

Spin Dressing

Thesis by
Raymond Tat

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy in Physics



CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2026
Defended June 11, 2025

© 2026

Raymond Tat

ORCID: 0000-0002-5648-0360

All rights reserved except where otherwise noted

ACKNOWLEDGEMENTS

I would firstly like to thank my parents, Vinh and Miyo Tat, for their unwavering encouragement and support. Words cannot express my gratitude for everything they have given me over the years. Thank you to Eric Fries, Xuan Sun, and Marie Blatnik, for welcoming me to this research group with open arms. Thank you to the many researchers and staff at Duke for generously lending their time, resources, and insight to get the SEOP apparatus operational. I am particularly grateful for the help of Zhiwen Zhao and Haiyan Gao, who graciously allowed me to set up shop in Physics 014. Thank you also to Adam Dipert, the previous operator of the SEOP system, for helping us acquire all the necessary tools and equipment for the lab. Adam's notes were an invaluable resource for this project. Thank you to Bob Golub, who shared his vast knowledge of all things NMR over many long discussions. A special thank you to Wolfgang Korsch for pointing out that the SEOP cell should be rotated to more efficiently heat the rubidium. This comment likely saved me from several months of trial and error. Thank you to Chris Swank, who taught me the fundamentals of spin dressing, and who guided me through the early years of the Ph.D. program. And of course, thank you to my advisor, Brad Filippone, without whom this entire project would have been impossible. It is no mean feat to steer a ship through such uncertain waters, but I am grateful to have Brad at the helm.

ABSTRACT

Measurement of a neutron electric dipole moment (nEDM) is a sensitive probe of CP-violating physics beyond the Standard Model. Such a measurement could provide clues as to the origins of the matter-antimatter asymmetry. Spin dressing is a technique where an off-resonant oscillating magnetic field is applied to a spin or ensemble of spins to modify their effective gyromagnetic ratio. This technique was proposed for use in the neutron electric dipole moment experiment at the spallation neutron source (nEDM@SNS) to increase statistical sensitivity and mitigate systematics. Here we present a comprehensive study of spin dressing in the context of EDM measurement from theoretical, computational, and experimental perspectives.

PUBLISHED CONTENT AND CONTRIBUTIONS

Below is a list of the author's published papers that contributed to this thesis. "Quantum control of two critically dressed spin-1/2 species in magnetic fluctuations," published in *Physical Review A*, was authored by Chris Swank and me. I contributed to the simulations for this work, as well as the theoretical calculation of the frequency shifts and relaxation in the spin-dressed case. I am the sole author of "The semi-classical Floquet-Markov master equation for Monte Carlo spin integration," which has been published in *Physical Review A*.

- [Tat25] R. Tat. "Semiclassical Floquet-Markov master equation for Monte Carlo spin integration." In: *Phys. Rev. A* 112 (1 July 2025), p. 013113. DOI: 10.1103/2bp2-hs9d. URL: <https://link.aps.org/doi/10.1103/2bp2-hs9d>.
- [TS22] R. Tat and C. M. Swank. "Quantum control of two critically dressed spin-1/2 species in magnetic fluctuations." In: *Phys. Rev. A* 105 (5 May 2022), p. 053120. DOI: 10.1103/PhysRevA.105.053120. URL: <https://link.aps.org/doi/10.1103/PhysRevA.105.053120>.

TABLE OF CONTENTS

Acknowledgements	iii
Abstract	iv
Published Content and Contributions	v
Table of Contents	v
List of Illustrations	ix
List of Tables	xiv
Chapter I: Introduction	1
Chapter II: CP violation and the Baryon Asymmetry	3
2.1 The Baryon Asymmetry (Or the Absence of an Anti-Elephant in the Room)	3
2.2 Baryon Number Violation	3
2.3 Out-of-equilibrium Interactions	4
2.4 C- and CP-Violation	5
2.4.1 Symmetry Breaking in the Standard Model	5
2.5 Electric Dipole Moments (EDMs) and CP Violation	6
Chapter III: Measuring an nEDM	8
3.1 Ultracold Neutrons	8
3.1.1 Why UCN?	8
3.1.2 What is a UCN, Anyway?	9
3.1.3 Reflection from Surfaces	9
3.2 UCN Production	10
3.3 The State of nEDM Experiments	11
3.4 nEDM@SNS Experiment Overview	12
3.4.1 Measurement Principle	12
3.4.2 Overview of the Apparatus	16
3.4.3 UCN Storage	16
3.4.4 Signal Detection	18
3.4.5 Electric Field System	18
3.4.6 Magnetic Field System	19
Chapter IV: Spin Dressing Theory	22
4.1 Quantum Control of Two Critically-Dressed Spin-1/2 Species	22
4.1.1 Spin Dressing Hamiltonian	24
4.1.2 Theory	24
4.1.3 Simulation	38
4.1.4 Mitigation Strategies	41
4.1.5 Robust Dressing	43
4.1.6 Relaxation under Robust Dressing	50
4.1.7 Conclusion	56
4.1.8 Appendix to [TS22]	58

4.2	Magic Dressing	65
4.2.1	Magic Dressing for a Single Species	65
4.2.2	Magic Dressing for Two Species	65
4.3	Considerations for Spin Dressing Data Analysis	69
4.3.1	Magnetic Field Fluctuations as a Random Walk	70
4.3.2	Matrix Formulation	71
4.3.3	Continuum Limit	74
4.3.4	Application to nEDM	77
Chapter V:	Spin Integration - Computational Techniques	81
5.1	Commutator-Free Exponential Time-Propagator (CFET) Method	81
5.1.1	The Magnus Expansion	82
5.1.2	The Baker-Cambell-Hausdorff (BCH) Formula	84
5.1.3	Gauss Quadrature	85
5.1.4	Performance	85
5.2	The semi-classical Floquet-Markov Master Equation for Monte Carlo Spin Integration	87
5.2.1	Introduction	87
5.2.2	Spin Physics and Frequency Shifts	88
5.2.3	Floquet Theory and the Floquet-Markov Master Equation	89
5.2.4	Floquet Diagonalization	90
5.2.5	Spectrum Estimation	91
5.2.6	Acceleration with Graphics Processing Units (GPUs)	95
5.2.7	Comparison with Theory	96
5.2.8	Performance	99
5.2.9	Conclusion	101
5.2.10	Acknowledgments	102
5.2.11	Simulation Parameters	102
5.2.12	Floquet-Markov Master Equation Derivation	103
5.2.13	Cross-Spectrum Properties	110
Chapter VI:	SEOP Studies of Spin Dressing Techniques at Duke University	111
6.1	Spin Exchange Optical Pumping (SEOP) Theory	111
6.1.1	Buffer Gas	112
6.1.2	Hybrid SEOP	112
6.2	Description of SEOP Apparatus	113
6.2.1	Laser Optics	114
6.2.2	Heating System	114
6.2.3	Magnetic Field System	115
6.2.4	SEOP and NMR Electronics	119
6.3	NMR Techniques	121
6.4	B_0 and Ambient Field Characterization	123
6.5	Polarization Estimate	125
6.6	Demonstration of Magic Dressing	127
6.6.1	Dressed Tipping	127
6.6.2	Measurement Scheme	128
6.6.3	Analysis of Magic Dressing Data	129

