattering in superfluid He₄,” *Phys. Lett. Sect. A Gen. At. Solid State Phys.*, vol. 301, no. 5–6, pp. 462–469, Sep. 2002.
- [49] R. E. Hill *et al.*, “Performance of the prototype LANL solid deuterium ultra-cold neutron source,” *Nucl. Instruments Methods Phys. Res. Sect. aAccelerators Spectrometers Detect. Assoc. Equip.*, vol. 440, no. 3, pp. 674–681, 2000.
- [50] A. Frei *et al.*, “First production of ultracold neutrons with a solid deuterium source at the pulsed reactor TRIGA Mainz ★,” *Regul. Artic. Phys. Eur. Phys. J. A*, vol. 34, pp. 119–127, 2007.
- [51] A. Saunders *et al.*, “Demonstration of a solid deuterium source of ultra-cold neutrons,” *Phys. Lett. Sect. B Nucl. Elem. Part. High-Energy Phys.*, vol. 593, no. 1–4, pp. 55–60, Jul. 2004.
- [52] P. W. Lisowski and K. F. Schoenberg, “The Los Alamos Neutron Science Center,” *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.*, vol. 562, no. 2, pp. 910–914, 2006.
- [53] C. L. Morris *et al.*, “Measurements of Ultracold-Neutron Lifetimes in Solid Deuterium.”
- [54] A. Saunders *et al.*, “Performance of the Los Alamos National Laboratory spallation-driven solid-deuterium ultra-cold neutron source,” *Rev. Sci. Instrum.*, vol. 84, no. 1, Jan. 2013.

- [55] T. M. Ito *et al.*, “Performance of the upgraded ultracold neutron source at Los Alamos National Laboratory and its implication for a possible neutron electric dipole moment experiment,” vol. 012501, pp. 1–6, 2018.
- [56] A. Frei *et al.*, “Understanding of ultra-cold-neutron production in solid deuterium,” *Epl*, vol. 92, no. 6, 2010.
- [57] C. Liu, “Cold Neutron and Ultracold Neutron Sources,” *Physics (College. Park. Md.)*, no. June, 2009.
- [58] C.-Y. Liu, S. K. Lamoreaux, a. Saunders, D. Smith, and a. R. Young, “An apparatus to control and monitor the para-D2 concentration in a solid deuterium, superthermal source of ultra-cold neutrons,” *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.*, vol. 508, no. 3, pp. 257–267, Aug. 2003.
- [59] C.-Y. Liu, A. R. Young, and S. K. Lamoreaux, “Ultracold neutron upscattering rates in a molecular deuterium crystal,” 2000.
- [60] C. Liu, “A Superthermal Ultra-Cold Neutron Source,” 2002.
- [61] A. Anghel *et al.*, “Solid deuterium surface degradation at ultracold neutron sources.”
- [62] F. Atchison *et al.*, “Diamondlike carbon can replace beryllium in physics with ultracold neutrons,” *Phys. Lett. Sect. B Nucl. Elem. Part. High-Energy Phys.*, vol. 642, no. 1–2, pp. 24–27, Nov. 2006.
- [63] A. T. Holley *et al.*, “A high-field adiabatic fast passage ultracold neutron spin flipper for the UCNA experiment.,” *Rev. Sci. Instrum.*, vol. 83, no. 7, p. 073505, Jul. 2012.
- [64] Schaerpf, “PROPERTIES OF BEAM BENDER TYPE NEUTRON POLARIZERS USING SUPERMIRRORS.”
- [65] R. T. Robiscoe, “A Spin Flip Problem,” *Am. J. Phys.*, vol. 39, no. 2, pp. 146–150, 1971.
- [66] A. Abragam, *The principles of nuclear magnetism*. Oxford : Clarendon Press, 1961., 1961.
- [67] D. Mund *et al.*, “Determination of the weak axial vector coupling $\lambda=g A/gV$ from a measurement of the β -asymmetry parameter A in neutron beta decay,” *Phys. Rev. Lett.*, vol. 110, no. 17, pp. 1–5, 2013.
- [68] H. Abele and H. Leeb, “Gravitation and quantum interference experiments with neutrons,” *New J. Phys.*, vol. 14, no. 5, p. 055010, May 2012.
- [69] T. Jenke, P. Geltenbort, H. Lemmel, and H. Abele, “Realization of a gravity-resonance-spectroscopy technique,” *Nat. Phys.*, vol. 7, no. 6, pp. 468–472, Apr. 2011.
- [70] H. Abele, T. Jenke, H. Leeb, and J. Schmiedmayer, “Ramsey’s method of separated oscillating fields and its application to gravitationally induced quantum phase shifts,” *Phys. Rev. D*, vol. 81, no. 6, p. 065019, Mar. 2010.
- [71] M. Kreuz *et al.*, “A method to measure the resonance transitions between the gravitationally bound quantum states of neutrons in the GRANIT spectrometer,” *Nucl. Instruments Methods Phys. Res. Sect. A Accel. Spectrometers, Detect. Assoc. Equip.*, vol. 611, no. 2–3, pp. 326–330, 2009.
- [72] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, “Phenomenology, astrophysics, and cosmology of theories with submillimeter dimensions and TeV scale quantum gravity,”

1999.