6.6.4 Discussion of Results	133
Chapter VII: Conclusion and Outlook	135
7.1 Spin Dressing Theory: What's Next?	136
7.2 Bigger Computers, Faster Algorithms	136
7.3 Further SEOP Studies	137
7.4 nEDMSF: Where We Stand	138
Bibliography	139

LIST OF ILLUSTRATIONS

<i>Number</i>	<i>Page</i>
3.1 Phonon dispersion curve (blue) [DB98] and free neutron dispersion curve (orange). UCNs with energy near the intersection can be brought to rest by scattering a single phonon.	11
3.2 nEDM 90% confidence level upper bound over time. The current limit is shown as a dotted line. The Standard Model nEDM (not shown) is below the bottom of the y-axis, at $d_n = 10^{-31} \text{ ecm}$. Plot is taken from the 2022 Snowmass proceedings [Ala+22].	12
3.3 Full-scale prototype cell illuminated under UV light. Image from [Leu25].	17
3.4 3D COMSOL simulation of the Cavallo voltage multiplier in the nEDM@SNS central volume, in which the central electrode is charged to 650 kV. Figure copied from [Cla+18].	19
3.5 Schematic of the inner magnet volume (IMV). G10 frame not shown.	19
3.6 Partially wound B_0 coil with G10 frame. Image from [Sun22].	20
3.7 COMSOL model of the dressing coil used to simulate Joule heating due to the dressing field.	21
4.1 Run-to-run uncertainty in the measured frequency due to dressing field noise for $57 \text{ pT}/\sqrt{\text{Hz}}$ and $5.7 \text{ pT}/\sqrt{\text{Hz}}$ dressing field noise and a perfect band-pass filter around the dressing frequency, calculated from equation (4.76). The predicted per-run sensitivity for the nEDM@SNS, formulated from reference [Ahm+19], is shown for comparison. It is shown that $57 \text{ pT}/\sqrt{\text{Hz}}$ dressing field noise will significantly reduce the ultimate sensitivity of the measured frequency, while $5.7 \text{ pT}/\sqrt{\text{Hz}}$ dressing field noise is acceptable.	37
4.2 Predicted and simulated values of Γ_ϕ for various high-pass filter cutoffs. The theory line is calculated from equations (4.70) and (4.83). Each simulated point represents 200 neutrons. Note that cutoff values beyond the dressing frequency (1000 Hz) are only achievable in simulation, and are shown here for comparison to the theory.	40

4.3	Predicted and simulated frequency shifts, with magnetic field noise given by equation (4.83), for various high-pass filter cutoffs. The theory line is calculated from equations (4.71) and (4.83). Each simulated point represents 100,000 neutrons.	41
4.4	The root mean squared deviation of the angle between the neutron and ^3He spins (ϕ_{n3}) as a function of time, with dressing field powered by a Kepco four-quadrant 400 36-12 power supply, for magnetic fields generated with (solid line) and without (dashed line) feedback control. The dataset without feedback consists of 28 runs, while the dataset with feedback consists of 50. Data are smoothed for clarity.	44
4.5	Phase standard deviation with feedback control applied from the field powered by a Kepco four-quadrant 400 36-12 power supply (solid line). The rms phase deviation rises for a period of less than five seconds, then appears to settle near a fixed value. For comparison, the theoretical estimate (dotted line) is computed assuming $400 \text{ pT}/\sqrt{\text{Hz}}$ noise high-pass-filtered at 500 Hz. Data are smoothed for clarity. . . .	45
4.6	A classical visualization of the neutron spin trajectory undergoing robust dressing on the Bloch sphere in the laboratory frame, with arrows representing the static uniform field, B_0 , the dressing field, B_1 , and the robust dressing effective field B_{eff}	49
4.7	Relaxation after optimization of robust dressing parameters for 200 Hz robust dressing and a wide range of gradient magnitudes. The gradients in this simulation arise from the divergence theorem, $\frac{dB_z}{dz} = -\frac{dB_x}{dx}$	54
4.8	1- Polarization for robust dressing modulated according to equation 4.86, where $B_1 = 40.2497 \text{ } \mu\text{T}$ and $\omega_m = 200 \text{ Hz}$, with simulated noise power varied from 20 dB SNR to 80 dB SNR	55
4.9	Simulated signal for highly tuned RCD modulated at 200 Hz (red) and CSD (green), with theoretical estimation for RCD (blue). The starting phase for ^3He is $\phi_3 = \pi/8$, and it is the same for neutrons, $\phi_n = \pi/8$. Signal is observable for a long time in the RCD scenario, while the signal is indistinguishable from noise in a traditional critical dressing scenario.	57

4.10	Relative precession frequency for neutron and ^3He as a function of dressing amplitude (B_1). Dressing amplitudes where the frequency shifts due to changes in dressing amplitude are equal for neutron and ^3He are shown as dashed and dotted lines.	66
4.11	Two-species modulated magic dressing pulse with optimized parameters.	68
4.12	Heatmap of the frequency difference between neutron and ^3He as holding field strength and dressing amplitude are varied about the magic dressing parameters.	68
4.13	Comparison of the response of magic dressing and unmodulated critical dressing to changes in the dressing field amplitude.	69
4.14	Dot product between the neutron and ^3He spin vectors undergoing two-species magic dressing. The product is evaluated at integer multiples of the magic dressing modulation period $\tau + t_{\text{rest}}$, and the initial angle between the two spin vectors is $\pi/2$	70
4.15	Optimal weights for selected values of α . The upper plot shows the weights in the x_i basis, while the lower shows the same weights in the z_i basis.	73
4.16	nEDM Uncertainty for a single cell, single cycle, versus phase noise due to power supply.	79
4.17	Plot of naive and exact predictions of nEDM uncertainty.	79
4.18	Ratio of predictions (naive/exact).	80
5.1	Performance of the 8th-order Magnus CFET integrator compared to DOP853 for integrating the spins under free precession.	86
5.2	Performance of the 8th-order Magnus CFET integrator compared to DOP853 for integrating the spins in critical spin dressing mode. . . .	86
5.3	Error comparison of Magnus CFET, MagnusGL8 (an integrator available in the DifferentialEquations.jl [RN17] library), and DOP853. The simulation parameters used are the critical dressing parameters, with no gradients or electric field applied.	87
5.4	Comparison of DOP, FMME Goertzel, and FMME FFT methods with theoretical predictions for the B^2 frequency shift in the adiabatic limit. The data points for the three integrators overlap.	97
5.5	Comparison of DOP, FMME Goertzel, and FMME FFT methods with theoretical predictions for the E^2 frequency shift in the adiabatic limit. The data points for the three integrators overlap.	97

5.6	Comparison of DOP, FMME Goertzel, and FMME FFT methods with theoretical predictions for the linear-in-E frequency shift in the adiabatic limit. For these simulations the electric field is varied, while the gradient is fixed at $\frac{dB_z}{dz} = 3 \times 10^{-10} \text{ T m}^{-1}$. A relatively small gradient was chosen to minimize the effect of the B^2 shift in this calculation.	98
5.7	Comparison of DOP, FMME Goertzel, and FMME FFT methods with theoretical predictions for the B^2 frequency shift in the adiabatic limit for spin dressing. The data points for the three integrators overlap.	99
5.8	The B^2 shift estimated by the Floquet Goertzel method varies with window length. The data for 0.1 s and 0.01 s overlap.	100
5.9	The B^2 shift estimated by the Floquet FFT method varies with window length.	100
5.10	Comparison of time taken to simulate various numbers of particles with DOP853, Floquet Goertzel method, or Floquet FFT method. . .	101
6.1	Level diagram of rubidium valence electron, used for optical pumping.	113
6.2	Optics elements to produce two beams of circularly polarized light to optically pump the rubidium vapor. The orientation of the three quarter-wave plates was tuned by mounting a linear polarizing filter in the optics box in conjunction with a laser power meter downstream.	114
6.3	Schematic of optics used for the SEOP system. Having two beam paths that converge at the cell allows us to utilize all 60 W of laser power. It is important to ensure that both beams have the same handedness.	115
6.4	Side view of magnetic field components used for in-situ NMR with the SEOP system, with the oven shown in the center. The outer red coils are the B_0 , the circular black coils are shim coils, and the coils around the oven are the dressing field coils.	116
6.5	One octant of the SEOP dressing coil, simulated in COMSOL with arrows showing the magnetic field direction and magnitude. The sphere octant in the lower right corner represents the SEOP cell. . .	117
6.6	Front view of the SEOP oven, showing the four square dressing coils around in the oven.	118
6.7	Figures showing the design and construction of the pickup coil. . .	119
6.8	Simplified diagram of electronics used for NMR in the SEOP Lab . .	119
6.9	TNMR software interface, showing a sample FID signal.	120

6.10	Electronics cabinet, including the Redstone NMR system used to perform data acquisition and arbitrary waveform generation. The LVC2016 amplifier and SR560 pre-amplifier are at the top of the cabinet.	120
6.11	Pickup coil response during frequency sweep AFP applied to the SEOP cell with polarized ^3He (top) and unpolarized ^3He (middle). The difference between the two (bottom) is also shown.	122
6.12	Aluminum frame used for ambient field and B_0 monitoring. Some of the probe locations visible have been circled.	123
6.13	Frequency response of the pickup coil located in the SEOP oven. . .	125
6.14	Amplitudes of a sequence of FID measurements, along with a fit to the exponential decay in polarization. The angle is given in radians. .	126
6.15	Sample plot demonstrating the shim coil calibration process. The shim coil current is adjusted until the size of the signal caused by the dressed tipping effect is minimized.	128
6.16	Polarization of the ^3He spins after a 10-millisecond dressing pulse at 15 kHz. Four gradient coils are active on each side of the other, each in reverse orientation.	131
6.17	Polarization of the ^3He spins after a 10-millisecond dressing pulse at 15 kHz. Two gradient coils are active on each side of the oven, each in reverse orientation.	131
6.18	Polarization of the ^3He spins after a 10-millisecond dressing pulse at 15 kHz. All gradient coils are disconnected.	132
6.19	Polarization of the ^3He spins after a 10-millisecond dressing pulse at 15 kHz. Two gradient coils are active on each side of the oven, each in standard orientation.	132
6.20	Polarization of the ^3He spins after a 10-millisecond dressing pulse at 15 kHz. Four gradient coils are active on each side of the oven, each in standard orientation.	133

LIST OF TABLES

<i>Number</i>	<i>Page</i>
3.1 Effective Fermi potentials for some UCN bottle materials. Values for non-deuterated materials are taken from [R G91]. Values for dPMMA and dTPB are calculated based on the scattering lengths given in [Sea92] using equation 3.2	9
3.2 Proposed nEDM experiments, their approximate 90% confidence level upper bound on the nEDM after 300 live-days of measurement, and their estimated year for finishing data taking for this precision. Values taken from [Fil22].	13
4.1 Optimized parameters for the starting position on the Bloch sphere for neutrons and ^3He to achieve cosine dressing with a sine dressing field described in section 4.1.2.5. $\phi_{n3} = \phi_n - \phi_3$ is the relative azimuthal angle of the neutron and ^3He that is desired during critical dressing, $\phi_{n,3}$ and $\theta_{n,3}$ are the absolute azimuthal angle and absolute polar angle of the neutron and ^3He , respectively. When the neutron and ^3He start at these positions with sine dressing, the result is cosine dressing dynamics for the respective relative phase ϕ_{n3} . All units are in radians. Values are for the fiducial nEDM@SNS experiment parameters for critical dressing, where the dressing field amplitude is $40.2497 \mu\text{T}$ oscillating at 1 kHz with a static uniform field strength of $5.2 \mu\text{T}$ in the $\hat{z}$ direction. Optimization of the fiducial nEDM@SNS parameters are discussed in reference [Ahm+19].	38
4.2 Optimized parameters for robust critical dressing. All values are calculated for a static uniform field of $5.2 \mu\text{T}$ and amplitude modulation of a nominal critical dressing field of $40.2497 \mu\text{T}$ oscillating at 1000 Hz.	56
4.3 Optimized magic dressing parameters for neutron and ^3He , assuming a dressing frequency of $\omega = 1 \text{ kHz}$ and a holding field strength of $B_0 = 50 \text{ mG}$	67
5.1 Parameters used for comparisons with [Pig+15].	103
6.1 Optimized parameters used for dressing coil design. The actual design parameters will differ from these slightly owing to the thickness of the coil wires.	117
6.2 Basis functions used for static field reconstruction.	124