- [73] I. Antoniadis, N. Arkani-Hamed, S. Dimopoulos, and G. Dvali, “New dimensions at a millimeter to a fermi and superstrings at a TeV,” 1998.
- [74] V. V Nesvizhevsky and K. V Protasov, “Classical and Quantum Gravity Constraints on non-Newtonian gravity from the experiment on neutron quantum states in the earth’s gravitational field Related content Constraints on non-Newtonian gravity from the experiment on neutron quantum states in the earth’s gravitational field,” *Class. Quantum Grav*, vol. 21, pp. 4557–4566, 2004.
- [75] A. Westphal, H. Abele, and S. Baeßler, “Analytically derived limits on short-range fifth forces from quantum states of neutrons in the Earth’s gravitational field,” 2007.
- [76] P. Callin and C. P. Burgess, “Deviations From Newton’s Law in Supersymmetric Large Extra Dimensions,” 2005.
- [77] J. E. Moody’ and F. Wilczek, “PHYSICAL REVIEW D VOLUME.”
- [78] D. L. Kendall, “Vertical Etching of Silicon at Very High Aspect Ratios,” *Ann. Rev. Mater. Sci.*, vol. 9, pp. 373–403, 1979.
- [79] P. M. Zavracky, T. Earles, N. L. Pokrovskiy, J. A. Green, and B. E. Burns, “Fabrication of vertical sidewalls by anisotropic etching of silicon (100) wafers,” *J. Electrochem. Soc.*, vol. 141, no. 11, pp. 3182–3188, 1994.
- [80] P. Oliver and H. B. Harrison, “Anisotropic etching of {100} and {110} planes in (100) silicon,” *J. Micromechanics Microengineering*, vol. 11, no. 3, p. 217, 2001.
- [81] A. S. Rollier, B. Legrand, D. Deresmes, M. Lagouge, D. Collard, and L. Buchailot, “Tensile stress determination in silicon nitride membrane by afm characterization,” *Dig. Tech. Pap. - Int. Conf. Solid State Sensors Actuators Microsystems, TRANSDUCERS ’05*, vol. 1, pp. 828–831, 2005.
- [82] A. Holke and H. T. Henderson, “Ultra-deep anisotropic etching of (110) silicon,” *J. Micromech. Microeng.*, vol. 9, pp. 51–57, 1999.
- [83] W.-T. C. Chien, C. O. Chang, P.-C. Chen, and C.-S. Chou, “On the Determination of Si(110) Convex Corner Undercut Planes by the Zoning Model,” *J. Electrochem. Soc.*, vol. 155, no. 3, p. D209, 2008.
- [84] H. K. Jung, Y. S. Hwang, I. J. Hyeon, Y. K. Kim, and C. W. Baek, “Silicon/quartz bonding and quartz deep RIE for the fabrication of quartz resonator structures,” *3rd IEEE Int. Conf. Nano/Micro Eng. Mol. Syst. NEMS*, pp. 1172–1176, 2008.
- [85] A. Y. N. Sofla and C. Martin, “A vapor-assisted method for adhering polydimethylsiloxane and glass,” *Lab Chip*, vol. 10, no. 2, pp. 250–253, 2010.
- [86] J. N. Lee, C. Park, and G. M. Whitesides, “Solvent Compatibility of Poly (dimethylsiloxane) -Based Microfluidic Devices,” vol. 75, no. 23, pp. 6544–6554, 2003.
- [87] A. L. Thangawng, R. S. Ruoff, M. A. Swartz, and M. R. Glucksberg, “An ultra-thin PDMS membrane as a bio/micro-nano interface: Fabrication and characterization,” *Biomed. Microdevices*, vol. 9, no. 4, pp. 587–595, 2007.
- [88] K. C. Tang *et al.*, “Evaluation of bonding between oxygen plasma treated polydimethyl siloxane and passivated silicon,” *J. Phys. Conf. Ser.*, vol. 34, no. 1, pp. 155–161, 2006.

- [89] S. H. Tan, N. T. Nguyen, Y. C. Chua, and T. G. Kang, "Oxygen plasma treatment for reducing hydrophobicity of a sealed polydimethylsiloxane microchannel," *Biomicrofluidics*, vol. 4, no. 3, pp. 1–9, 2010.
- [90] R. R. Mammei, "Thin Films for the Transport of Polarized Ultracold Neutrons for Fundamental Symmetry Study," 2010.

APPENDICES

Appendix A

Procedure for Microfluidic Array Production

Block Processes:

- 0) Block (and Wafer) PDMS/contaminant removal
- 1) Hydrophilic preparation of block and wafer
- 2) Silanization of wafer
- 3) PDMS coating
- 4) PDMS preparation

Materials needed:

Quartz 4" square block, 6" diameter fused silica wafer, All PPE, training, and nearby associate for concentrated acid handling. **BE EXTREMELY CAREFUL.**

Glassware: 2 hotplates, 2 glass thermometers, 2 watchglasses (for covering beakers), 3 1500mL beakers, 2 8" diameter Crystallizers.

Chemicals: Waste disposal bottle(glass) for concentrated H_2SO_4 , 2L 97% H_2SO_4 , Baking soda (acid neutralization), 200 mL 30% H_2O_2 , 200mL NH_4OH , 200 μ L FOTS (flurooctatrchlorosilane), 1g Elastomer, 200mg Curing agent.

Assorted equipment: 2 10-32 threaded rods (PTFE preferable, stainless steel acceptable in short term), 2 large KimWipes, 2 >6" dessicators, PTFE block as wafer holder, P1000 micropipette and tips, 2 plastic syringes(o-ring, compression fitting, 10-32 pink screw, possibly Teflon tape), Pitri dish.

Steps:

Procedure 0: RCA1 clean and PDMS removal

- 1) Place block in beaker.
- 2) Add water to rinsing beaker, keep nearby the experiment.
- 3) Pour out 800 mL, 97% H₂SO₄ into each of two 1500 mL beakers . Insert thermometer, cover beakers with watchglasses.
- 4) Set hotplates to setting 7, allow baths to rise to ~100 C (roughly 30 mins). Increase setting to 8, and let rise to ~130 C. Maximal suggested temperature 160C, ideally no need to go this high.
- 5) Let sit until PDMS dissolved from block. Film of PDMS should float on acid

NOTE: Vaporized acid creeps up thermometer, so acid will be further up than contact point with the acid.

- 6) Remove both watchglasses, insert threaded rods into block. Lift block from solution, tilt to one side, and allow drops to become spaced by roughly 1-2 seconds. Straighten block, and make transfer, being very careful not to spill. Remove rods, to allow channels to clear. If spill occurs, use baking soda, call somebody. **BE EXTREMELY CAREFUL.**
- 7) Turn off both hotplates and allow to cool to room temperature.
- 8) Pour first solution of acid into waste container. Rinse thoroughly. Then pour 800 mL of DI water into beaker.
- 9) As with step 4, transfer block to room temperature water bath. We will keep the 2nd acid bath, which was used as a rinse step, for the first step of the next procedure.
- 10) Remove from water, rinse in DI water, dry with N₂.

If wafer needs cleaning too: Use crystallizers instead of beakers, metal tongs instead of rods, and use recycled 2nd step acid from previous step for both crystallizers. No cooldown step (7) is required, since the thinness of the wafer prevents cracking from thermal shock.

Procedure 1: Block+Wafer RCA2 clean and Hydrophilic preparation

- 1) Place block in beaker and wafer in crystallizer.
- 2) Use 5:1:1 ratio mixture of H₂O:NH₄OH:H₂O₂. For beaker, this should be 600mL:120mL:120mL, and for crystallizer, 300:60:60. Pour over block and wafer.
- 3) Add thermometers and watchglass. Heat both to 70C. Leave at 70 for 15-30 minutes.
- 4) Turn off heat. Allow block to cool to room temperature (no need for wafer).
- 5) Transfer to DI water rinse bath.
- 6) DI Water rinse, N₂ dry. Check wetting angle and spread for measure of hydrophilicity.