6.3	Ambient field gradient in $\mu\text{T m}^{-1}$	124
6.4	B_0 coil gradient in $\mu\text{T m}^{-1} \text{ A}^{-1}$	124

Chapter 1

INTRODUCTION

In this dissertation, we discuss the measurement of the neutron electric dipole moment (nEDM), and in particular an experimental technique called spin dressing which may be used to improve the sensitivity of such a measurement. Spin dressing refers to the application of an off-resonant AC magnetic field to a spin to modify its gyromagnetic ratio. The nEDM experiment at the Spallation Neutron Source (nEDM@SNS) proposes to use this technique to improve its statistical sensitivity by a factor of two. The application of such a field presents a variety of opportunities and challenges, which we thoroughly explore in this work. In Chapter 2, we introduce the concepts of CP violation and the baryon asymmetry, which are the primary motivations for conducting measuring the nEDM. In Chapter 3, we give an overview of past, present, and future efforts to measure the nEDM. We describe the various components of the nEDM@SNS apparatus, which leverages the properties of superfluid ^4He to greatly increase its sensitivity. When completed, this experiment would provide a world-leading measurement of the nEDM. Chapter 4 analyzes spin dressing from a theoretical perspective. In the first section of this chapter, which is adapted from the article [TS22], we use time-dependent perturbation theory to derive expressions for the frequency shifts and relaxation rate due to magnetic field fluctuations parallel to the dressing field directions. We then discuss several mitigation strategies for this phenomenon. In the next section, we introduce another way of mitigating this type of noise: magic dressing. This is a modulation scheme that, if applied to an nEDM measurement, could insulate the spins under study from fluctuations in the dressing field amplitude. We conclude this chapter with an analysis of how magnetic field fluctuations could affect data analysis for a nEDM measurement. In Chapter 5, we look at spin dressing from a computational perspective. Simulating spin dressing requires numerical integrators to handle strong time-varying magnetic fields, a situation which is unfavorable for most conventional algorithms. We develop two methods, a Magnus expansion-based integrator and a Floquet-Markov master equation (FMME)-based integrator, which when applied to this problem can produce dramatic improvements in performance. The analysis of the latter method has been published in Physical Review A [Tat25]. In Chapter 6, we use the spin-exchange optical pumping (SEOP) apparatus at Duke

to experimentally demonstrate the magic dressing scheme proposed in Chapter 4. We find that applying the method results in enhanced transverse relaxation times, even in the presence of an artificially large dressing field gradient. We conclude in Chapter 7 with a summary of these results and a discussion of how these endeavors might be continued into the future.

Chapter 2

CP VIOLATION AND THE BARYON ASYMMETRY

2.1 The Baryon Asymmetry (Or the Absence of an Anti-Elephant in the Room)

A long-unsolved problem in physics is that of the baryon asymmetry of the universe (BAU), also known as matter-antimatter asymmetry. The problem is thus: the known laws of physics (general relativity and the Standard Model of particle physics) indicate that the universe should have been created with nearly equal parts matter and antimatter; yet measurements of the cosmos show that the observable universe today contains almost exclusively matter, with very little antimatter. So where does this imbalance come from? Sakharov [Sak] showed that three conditions are necessary to cause a matter-antimatter asymmetry:

1. Baryon number violation,
2. Out-of-equilibrium interactions,
3. C- and CP-symmetry violation.

In the following sections, we discuss potential models that can satisfy these conditions with particular attention to CP violation, as this is the condition that is most relevant to a neutron electric dipole moment (nEDM) search.

2.2 Baryon Number Violation

Baryon number refers to the number of baryons minus the number of anti-baryons; so for example, a set of three protons and two anti-protons would have a total baryon number of one. Baryon number violation — the existence of processes which create or destroy a baryon without an accompanying anti-baryon — is necessary in order to produce more baryons than anti-baryons; after all, if the universe began with baryon number $B = 0$, and baryon number were conserved, then the baryon number would still be zero today, and there could be no matter-antimatter asymmetry. To date no experimental evidence of baryon number violation has been found (though not for lack of trying — see the efforts at Super-Kamiokande to observe proton decay, for example [Tak+20]). Nonetheless, the electroweak sector of the Standard Model

contains a hypothetical “sphaleron” process which can turn three baryons into three anti-leptons, or vice versa, conserving the difference between the baryon and lepton numbers $B - L$. The energy scale for the sphaleron is thought to be of order 10 TeV [HK92], and so is suppressed at the energies accessible with modern particle accelerator technology; however, in the temperatures of the early universe the process could have been prevalent enough to play a significant role in the creation of the BAU. One prominent model of baryogenesis — electroweak baryogenesis — uses the sphaleron process to generate a baryon asymmetry. In electroweak baryogenesis, the asymmetry is created at the electroweak phase transition, when “bubbles” of broken electroweak symmetry form in a previously symmetric universe. In the symmetric phase, the sphaleron process is no longer suppressed, and CP-violating processes at the boundaries of the expanding bubbles of broken phase can create a baryon excess inside of the bubbles, which could persist until the current day if the rate of expansion is rapid enough [HS95]. An alternative model is leptogenesis, where instead of creating a baryon asymmetry outright, a lepton asymmetry is first created, and subsequently this asymmetry is transferred to the baryons via the sphaleron process, which conserves $B - L$. A lepton asymmetry can arise from the decay of heavy sterile neutrinos [DNN08], a particle theorized to explain the smallness of the neutrino masses.

2.3 Out-of-equilibrium Interactions

Even if a baryon asymmetry can be created, the asymmetry needs to persist until the present day. For this to occur, these interactions must occur out of thermal equilibrium, because in thermal equilibrium the processes which produce baryons will be balanced by their respective reverse processes. More formally, this can be shown by examining the baryon asymmetry at equilibrium. If we denote the thermal average of the baryon number at temperature $T = 1/\beta$ as $\langle B \rangle_T$, we have that

$$\langle B \rangle_T = \text{Tr}[\exp(\beta H) B] \quad (2.1)$$

$$= \text{Tr}[(CPT)(CPT)^{-1} \exp(\beta H) B] \quad (2.2)$$

$$= \text{Tr}[\exp(\beta H) (CPT) B (CPT)^{-1}] \quad (2.3)$$

$$= -\langle B \rangle_T \quad (2.4)$$

where we have used the fact that the Hamiltonian H commutes with the CPT operator as required by the CPT theorem [Lud54], and that the CPT operator applied to the baryon number reverses its sign. This derivation is taken from [TC04]. Both of the aforementioned models (electroweak baryogenesis and leptogenesis) contain

out-of-equilibrium processes, though not without some caveats. For electroweak baryogenesis, whether this condition is met depends on whether the electroweak phase transition is first- or second-order, with the Sakharov condition only being satisfied in the former case [TC04]. This poses a problem, as lattice simulations indicate that the transition is only first order for Higgs mass $\lesssim 80$ GeV [Far+94] — far below the value of 126 GeV measured at the Large Hadron Collider [Aad+12]. With leptogenesis, the situation depends on how the rate of the lepton-asymmetry generating interactions compares with the expansion rate of the universe at that time, a more detailed discussion of which can be found in Chapter 4 of [DNN08]. In either case, it is clear that both proposals contain out-of-equilibrium interactions, but it is a model-dependent question whether this departure from equilibrium takes place in a way that could result in the asymmetry of the magnitude we see today.

2.4 C- and CP-Violation

Symmetries are a key concept in our modern understanding of physics. We often learn about symmetries when applied to geometric shapes; in this context, a symmetry is a transformation that, when applied to an object, leaves that object unchanged. For example, if we rotate a square clockwise by 90 degrees about its center, this leaves the square in the same position as when we started. An object may have multiple symmetries; for example, rotation by 180 degrees, or reflection across a diagonal are also symmetries of the square. The concept of a symmetry in physics is more general, as here we can apply transformations to the laws of physics themselves. Symmetries come in two types: continuous and discrete. Continuous symmetries include those such as time translation, spatial translation, or spatial rotation. In the Standard Model of particle physics, three discrete symmetries play a major role. These are charge conjugation (C) which replaces every particle with its antiparticle; parity reversal (P) which reverses the direction of vector quantities; and time-reversal (T) which reverses the direction of time. Regardless of what form baryon number violation and out-of-thermal equilibrium dynamics take in our model of baryogenesis, C- and CP-asymmetry are needed so that interactions which produce a baryon excess are not canceled out by a reverse interaction.

2.4.1 Symmetry Breaking in the Standard Model

So the Sakharov conditions mandate that there be a violation of discrete symmetry in the early universe. Could this symmetry violation be found in the Standard Model? There is currently no experimental evidence of CP violation in the electro-

magnetic force. The same is true for the strong force as well, though the situation is more interesting. Theoretically, quantum chromodynamics (QCD) allows for a CP-violating term, whose strength is characterized by a parameter $\bar{\theta}$. However, this parameter is heavily constrained by measurements of the neutron electric dipole moment (nEDM). Naively, the $\bar{\theta}$ term in the QCD Lagrangian affects the way the constituent quarks are arranged inside of the neutron, leading it to develop a dipole moment. Theoretical calculations [PR00] show that this EDM is

$$d_n = \bar{\theta} \times 10^{-16} \text{e cm}. \quad (2.5)$$

With current limits on the nEDM being $d_n \lesssim 10^{-26} \text{e cm}$, this constrains $\bar{\theta}$ to be less than of order 10^{-10} [PR00]. This is extremely small for a parameter that one would naturally expect to be of order 1. This unnaturalness is known as the *strong CP problem*, and has spurred advancements in the study of axions to attempt to resolve it [Hoo19]. However, as of yet strong sector CP violation does not appear sufficient to resolve the baryon asymmetry. Weak force parity violation was discovered experimentally as early as 1957, when C.S. Wu measured the beta decay asymmetry in the decay of spin-polarized ^{60}Co [Wu+57], and shortly thereafter Garwin, Lederman, and Weinrich [GLW57] found that in the decay $\pi^+ \rightarrow \mu^+ \nu_\mu$, the muon was preferentially polarized along its direction of travel. These experiments indicated that the weak interaction violated P, but for some time it was conjectured that the weak interaction might be invariant under the joint transformation CP. It was not until the work of Cronin and Fitch [Chr+64] that we had experimental confirmation of CP violation in the weak interaction, specifically in the kaon system. Essentially, they found that the weak eigenstate K_L , thought to be a CP-odd eigenstate, could occasionally (with probability around 1/500) decay into the CP-even final state of two pions. This implied that K_L was actually a superposition of even and odd CP eigenstates, proving that CP was indeed violated. The experimentally measured CP violation in the quark sector is expressed in the Standard Model through a complex phase in the Cabibbo-Kobayashi-Maskawa (CKM) matrix describing the mixing between the quark flavors. However, it has been shown that the baryon asymmetry generated by the CKM matrix is orders of magnitude lower than what is observed cosmologically [Gav+94].

2.5 Electric Dipole Moments (EDMs) and CP Violation

The current situation is thus: we know there must be CP violation in order to explain the baryon asymmetry of the universe (BAU), but calculations show that

the CP violation in the Standard Model (through the CKM matrix) is insufficient to generate the observed asymmetry. In light of this, a great variety of extensions to the Standard Model have been proposed, such as supersymmetry (SUSY) [BLS11], extended Higgs sector [Bra+12; OOR18], and Grand Unified Theories (GUTs) [Ber02]. From the perspective of an experimentalist, these theories pose a challenge, as they necessarily propose new interactions and particles at energy scales exceeding what we can probe directly with current particle colliders. Fortunately, many theories of beyond-the-Standard-Model (BSM) CP violation result in particles such as the neutron or electron acquiring an EDM much larger than what the Standard Model would predict. For reference, the nEDM resulting from the CKM matrix CP violating phase is of the order 10^{-31} e cm, while the current upper bound on d_n is 10^{-26} e cm, measured at the Paul Scherrer Institute (PSI) [Abe+20]. To see that an EDM of an elementary particle violates CP, we need only examine the Hamiltonian for this system. A neutron in a static electric and magnetic field has a Hamiltonian of the form

$$H = -\frac{2d_n}{\hbar}(\vec{S} \cdot \vec{E}) - \gamma(\vec{S} \cdot \vec{B}), \quad (2.6)$$