Procedure 2: Wafer silanization

- 1) Place wafer in Teflon holding block, and both in dessicator.
- 2) Crumple 2 large KimWipes, place on either side of wafer in dessicator.
- 3) Pipette 100 microL FOTS onto each of the KimWipes.
- 4) Turn on vacuum, let sit at minimum 2 hours.
- 5) After silane deposition, wafer should be hydrophobic; test by observing wetting angle.

Procedure 3: Wafer-PDMS deposition

- 1) Place silica wafer (or 2) in dessicator, with coating side totally exposed to air. Put 200 microliters of FOTS onto kimwipe for each wafer, leave for ~24 hours.
- 2) Mix 7g PDMS 1:10 curing agent for 0.7g curing(I used aluminum trays, since the plastic ones reacted when I tried hexane dilutions, but this may no longer be the best choice), mix very well, and dessicate for half hour, with occasional repressurization to pop bubbles. There must be NO bubbles. You can gently wick away few remaining bubbles with glass stirring rod, but be careful, and always end with dessication to be safe.

- 3) Pour entirety of PDMS mixture onto the center of silanated silica wafer; pour very close to the wafer, to achieve laminar flow, to avoid bubbles, which are the wafer killer.
- 4) Spin coat with ramp rate of 150, up to spin speed of 1000 rpm, for 75s total time.
- 5) If there are ANY bubbles in the region of interest, clean with spin speed of 3500rpm, A) n-hexane, removes bulk of PDMS, b) Acetone c) Methanol d) Isopropanol. If wafer was dirty to begin with, only the last 3 are necessary.
- 6) If wafer coat went well, bake at 80C for 1.5-2 hours.

Procedure 4: Block-PDMS preparation

- 1) Descum quartz block in O₂ plasma asher at 300W for ~10mins first. Rinse in DI water.
- 2) Place both PDMS wafer and quartz in O₂ plasma asher for 60s at 60W. (Low power, and short time, required for surface activation; longer times and higher powers eliminate this effect).
- 3) Remove, place PDMS wafer on table, then gently lay quartz on it; Wet back of quartz with water to see bonding, after rings are done moving, apply heavy pressure in problem areas to maximize bonding surfaces, working outwards from the center.
- 4) ******(May actually need higher temperatures or longer times, perhaps 150C for ~1hr) Bake wafer side down at 110C for ~20 mins, ideally with additional weight on top.
- 5) Cool for ~10 mins, until holdable.

Array Processes:

- 1) Turn Silicon Wafer into Microchannel image
- 2) Etch Wafer

3) Siphon

Materials Needed:

Silicon wafer, <110> , TTV <2 micron, N Doped (Phosphorous), thickness currently 600 micron, prime grade, DSP, diameter 150mm, well controlled PFL.

Steps:

Procedure 1: (Tensile) Nitride Deposition

- 1) RCA2 (Base step clean with 1:1:5 H₂O₂:NH₄OH:H₂O).
- 2) BOE (Buffered oxide etch) (HF) for ~ 1 minute.
- 3) 2.5(~500nm) – 4(~800nm) hour nitride deposition. 22 sccm Dichlorosilane to 100 sccm NH₃ (1:5) for tensile. At 17 sccm DCS (1:6), get compressive stress, which causes wrinkles, which are highly undesirable. For 4 hour nitride, process took 6 hours 17 minutes, to include warm up time.
- 4) Coat one side of silicon wafer in photoresist. Used S1813 positive photoresist (+ means the UV cured area is removed, so the ion bombardment will eat nitride in exposed area, while it eats just photoresist in nonexposed).
- 5) Spin coat at 1000 RPM for 45 seconds. Be careful to use correct size chuck. Bake on hotplate at 120C for only ~2 minutes, since heat conduction much better than oven.
- 6) Look at PR coating under microscope. If polarizer available, can tell difference between dust and bubbles easily. Dust: asymmetric, small, dark. Bubbles: Exhibit dipole pattern of birefringence, colorful, larger, multiple different images in different lightings. Bare wafer: Diffraction rings, grey in color

Procedure 2: Photolithography

- 1) Used Karl Suss TSA MA-6 mask aligner. Place mask in position, assure vacuum is good. Use 50 micron alignment gap, and 60 second exposure time. Make sure to align as well as possible to alignment flat.
- 2) Possible to use microprime primer p-20. This prevents developing solution from creeping under photoresist, and blurring resolution of development. Should only be necessary for features <10microns.
- 3) Expose aligned mask for 60 seconds in UV.
- 4) Bring exposed mask rather quickly to developing solution. Used MF-319, important part being TMAH based solution. Leave for ~2 minutes, but will notice coloration of channels changing until finished. If not complete, re-expose and develop again. If very thick photoresist, may need to wait up to 10 minutes before developing, for photoacid to effectively be removed. Don't wait too long, or features will be smeared as acid diffuses.
- 5) Reactive Ion Etch nitride from developed wafers. For 2.5 hour deposition, do 20 minutes etching, for 4 hours, do 35 minutes. RIE gas ratio formula 45 sccm CHF₃ : 5 sccm O₂, 250 Watts, 20 millitorr.

Procedure 3: Array wet-etch

- 1) Get or make 30% KOH solution by volume (8 M). Using current 87.5% pellets, and making sure to account for added volume, ratio has been 1000 mL water to 384g pellets for the 1500 mL beaker. Make sure to use magnetic stirrer until completely clear. Note: Dissociation of salt increases heat dramatically, be slow.
- 2) Rinse wafer in 1165 (NMP) bath for ~30 minutes, agitating every few minutes for removal of photoresist.
- 3) Place Teflon holding block in KOH beaker, then place wafer into holder.

4) Heat to ~70 C. Bubbling rate is extremely temperature dependent, so after time you can tell the temperature by amount of bubbles

5) Let etch. For 600 microns, appears to take ~8 hours if aggressive for duration of etch.

Note signs of etch nearing completion; A) Ridges have visible thickness, B) Windows on periphery typically clear first. C) Cleared membrane is completely see through, and yellow in hue. It will spread from the first breakthrough across whole membrane in roughly 10 minutes or less.

Procedure 4: Array-Block bonding

1) Cut wide margin around block with razor on PDMS of silica wafer, such that large flaps are grippable, then peel off outer layer of PDMS on wafer to make it easier to grab flaps.