where d_n is the nEDM, E and B are the electric and magnetic fields, respectively, and S is the neutron spin vector. The term $(\vec{S} \cdot \vec{E})$ is T-odd ($\vec{S} \rightarrow -\vec{S}$, $\vec{E} \rightarrow \vec{E}$), and so by the CPT theorem [Lud54], which says that the combined symmetry CPT holds for all systems, this term must also violate CP. Thus observation of a nEDM would be clear evidence of CP violating BSM physics, and conversely non-observation of a nEDM would rule out certain extensions of the Standard Model — in fact, nEDM measurements have already excluded a great number of such theories. As said by Golub and Lamoreaux in 1994: “It is fair to say that the neutron EDM has ruled out more theories (put forward to explain K_0 decay) than any experiment in the history of physics” [GL94]. Because of the promise these experiments hold in expanding our understanding of how matter came to be, much effort has been expended over the past decades to make increasingly precise measurements of the nEDM. In the next section we detail some past, present, and future endeavors, with particular focus on those making use of ultracold neutrons (UCNs).

Chapter 3

MEASURING AN NEDM

3.1 Ultracold Neutrons

3.1.1 Why UCN?

Improving the precision of a neutron electric dipole moment (nEDM) experiment is a complex affair, but in essence there are three primary variables to consider:

1. The strength of the electric field. Higher is better.
2. The number of neutrons. More is better.
3. The length of measurement. Longer is better.

These considerations are summarized in the relation

$$\sigma_{d_n} \propto \frac{\hbar}{E\tau\sqrt{N}} \quad (3.1)$$

Where E is the magnitude of the electric field, τ is the measurement time, N is the number of neutrons, and we ignore $O(1)$ constants. The advantage provided by ultra-cold neutrons (UCN) is that they can be stored in materials, thus allowing for a measurement of the nEDM that is constrained only by the inherent β -decay lifetime of the neutron (roughly 900 seconds). Contrast this with a beam experiment, where a typical cold neutron beam may have neutron velocities of order 400 m/s. In this case only very short observation times are possible, and τ is limited by the length of the beamline. Another less obvious benefit of using UCN is that it limits the systematic shift caused by the motional magnetic field [LG09]. This is due to two effects — first, the motional magnetic field is proportional to $\mathbf{v} \times \mathbf{E}$, and so the lower velocities involved with UCN diminish the strength of this term; second, if the neutron trajectories are symmetrically distributed in velocity (as would be the case in a UCN container), then the motional field experienced by neutrons going one way will tend to cancel the field experienced by those traveling in the opposite direction. This cancellation is incomplete, but it still helps to suppress this systematic effect. Beam experiments have other advantages, such as having a large neutron flux, and not needing a dedicated cryogenic storage vessel for the neutrons. Nonetheless, the appeal of ultra-cold neutrons for nEDM measurement is clear.

Material	V_{eff} (neV)
Ni ⁵⁸	335
BeO	261
Ni	252
Be	252
PTFE (Teflon)	123
dPMMA (deuterated Plexiglas)	180
dTPB (deuterated TPB)	157

Table 3.1: Effective Fermi potentials for some UCN bottle materials. Values for non-deuterated materials are taken from [R G91]. Values for dPMMA and dTPB are calculated based on the scattering lengths given in [Sea92] using equation 3.2

3.1.2 What is a UCN, Anyway?

The defining property of the UCN is that it can be contained in a material bottle. This means that at the material surface the UCN must experience total reflection (i.e., reflection at any glancing angle θ). For most materials, this amounts to speeds less than around 8 m/s, or kinetic energies less than around 300 neV [Sun22]. At these speeds, gravity becomes relevant to the neutron's motion, as the gravitational potential is approximately 1 neV/cm.

3.1.3 Reflection from Surfaces

The scattering of neutrons from matter is a rich subject, one which we cannot hope to cover here. However, for UCN in particular, several simplifications are possible. The fact that the neutron de Broglie wavelength is long relative to the nuclear radius implies that we should primarily consider s-wave scattering, and the fact that the wavelength is also longer than the distance between neighboring atoms in a material means that we need not be concerned with diffraction effects. The net result of these simplifications is that in many cases a homogeneous material looks to an incoming UCN like an effective potential (also known as the Fermi Potential) [R G91]

$$V(r) = \frac{2\pi\hbar^2}{m} [an(r)] \quad (3.2)$$

where m is the neutron mass, a is the bound atom coherent scattering length, and $n(r)$ is the local density of the scatterer. This potential can alternately be viewed as a wavelength-dependent index of refraction, and the phenomenon of neutron reflection can then be seen as a kind of total internal reflection. Some common materials for containing UCN and their Fermi potentials are shown in table 3.1.

3.2 UCN Production

Several methods can be used to produce UCN. Neutron sources come primarily in two varieties. The first of these are spallation sources, such as those at Los Alamos, Argonne, Oak Ridge, or the European Spallation Source (ESS) currently under construction. The second type are those that produce neutrons by fission in a nuclear reactor. Whatever the method of production, neutrons are generally produced with energies in the MeV range [R G91], and so must first be cooled before UCN can be effectively produced. This is usually accomplished by a material called a moderator, whose function is to repeatedly scatter the neutrons until thermal equilibrium is reached. Common material choices for this include liquid hydrogen and deuterium at 20K [R G91]. Ideally, the neutrons exiting the moderator would be at thermal equilibrium with the moderator, though in general this will not be the case, and the neutron velocity distribution following moderation will be determined by a complicated function of the moderator condition and its scattering cross section [R G91]. Nonetheless, once the neutrons have been cooled by the moderator, UCN may be extracted in various ways. Some techniques include sending the neutrons up a vertical tube, thus reducing the neutron velocity, or bouncing neutrons off of moving surfaces. The UCN source at the Institut Laue-Langevin [Ste+86] employs both of these techniques, having the cold neutrons exiting the moderator navigate a curved vertical pipe, followed by a chamber with rotating blades. While these methods can modify the phase space distribution of the neutrons to a form more favorable for UCN experiments, they cannot increase the phase space density above that which was originally present in the moderator. This is a consequence of Liouville's theorem, which states that phase space density is constant along particle trajectories [LL80]. This limit was calculated by [R G91], and is given by

$$\rho_{UCN} = \frac{2}{3} \frac{\Phi_0}{\alpha} \left(\frac{V}{k_B T} \right)^{3/2} \quad (3.3)$$

where T is the moderator temperature, $\alpha = \sqrt{2k_B T_n/m}$ with T_n being the neutron velocity pre-moderation (α is $\sqrt{2/3}$ times the neutron root mean squared speed), V is the potential of the container walls, and Φ_0 is the neutron flux. In practice, neutron losses will lead to UCN densities far below the value given in 3.3. To increase the UCN density further, we must employ so-called superthermal UCN production. Essentially, we need to engineer a situation in which thermal neutrons can scatter into UCN, but the rate for the various loss processes (up-scattering, absorption, wall losses, etc.) can be made very small. In such a case, the density of UCN will build up until the production rate equals the loss rate. An ideal candidate




Figure 3.1: Phonon dispersion curve (blue) [DB98] and free neutron dispersion curve (orange). UCNs with energy near the intersection can be brought to rest by scattering a single phonon.

for in-situ production of UCN is superfluid ^4He . In superfluid ^4He , the primary excitations under consideration are phonons, whose dispersion relation is shown in figure 3.1 (rotons can also be excited in the fluid, but these do not play a role in this technique). Conservation of energy and momentum dictate that a neutron at rest can only scatter a phonon if that phonon has a particular momentum and energy, given by the intersection between the free neutron energy and the phonon dispersion curves. This momentum turns out to correspond to 8.9 angstrom neutrons. Because the number density of these phonons decreases exponentially as temperature decreases [Wil67], the upscattering rate can be made almost arbitrarily small provided that the superfluid bath can be sufficiently cooled. In addition, ^4He has virtually no neutron absorption cross section, and has a lower UCN potential much smaller than most wall materials [R G91]. This is the motivation behind nEDM@SNS's in-situ UCN production scheme. At the Spallation Neutron Source (SNS), nominal parameters for the experiment would have achieved UCN densities of up to 150/cc, or about 9×10^5 total UCN [Ahm+19].

3.3 The State of nEDM Experiments

Figure 3.2 shows the progress in nEDM measurements over the past several decades. It is interesting to note that following the 1970s, nEDM measurements transitioned from primarily beam experiments (applying an electric field to a beam of cold neutrons) to UCN-based experiments. This is in large part due to a systematic




Figure 3.2: nEDM 90% confidence level upper bound over time. The current limit is shown as a dotted line. The Standard Model nEDM (not shown) is below the bottom of the y-axis, at $d_n = 10^{-31} ecm$. Plot is taken from the 2022 Snowmass proceedings [Ala+22].

effect known as the geometric phase effect, the $\mathbf{v} \times \mathbf{E}$ (pronounced “v cross E”) effect, or the linear-in-E frequency shift. Essentially, neutrons moving through a static electric field experience an effective magnetic field, and this magnetic field — being linear in electric field strength — can mimic an electric dipole moment. This effect is discussed further in the next chapter, as even though this phenomenon is mitigated in UCN experiments due to the lower velocities involved, it is still the dominant systematic effect. The current generation of nEDM experiments are either in R&D or currently underway at neutron sources around the world. These efforts are summarized in table 3.2.

3.4 nEDM@SNS Experiment Overview

The (former) neutron electric dipole moment experiment at the Spallation Neutron Source (nEDM@SNS experiment) proposes to measure the nEDM using the spin-dependent capture cross section of neutrons on ^3He .

3.4.1 Measurement Principle

nEDM@SNS uses nuclear magnetic resonance (NMR) to measure the precession frequency of the neutron in a static magnetic and electric field. Recall that the Hamiltonian for a spin-1/2 particle in these fields is given by 2.6. If the electric

Experiment : Facility	90% C.L. after 300 live-days (10^{-28} e cm)	Year for acquisition (estimated)
Crystal: JPARC	< 100	R&D
Beam: ESS	< 50	2030
PNPI: ILL	< 10	R&D
n2EDM: PSI	< 15	2026
PanEDM: ILL/Munich	< 30	2026
TUCAN: TRIUMF	< 20	2027
nEDM: LANL	< 30	2026
nEDM@SNS: ORNL	< 3	2031

Table 3.2: Proposed nEDM experiments, their approximate 90% confidence level upper bound on the nEDM after 300 live-days of measurement, and their estimated year for finishing data taking for this precision. Values taken from [Fil22].

and magnetic field are parallel (we use the convention where these fields are applied along the $\hat{z}$ axis) and the neutron's spin vector is oriented perpendicular to these fields, then this Hamiltonian results in the neutron spin vector precessing about the electric and magnetic fields, with frequency given by

$$\omega_0(E, B) = \frac{2d_n}{\hbar}E + \gamma B. \quad (3.4)$$

The basis of an nEDM measurement is thus to look for a slight change in frequency proportional to the electric field strength (in practice, one often reverses the direction of the electric field in order to control for systematics, thus the relevant quantity is $\omega_0(+E, B) - \omega_0(-E, B) = \frac{4d_n}{\hbar}|E|$, which is a quantity that is directly proportional to the EDM d_n). In most EDM experiments, this frequency measurement uses the Ramsey technique, also known as the separated oscillatory fields technique [Ram50]. With this method the spins start parallel to $\vec{B}$ and are tipped into the plane of precession by application of a $\pi/2$ pulse (the “oscillatory field” in the oscillatory fields method). Then after waiting for a period of time T , a second $\pi/2$ pulse is applied, and the polarization of the spins along the z axis is measured. By repeating this experiment for different waiting times T , one can build the characteristic Ramsey fringes, from which the frequency can be inferred. nEDM@SNS uses a modified version of the Ramsey technique owing to its unique method of observing the neutron spin direction. With nEDM@SNS, a small number of ^3He atoms (concentration of $\sim 10^{-10}$ relative to ^4He [Ahm+19]) occupy the same space as the neutrons. These ^3He atoms serve two purposes. First, they serve as a co-magnetometer. Because the EDM interaction we seek to measure is so small compared to the magnetic field interaction, we need very precise ways to stabilize and measure the magnetic field

at the cell. ^3He has virtually no electric dipole moment owing to shielding of the nucleus by its orbital electrons, and so its precession frequency will depend only on the ambient magnetic field. Second, ^3He has a large neutron capture cross section, and this cross section is highly spin-dependent. The capture rate is given by