2) Block input channels of block with o-ring and teflon screws, to prevent air from being allowed to flow in during wafer removal.

3) Gently roll edge of PDMS until you can fold each of the PDMS 'drapes' completely around the quartz, and try to get them to stick mildly and lay flat against the walls; we need to effectively be pulling the PDMS down and out onto the quartz while we remove wafer.

4) Gently remove wafer from corners and edges; we are slowly peeling off the PDMS, which can also be torn from the quartz if not sufficiently delicate. The idea is to have minimal separation, to avoid air bubbles.

5) Once wafer is removed, puncture input ports on top through the PDMS, to permit flow later.

6) Remove input channel blocks.

7) Cut off excess drapes of PDMS with razor.

8) Insert both microfluidic array, bonding side up, and quartz block into O₂ asher, and run 60s at 60W again to activate surfaces for bonding.

- 9) Gently place array onto quartz, aligning input channels. Apply no pressure yet.
- 10) Wet the bottom of the quartz (you'll have to look up at it), and then see if any seal has begun.
If not, press gently at a corner to begin. Wait until seal has stopped moving. Then, apply slow and continuous pressure in unsealed areas, to create the necessary seal around the array
(Once an area is sealed for ~5 seconds, it should not become unsealed; if it does, activation was not successful, and it will not completely seal). Continue to apply pressure to bond as much as possible, being very smooth and careful to not destroy array; even bending the unbonded edges of the wafer too far will rupture the membrane, so stay on the areas with quartz backing.
- 11) Bake at 110C for ~2-3 hours in Blue M Oven. Then, turn off oven and let cool down slowly; this prevents rapid gas movement from the excessively quick cooling, which may rupture the membrane.

Appendix B

General instructions for use of parallelized transport code.

Getting Started

This is a quick, step-by-step guide for how one should go about configuring the code, running it, and beginning analysis, both confirming its functionality, and exploring those aspects of interest. Steps denoted with **(P)** will represent those which only apply to parallelized implementation, and at NERSC specifically. Older scripts utilizing the QSub framework are available upon request, but functionally behave almost identically.

0) (NERSC) In your home directory, append the following statements to your (hidden) `.bashrc.ext` file, where `\n` indicates physically starting a new line: `module load root \n module load gcc \n module swap PrgEnv-intel PrgEnv-gnu`. This will load the modules required for using the code, and change environments to the appropriate one as well, all automatically at each login. This need be done only once. For evaluating performance of parallelization, additional modules should be loaded (if available); `module load scalasca \n module load scorep \n module load cube`.

1) (P) configure number of nodes to be used. This value must be changed in multiple places.

a) Inside `mpijobfarmer.sh`, near the top, is a variable named “nodes” to be set. This establishes the numbering convention for batches and specifies how the summing routine should perform.

b) Inside `mpijobrunner.sh`, in the header, change `#SBATCH -N nodes` so that the value for “nodes” matches that defined inside the farmer script. This sets the number of nodes that the script will request from the cluster.

c) Inside `MPIOMPtransportcode.c`, at the beginning of “main”, change `corenum` to be

the product of nodes * cores per node. Cores per node will be set by the cluster's specific architecture (for example, Edison has 24 cores/node, while Cori has 32 cores/node), along with your choice for intranodal parallelization (utilizing OpenMP). Cores per node are decided by running test cases and determining ideal performance. For example, specifying 5 nodes, with 2 cores allocated per job on Edison would result in setting $corenum=5*(24/2)$. This should not be required for implementing hyperthreading.

2) **(P)** Set the amount of time you expect your request to take. As the farmer script utilizes the shared partition, these times can be made as long as allowable (currently 36 hours at NERSC), as the usage cost associated with a single core is minimal. However, the runner script should be made as low as achievable, both to contain the impact of any errors when running, and to expedite allocation of cluster resources to your job. These are adjusted in the Slurm headers, by altering `#SBATCH -t hh:mm:ss`

3) Set the purpose of the run. Choose whether you're running a simulation, or summing previous runs. Then, in the farming script, set "runsum" to 0 for running, or 1 for summing. If summing, make sure to set the appropriate way to iterate over runs to be summed.

4) Set the runlist to be performed. In the farming script, near the bottom, set the nested for-loops to iterate over whichever variables you wish to vary. Special note; "sfload" chooses run type, which can be chosen from an already populated list, or else new run types can be constructed. Make sure to create these run types inside the "writing()" subroutine, with unique "sfload" identifiers. This is also where you may adjust the number of particles per pulse (simply "n"), and the pulse parameters.

5) Submit the job. Once everything is configured properly (make sure to run test cases with small numbers of nodes, like 1 or 2, before committing to large statistic run cycles), you

must submit the job. On the cluster, this is done with “sbatch mpijobfarmer.sh”, whereas the non-parallel version is simply “./Nonparrun 0 0”, where the 0 0 can be changed to any numbers, and represents batch numbering conventions.

- 6) **(P)** Check the job. Typing “sqs” will show you the current state of your jobs, if they’re running, or the ETA until run begins (though this is highly unreliable). When complete, “sqs” should return nothing.
- 7) **(P)** Sum the runs. Sums of all histograms must be taken over every run, for achieving the intended statistics. It is typically best to leave all parameters as they were for running, when changing to summing, though it is not required after the independent numbering convention file has been generated. In the farming script, change “runsum” to 1, and pay attention to the other summing parameters, such as “sumbybatch” and “resum” if particular sums must be taken again. With the summing designation set, submit the job again, with “sbatch mpijobfarmer” as before.
- 8) View the results. Change directory to the highest-order sum reached, and open either the plain “.root” file (the detectors file), or the “...hists.root” file, which consists of the simulation’s various tallies, which you may add to. To open in root, simply type “root filename.root”. Once opened, any given tree can be viewed with the command “.ls”, which will list all available histograms to view.
 - a) For the detector file, pick a detector signal you’d like to view (as can be seen from the command “.ls”). Next, draw the figure; for example, “swthist→Draw()” for the swthist detector. Since binned in 0.1s intervals, it may be desirable to widen the bins for better statistics per bin, with a command such as “swthist→Rebin(10)”. This

change will only persist within the given root session, and will revert to normal binning once the session is exited (with “.q”).

- b) For the histograms file, more complicated histograms (higher dimensionality) are involved, and as such the commands can be slightly more complicated. The list of available histograms can be accessed with the “.ls” command as usual. Typically the most helpful first view is by checking the spatial distribution of all neutrons over all time; with even modest statistics, this generally gives a helpful visualization of the geometry, and is useful for trouble shooting.

To do so, we utilize the 3-d histogram plusXZcounts or minusXZcounts, where the plus/minus refers to the spin flipper loading state (respectively on, or off). Then, the time window may be chosen, with

```
“plusXZcounts→GetXaxis()→SetRangeUser(t1,t2)”
```

Then, a 2-D visualization for this time range can be made by projecting onto a 2-D histogram, with

```
“TH2F *hplus=(TH2F*)plusXZcounts→Project3D(“zy”)”
```

And drawn with “hplus→Draw(“COLZ”)”. Note: you may find it preferable to change the default color palette. Simply use “gStyle→SetPalette(1)” for better coloring scheme.

Files

List of files and their purposes.

Reg_ucna_5.0

The regions file (for UCNA). This file includes the definition of each region making up the geometry of the apparatus, and the properties/parameters to be used for each.

Conn_ucna_5.0

The connections file (for UCNA). This file specifies how the components defined in Reg_ucna_5.0 are linked together.

Config_testing.txt

The configuration file for each run. Has a multitude of parameters that are read in by the main transport code for defining the type of run simulated. This file is the one most often changed.

MPIOMPtransportcode.c/Nonparcode.c

The transport code itself. MPIOMP refers to the parallelized implementation of the MPI and OpenMP protocols for inter- and intra-nodal parallelization, whereas Nonparcode is the de-parallelized version.

Optdetread.c

Optimized detector read. This file is run after the simulation is completed, and converts the output from each run into a consolidated root file, and constructs histograms for each detector signal, and sums them across runs. This is almost never changed

Treemerger.L

A root header file that allows implementing various subroutines/functions within the code. While within a root session, added with “.L treemerger.L”, and then all functions defined

within the file are accessible from the command prompt, while the session persists. Can also be loaded into standard C-programs, and any of the functions employed during the run of the code itself. Good for doing quick analysis across run types.

Makefile

Standard makefile, which compiles the various codes, using generalized environmental variables. This is almost never changed.

Makechain.C

A simple C routine that performs a consolidation of the various histograms produced by running the code, through a ROOT function utilizing a TChain structure of TTree's. Found to be significantly faster than alternative methods of synthesizing individual TTree's. This is never changed.

Parallelized batch scripts

Mpijobfarmer.sh

Slurm script for creating multiple jobs to be submitted to the cluster. Contains loops for iterating over many different variables that may be scanned over to explore parameter space. Responsible for writing configuration files, and submitting jobs to the cluster for either running the transport code, or summing it. Also defines a running list of batch numbering based on all altered parameters for later reference. This is changed most frequently of all files.

Mpijobrunner.sh

The Slurm script responsible for actively submitting jobs to the cluster for running the transport code under specific conditions. Called by the farmer script, this specifies the parallelization parameters (such as nodes, time requested, etc.), and runs the jobs. This is typically only changed when changing node counts.

Mpijobsummer.sh

The slurm script responsible for summing histograms after specific runs are completed. Uses a logarithmic reduction method by recursively summing in groups of 10. Reports if files are missing during the attempted sum, and then moves summed files to their proper naming convention (as specified in the farmer script).

Geometry input files:

Conn_ucna_5.0 details:

Can support up to 6 connections per piece. Each regions entry must be ended with a “/” character. The first entry in each row is the region # to be discussed, and the following entries are those to which it is connected. Special handling codes can be included in the first entry, along with the region number, for describing unique cut-plane properties.

1→ Cut-plane is 100% absorptive.

2→ Special instruction for specifying an input value for a given cut-plane. Example given;

5 4,2(38) 6

Would describe that region 5 is connected to region 4, and that region 5’s cut-plane should be associated with some input value of 38.

3→ Cut-plane is a spin flip resonance.

4→ Cut-plane is a detector

5→ Cut-plane separates dissimilar-sized, same-type guides (i.e., with a lip).

Reg_ucna_5.0 parameterization is as follows:

Reg#: Region number, required to be listed in sequential numeric order. Used to later reference particular region

RType: Region type(shape). Currently supports only rectangular prisms (RType==1), and cylindrical systems (RType==2)

BP(x,y,z): Basepoint coordinate. Basepoints are defined as the geometric center of the face from which the component is extruded; in all but the initial guide components, the basepoint/baseplane is an imaginary surface representing transition from 1 guide section to the next. * represents the natural continuation from the previous segment. The “>” symbol can be used to shift the basepoint relative to a different guide sections basepoint, for beginning new contiguous regions.

Dim: Dimension(size) of each region. For rectangular prisms, represents length in the x,y, and z directions. For cylinders, represents “diameter, length”.

Orient: Orientation for extrusion of shape from baseplane. Represents angles ψ, θ, ϕ , corresponding respectively to rotations about the global y-axis, global x-axis, and then global z-axis, all in degrees.

Grad B: Specifies a magnetic field gradient in the region, with positive values corresponding to a positive gradient in the direction from the baseplane forward. Values are dB/dx, dB/dy, and dB/dz, with coordinates *relative to the component*, not global coordinate frames. E.g., field gradients directed longitudinally are represented by only a dB/dz value. Units in T/m

Spec: Specularity of wall interactions. Specifies the fractional chance of specular reflection from surfaces (i.e., that incident and reflected angle are the same). In case of failed probability roll, results in $\cos\theta$ distribution for a randomized reorientation.

Loss: Loss per bounce. Specifies the fixed probability of loss on reflection, if reflection occurs. Independent of the energy dependent loss rate from exponentially decaying tunneling produced by the transverse velocity component impinging on a given Fermi potential.

Depol: Depolarization per bounce. A fixed probability for depolarizing on a given wall interaction.

WPot; Wall potential. Represents the surface Fermi potential presented by the material, in neV.

BPot: Bulk Potential. The effective potential experienced by a neutron while inside a material. Also responsible for velocity boosts upon exiting the material. Units in neV.

Surf/Scat: Surface/scattering model for interactions in bulk medium. Allows specifying mean free path within a material.

Absorb: Absorption rate while within a bulk medium. Represents exponential decay probability a , such that the probability of loss from transiting for a duration of T is $1 - e^{-aT}$.

Det: Detector designation. 0 corresponds to a standard region, whereas a numerical entry designates that region as a numbered detector region, such that any absorptions in the region result in an increased count tally within the detectors output file.

PM: Propagation mode for bulk media.

0→ Medium has a material potential, found in BPot.

1→ Absorbing medium, such as detectors.

2→ Foil transmission. Includes formula for incoherent sum of reflections. Uses BPot, but assumes vacuum on either side of foil.

3→Scattering medium, using mean free path from Surf/Scat

4→Scattering+Absorbing medium.

5→ Trajectory tracking; events recorded for fixed interval of time. Use only for very small numbers of particles, as event files may become very large.

SM: Scattering Model. Allows implementing different surface scattering models.

0→ Specify probability using Spec

1→ Probability of nonspecular bounce, weighted by cosine of angle between velocity and surface normal.

2→ User defined.

LM: Loss Model.

0→ Fixed wall loss, no energy dependence

1→ Energy dependent loss-per-bounce

2→ Energy and angle dependent loss-per-bounce.

DM: Depolarization model.

0→ Specify probability of depolarization, as in Depol

1→ User defined

Appendix C

Regions parameters file

The following is the regions header file associated with the primary simulation code. It contains each region within the simulation, and its physical properties, such as the shape, angle, magnetic field, wall potentials, etc.. These represent the generally optimized values found throughout simulation campaigns, and were arrived at via hundreds of iterative tests, evaluating effects on dozens of output parameters. See Appendix B for information on how to interpret the specifications.

Particular Regions of Interest and their associated part numbers are provided, as an aid.

Part Numbers	Associated Geometry Component
0-33	UCN Source→Gate Valve
34-51	Gate Valve→ PPM→Guide Elbow
52-75	Guide Elbow→AFP
76-110	Electropolished Copper Guide
111-131	DLC Coated Copper
132-157	AFP
122	Spin Flip Location (1T in AFP)
158	Shutter
159	Rectangular Guide
160-165	Decay Trap
166+	Various Detectors and their guides


```

158 1 < 0.027305,0.05715,0.7112 -60.0,0.0,0 0.0,0.0,1.4 0.960 4.0e-4 2.0e-6 220. 0.0 0.0 0.0 0 0 0 1 0
159 2 [158]:-1.50368,0.0,0.77312 0.12383,0.003175 30.0,0.0,0.0 0.0,0.0,0.0 0.900 4.0e-4 2.0e-6 252. 252. 0.0 1.0e7 0 1 0 1 0
160 2 * 0.12383,1.0 30.0,0.0,0.0 0.0,0.0,0.0 0.960 3.0e-4 2.0e-6 252. 0.0 0.0 0.0 0 0 0 1 0
161 2 * 0.12383,5.0e-4 30.0,0.0,0.0 0.0,0.0,0.0 0.000 0.5 2.0e-6 168. 0.0 0.0 0.0 0 0 0 1 0
162 2 * 0.12383,1.0 30.0,0.0,0.0 0.0,0.0,0.0 0.960 3.0e-4 2.0e-6 252. 0.0 0.0 0.0 0 0 0 1 0
163 2 * 0.12383,5.0e-4 30.0,0.0,0.0 0.0,0.0,0.0 0.000 0.5 2.0e-6 168. 0.0 0.0 0.0 0 0 0 1 0
164 2 * 0.12383,1.0 30.0,0.0,0.0 0.0,0.0,0.0 0.960 3.0e-4 2.0e-6 252. 0.0 0.0 0.0 0 0 0 1 0
165 2 * 0.12383,0.003175 30.0,0.0,0.0 0.0,0.0,0.0 0.900 4.0e-4 2.0e-6 252. 252. 0.0 1.0e7 0 1 0 1 0
166 2 [28]:>0.0,-0.0508,0.0508 0.0102,0.0508 0.0,0.0,0.0 0.0,0.0,0.0 0.990 2.0e-4 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
167 2 * 0.1016,1.0 0.0,0.0,0.0 0.0,0.0,0.0 0.990 2.0e-4 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
168 2 * 0.1016,1.0e-4 0.0,0.0,0.0 0.0,0.0,0.0 0.300 2.0e-4 2.0e-6 70. 70. 0.0 3059 0 2 0 1 0
169 2 * 0.1016,0.1016 0.0,0.0,0.0 0.0,0.0,0.0 0.300 2.0e-4 2.0e-6 70. 0.0 0.0 1.0e7 1 1 0 1 0
170 2 [142]:>0.0,-0.035,0.08668 0.00635,0.005 0.0,0.0,0.0 0.0,0.0,0.0 0.960 4.0e-4 2.0e-6 168. 0.0 0.0 0.0 0 0 0 1 0
171 2 * 0.0762,2.5e-4 0.0,0.0,0.0 0.0,0.0,0.0 0.000 0.50 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
172 2 * 0.0762,2.5e-4 0.0,0.0,0.0 0.0,0.0,0.0 0.000 0.50 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
173 2 * 0.0762,1.0 0.0,0.0,0.0 0.0,0.0,0.0 0.990 2.0e-4 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
174 2 * 0.0762,4.0e-5 0.0,0.0,0.0 0.0,0.0,0.0 0.500 2.0e-4 0.0 100. 0.0 0.0 0.0 0 0 0 1 0
175 2 * 0.0762,1.0e-4 0.0,0.0,0.0 0.0,0.0,0.0 0.300 2.0e-4 2.0e-6 70. 70. 0.0 3059 0 2 0 1 0
176 2 * 0.0762,0.1016 0.0,0.0,0.0 0.0,0.0,0.0 0.300 2.0e-4 2.0e-6 70. 0.0 0.0 1.0e7 2 1 0 1 0
177 2 [162]:>0.0,-0.061915,0.5 0.00635,0.005 0.0,0.0,0.0 0.0,0.0,0.0 0.960 4.0e-4 2.0e-6 168. 0.0 0.0 0.0 0 0 0 1 0
178 2 * 0.0762,2.5e-4 0.0,0.0,0.0 0.0,0.0,0.0 0.0 0.5 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
179 2 * 0.0762,1.0 0.0,0.0,0.0 0.0,0.0,0.0 0.990 2.0e-4 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
180 2 * 0.0762,1.0e-4 0.0,0.0,0.0 0.0,0.0,0.0 0.300 2.0e-4 2.0e-6 70. 70. 0.0 3059 0 2 0 1 0
181 2 * 0.0762,0.1016 0.0,0.0,0.0 0.0,0.0,0.0 0.300 2.0e-4 2.0e-6 70. 0.0 0.0 1.0e7 3 1 0 1 0
182 2 [72]:0.0,0.0,0.0 0.0700,0.0254 120.,0.0,0.0 0.0,0.0,0.0 0.960 4.0e-4 2.0e-6 168. 0.0 0.0 0.0 0 0 0 1 0
183 2 * 0.0700,0.0508 120.,45.,0.0 0.0,0.0,0.0 0.960 4.0e-4 2.0e-6 168. 0.0 0.0 0.0 0 0 0 1 0
184 2 * 0.0700,0.004 120.,45.,0.0 0.0,0.0,0.0 0.01 0.990 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
185 2 * 0.0700,0.084667 120.,45.,0.0 0.0,0.0,0.0 0.990 2.0e-4 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
186 2 * 0.0700,0.002 120.,45.,0.0 0.0,0.0,0.0 0.01 0.990 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
187 2 * 0.0762,0.0762 120.,45.,0.0 0.0,0.0,0.0 0.990 2.0e-4 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
188 2 * 0.0762,0.0762 120.,75.,0.0 0.0,0.0,0.0 0.990 2.0e-4 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
189 2 * 0.0762,0.002 120.,75.,0.0 0.0,0.0,0.0 0.50 0.8 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
190 2 * 0.0762,0.470154 120.,75.,0.0 0.0,0.0,0.0 0.990 2.0e-4 1.0e-3 180. 0.0 0.0 0.0 0 0 0 1 0
191 2 * 0.0762,1.0e-4 120.,75.,0.0 0.0,0.0,0.0 0.300 2.0e-4 1.0e-3 180. 70. 0.0 3059 0 2 0 1 0
192 2 * 0.0762,0.1016 120.,75.,0.0 0.0,0.0,0.0 0.300 2.0e-4 1.0e-3 180. 0.0 0.0 1.0e7 4 1 0 1 0
193 1 [70]:0.0,0.0,0.00064 0.127,0.381,0.215 -60.0,0.0,0.0 0.0,0.0,0.0 0.010 0.1 2.0e-6 70. 0.0 0.0 0.0 0 0 0 1 0
194 1 * 0.127,0.381,0.215 -60.0,0.0,0.0 0.0,0.0,0.0 0.010 0.1 2.0e-6 70. 70. 0.0 1.0e7 0 1 0 1 0
/

```


Appendix D

Connections Files

The following is the connections header file associated with the primary simulation code.

This specifies for each region number, the other regions to which it is connected directly, to create the simulated geometry. See Appendix B for further information

Region#	Connects through its own cut-plane to:	AND also to:	AND also to:	AND also to:	AND also to:	AND also to:
0,1	0	1	/			
1	*					
2	*					
3	*					
4	*					
5	*					
6	5	7	8	/		
7	6	/				
8	6	9	/			
9	*					
10	*					
11	*					
12	*					
13	*					
14	*					
15	*					
16	*					
17	*					
18	*					
19	*					
20	*					
21	*					
22	*					
23	*					
24	*					
25	*					
26	*					
27	*					
28	27	29	166	/		
29	*					
30	*					
31	*					
32	*					
33,5	32	34	/			
34	*					
35	*					
36	*					
37	*					
38	*					
39	*					
40	*					
41	*					
42	*					
43	*					
44	*					
45	*					
46	*					
47	*					

48 *
49 *
50 *
51 *
52 *
53 *
54 *
55 *
56 *
57 *
58 *
59 *
60 *
61 *
62 *
63 *
64 *
65 *
66 *
67 *
68,5 67
69 *
70 *
71 *
72 *
73 *
74 *
75 *
76 *
77 *
78 *
79 *
80 *
81 *
82 *
83 *
84 *
85 *
86 *
87 *
88 *
89 *
90 *
91 *
92 *
93 *
94 *
95 *
96 *
97 *
98 *
99 *
100 *
101 *
102 *
103 *
104 *
105 *
106 *
107 *
108 *
109 *
110 *
111 *
112 *
113 *
114 *
115 *
116 *
117 *

69 /

118	*				
119	*				
120	*				
121	*				
122	*				
123	*				
124	*				
125,3	124	126	/		
126	*				
127	*				
128	*				
129	*				
130	*				
131	*				
132	*				
133	*				
134	*				
135	*				
136	*				
137	*				
138	*				
139	*				
140	*				
141	*				
142	141	143	170	/	
143	*				
144	*				
145	*				
146	*				
147	*				
148	*				
149	*				
150	*				
151	*				
152	*				
153	*				
154	*				
155	*				
156	*				
157	*				
158	157	162	/		
159,1	159	160	/		
160	*				
161	*				
162	161	163	177	158	/
163	*				
164	*				
165	164	/			
166	28	167	/		
167,5	166	168	/		
168	*				
169	168	/			
170	142	171	/		
171,5	170	172	/		
172	*				
173	*	/			
174	*				
175	*				
176	175	/			
177	162	178	/		
178,5	177	179	/		
179	*				
180	*				
181	180	/			
182	182	183	/		
183	*				
184	*				
185	*				
186	*				
187,5	186	188	/		

188 *
189 *
190 *
191 *
192 191 /
193 193 194 /
/ 194 193 /

Appendix E

Guide Parameters for UCNA

Below is a compilation of guide transport parameters either measured or used in simulations for UCNA relevant for 2011-2013 running. A thorough description of the guide system can be found in the thesis of B. Vorndick [3]. Data is quoted from individual guide measurements (M) in Vorndick’s thesis or XPS data from Vorndick’s or R. Mammei’s[90] theses, the “in situ” density measurements (D) performed in B. Vorndick’s thesis, or simulations (S) of UCN transport results in either Vorndick’s thesis or A. Holley’s thesis[2]. When a particular row has both measured data and parameters used in simulations, each field will have either an “M” or an “S”. A reference Cu guide (and an SS guide) were incorporated into all individual guide measurements in 2011 and 2012. The averaged Cu guide results are denoted by “M*” in below.

Name	In sys	Coating/Substrte	Dims (in)		Region	V _F (neV)	Spec	Lpb (10 ⁻⁴)	Meth M,D,S	Ref
			L	D (or A)						
GV-PPM	2010	Cu	26	3	PPM	168	0.930(15)	(1.1-2.4)	M*	1
PPM-Elbow	2010	Cu	44	3	PPM-Sw	168	0.930(15)	(1.1-2.4)	M*	1
0-60-Elbow	2010	Cu	18 CLR	3	PPM-Sw	168	0.930(15)	(1.1-2.4)	M*	1
	2010				PPM-Sw	168	0.96	4	S	2
	2014				PPM-Sw	168	0.85	4	D	3
Sw-Entrnc	2010	Cu	6.03	2.75	Sw	168	0.930(15)	(1.1-2.4)	M*	1
Sw-Thru	2010	Cu	5.47	2.75	Sw	168	0.930(15)	(1.1-2.4)	M*	1
Sw-Down	2010	Cu	5.57	2.75	Sw	168	0.930(15)	(1.1-2.4)	M*	1
Sw-Down	2010	Cu	6.03	2.75	Sw	168	0.930(15)	(1.1-2.4)	M*	1
Sw-Exit	2010	Cu	6.03	2.75	Sw	168	0.930(15)	(1.1-2.4)	M*	1
Sw-AFP	2010	Cu	26.94	2.75	Sw	168	0.92-0.94	(2.1-3.8)	M	2
	2014				Sw	>130		4	D, 4:	
	2014				Sw	168	>0.8	4	gps 1.2 – 1.7mm	

AFP-Upstr	2011	DLC/Cu	39.41	2.75	Sw-AFP	233	0.976(4)	(0.1-1.2)	M	1,5
RF-Quartz	2011	DLC/Q	39.37	2.75	Sw-AFP	264	>0.96 (D)	<3 (D)	M,D	1,5
AFP-Dwnstr	2011	DLC/Cu	39.53	2.75	Sw-AFP	210	0.978(4)	(0.6-1.4)	M	1,5
	2014				Sw-AFP	>170	>0.96	<3	D, 5: gps 1.2 – 1.7mm	
	2014				SW-AFP	M	>0.8	4		
AFP-SCS1	2011	DLC/Cu	8	2.75	AFP-SCS		0.898(11)	(0.5-1.0)	M	1,6
AFP-SCS2	2011	DLC/Cu	2	2.75	AFP-SCS		0.898(11)	(0.5-1.0)	M	1,6
Shutter-SG(2)	2011	DLC/Cu	23	2.75	AFP-SCS	201	0.980(4)	(1.1-1.7)	M	1,6
Shutter-plate	2011	DLC/Cu	¼ thick	>2.75	Spectrum					1
Shutter rd-rect	2011	DLC/Cu	0.125		Spectrum					1
Shutter-rect	2010	DLC/Cu	28	2.75x1.58	Spectrum	245	0.96 (S)	$4 \times 10^{-4}(S)$	M,S	2,7
DT center	2011	DLC/Cu	39.37	5.125	Spectrum	230	0.96 (S)	$4 \times 10^{-4}(S)$	M,S	2,7
DT end E	2011	DLC/Cu	39.37	5.125	Spectrum	241	0.96 (S)	$4 \times 10^{-4}(S)$	M,S	2,7
DT end W	2011	DLC/Cu	39.37	5.125	Spectrum	230	0.96 (S)	$4 \times 10^{-4}(S)$	M,S	2,7
DT-plug	2011	DLC/Cu			Spectrum	231	0.96 (S)	$4 \times 10^{-4}(S)$	M.S	2,7
Be Foil E	2011	Be/PET	150nm	500nm	Spectrum	252			S	2
BE Foil W	2011	Be/6F6F	150nm	500nm	Spectrum	252			S	2
DT center	2012	DLC/Cu			Spectrum	244			M	6,7
DT end E	2012	DLC/Cu			Spectrum	220			M	6,7
DT end W	2012	DLC/Cu			Spectrum	230			M	6,7
Be Foil E	2012	Be/6F6F	150nm	180nm	Spectrum	252			S	8
Be Foil W	2012	Be6F6F	150nm	130nm	Spectrum	252			S	8

1. Vorndick Thesis: using reference EP-Cu guide individual guide tests performed in both 2011 and 2012 for transmission and lifetime. The full range (lowest value – 1σ to highest value + 1σ) of 20.83 – 24.5s and any gap size in measurement below about 0.1mm (even 0.15mm gaps were not consistent with these lifetimes measurement, see p. 99) was consistent with a lpb from $1.1-2.5 \times 10^{-4}$. The specularities measured had a 1σ full range of from 0.914 – 0.944, see p. 65. Nominal Fermi potentials are assumed to be 168 neV, consistent with all measurements. The reference EP-Cu guide can be compared to the SW-AFP Cu guide, installed

in 2010 and possibly contaminated, which had a somewhat larger maximum loss per bounce), and the nominal value of 4×10^{-4} .

2. Holly Thesis: Holley successfully represented transport in the UCNA spectrometer guide system and the crossed polarizer analyzer using these Cu guide parameters (as per, for example, pg.359 – 360, and the plots after this). The same values of specularity and loss per bounce were also used for DLC guides. The Fermi potential for Be was set to 252 neV.

3. Vorndick Thesis: the density in the guides was measured in the “in situ”, 2014 density tests from the exit of the PPM to the spectrometer side of the AFP. The small reduction in density between the PPM and the switcher were reproduced by adjusting the specularity in the region between the PPM and the switcher region to 0.85. For details on the guide density test, see Table 4.5 and pages 74-76 for a description of the measurements, and pages 99-104 for analysis.

4. Vorndick Thesis: large drops (~15%) in the measured density in this region required the introduction of gaps. Gaps consistent with specularity of the guides greater than 0.6 required gaps between 1.2 and 1.7 mm for lpb with the nominal value of 4×10^{-4} for lpb and 168 for the Fermi potential. If the lpb and specularity were held at nominal values of 4×10^{-4} and 0.96, any Fermi potential greater than 130 neV was consistent with the measured density (or a greater value).

5. Vorndick Thesis: another large drop (~15%) in the measured density in this region required the introduction of gaps, as in 4, but somewhat smaller in size. Only specularities > 0.8 could be made to fit with loss per bounce at the nominal 4×10^{-4} value and Fermi potentials as

measured, with a gap range of 0.8 – 1.4mm (consistent with measured specularities). If the lpb and specularity were held at nominal (measured) values, the lowest Fermi potential consistent with the measured density was greater than 170 neV.

6. Vorndick Thesis: DLC densities were determined by XPS. XPS scans and thickness measurements exist for all DLC guides, but Fermi potentials were not reported for a few.

7. Mammei Thesis: guide fabrication procedures were established in this document. All DLC-coated Cu guides were mechanically polished, then electropolished for nominal 5nm Ra roughness to ensure good DLC adhesion (this included the decay trap guides, rectangular guides and shutter pieces). Most parts should be smooth and exhibit properties similar to the relatively well-characterized guides in the Sw-AFP region.