$$\Gamma \propto 1 - P_n P_3 \cos(\theta_{n3}) \quad (3.5)$$

where P_n and P_3 are the neutron and ^3He polarization, respectively, and θ_{n3} is the angle between their spin vectors. From this equation we can see that the capture rate is maximized when the spins are antiparallel and minimized when they are parallel. The gyromagnetic ratio γ_3 of ^3He is roughly 1.1 times the gyromagnetic ratio of the neutron γ_n ($\gamma_3 = -2.03 \times 10^8 \text{ T}^{-1} \text{ s}$ [Ahm+19], $\gamma_n = -1.83 \times 10^8 \text{ T}^{-1} \text{ s}$ [Gre+79]), so when the two spins precess in the same magnetic field, the ^3He will precess roughly 10 percent faster. If we now monitor the capture rate, then we will observe a signal which oscillates at a rate equal to the difference in precession in frequency between the two species. Since ^3He has negligible EDM owing to Schiff screening [LE07], change in this difference frequency proportional to an applied electric field would be indication of a neutron EDM. This measurement technique is distinct from the Ramsey method, as the neutron spins are continuously probed by the presence of the ^3He atoms, rather than the neutron spins being probed only at the end of the measurement period. This mode of operation is termed “free precession mode.” To analyze data from free precession mode, the capture events can be binned into time bins, and the resulting histogram can be fitted to a decaying sine wave (decaying as neutrons undergo beta decay, lost to walls, absorbed by ^3He , etc.). Comparing the fitted frequency between events with reversed electric field then yields a result proportional to the nEDM.

3.4.1.1 Critical spin dressing (CSD) mode — not frequency, but asymmetry

nEDM@SNS also has a secondary mode of operation using a technique called spin dressing. An off-resonant sinusoidal magnetic field with amplitude B_1 and frequency ω is applied, which “dresses” the gyromagnetic ratios of the precessing species (the neutron and ^3He). This situation has been analyzed theoretically [C69], and the factor by which the gyromagnetic ratio is modified was found to be given by equation 3.6

$$\gamma' = J_0(x)\gamma \quad (3.6)$$

$$x = \frac{\gamma B_1}{\omega}, \quad (3.7)$$

where J_0 is the Bessel function of order zero, γ is the undressed gyromagnetic ratio, γ' is the dressed gyromagnetic ratio, and x is termed the dressing parameter. Note that the dressing parameter will differ between the two species, and we use subscripts n and 3 to differentiate between the two species; for example, x_n refers to the neutron dressing parameter, while γ_3 refers to the ^3He gyromagnetic ratio. Because these parameters differ, there exist values of B_1 and ω where the effective gyromagnetic ratios γ'_3 and γ'_n are equal. The measurement mode in which such a magnetic field is applied is termed critical spin dressing (CSD) mode. This scenario has a variety of advantages for nEDM measurement. First, key systematic effects, in particular the geometric phase effect [Pen+04b; R G08], differ between free precession mode and CSD mode. Having a second measurement technique allows for a more robust measurement. In addition, CSD gives us better control over the relative angle between the two species. This is important because not all angles are equal in terms of sensitivity to the nEDM. As an example, consider a case where the neutron and ^3He spins are parallel. In this case, the capture cross section is negligible, and so the uncertainty due to counting statistics will be large. If the spins are antiparallel we get maximum counting statistics, but we have a different problem: in this configuration the system is insensitive to small changes in θ_{n3} , because $\cos(\theta_{n3})$ has a local extremum at $\theta_{n3} = \pi$. Furthermore, we must consider the fact that when counting statistics are maximized, so are neutron losses to the capture reaction. From balancing these three considerations — counting statistics, phase sensitivity, and UCN losses — we can determine the optimal angle at which to set θ_{n3} . This angle was found to be $\theta_{n3} = 0.48$ by Ahmed et al. [Ahm+19], and the corresponding statistical sensitivity was found to be nearly twice that of free precession mode ($\sigma_{d_n} = 1.6 \times 10^{-28}$ e-cm versus 3.3×10^{-28} e-cm). With CSD, the angle between neutron and ^3He can be maintained at this angle, thus maximizing sensitivity to the nEDM.

Whereas in free precession mode, the nEDM measurement is a frequency measurement, in spin dressing mode it is an asymmetry measurement. Recall that in critical dressing mode, the neutron and ^3He spins precess at the same rate. So assuming zero nEDM and neglecting neutron losses, the capture rate signal would then be constant in time. The measurement scheme proposed in [Ahm+19] is to modulate the amplitude and/or frequency of the dressing field with short pulses, such that the angle between the two spin species alternates between $+\phi_0$ and $-\phi_0$. If we now

reintroduce an nEDM d_n , then under a dressing field of amplitude B_1 and angular frequency ω , the angle θ_{n3} between the two species would behave as

$$\theta_{n3}(t) = \pm\phi_0 + \frac{2J_0(x)d_nEt}{\hbar}, \quad (3.8)$$

where $J_0(x_n)$ is the Bessel function of order zero evaluated at $x = \gamma_n B_1/\omega$, and is the factor by which the dressing field modifies the neutron gyromagnetic ratio. The capture rate signal is proportional to $1 - \cos(\theta)$, so for small d_n we find that

$$1 - \cos(\theta_{n3}(t)) \approx 1 - \cos(\phi_0) \pm \frac{2J_0(x)d_nEt \sin(\phi_0)}{\hbar}. \quad (3.9)$$

The key point here is that the signal is asymmetric, with an asymmetry that grows linearly with time at a rate proportional to d_n . In [Ahm+19], it is assumed that the uncertainty of this nEDM estimate would be dominated by counting statistics. In section 4.3 we will additionally consider how the presence of time-dependent magnetic field fluctuations, for example fluctuations from the spin dressing power supply, affect the uncertainty and data analysis strategies.

3.4.2 Overview of the Apparatus

The experiment consists of several layers, much like a Matryoshka doll. I will describe the innermost components first, as this is where the neutron EDM is measured. The central detection system (CDS) contains a superfluid bath of ^4He , with a small ($\sim 10^{-10}$) concentration of ^3He . Ultracold neutrons (UCN) will be produced by single-phonon downscattering of an $8.9 \text{ \AA}$ beam of cold neutrons.

3.4.3 UCN Storage

The UCN must be stored in a vessel with a high Fermi potential and low neutron loss cross section to ensure long storage times. In addition, the walls should be non-magnetic so as not to disturb the homogeneity of the magnetic fields. To meet these stringent requirements, nEDM@SNS will use a pair of PMMA (Plexiglass) cells, each with dimensions $40 \text{ cm} \times 10 \text{ cm} \times 7.5 \text{ cm}$. The walls of the cells will be coated with a thin layer of deuterated tetraphenyl butadiene (dTPB) in a deuterated polystyrene (dPS) matrix [Ahm+19]. dPS has a high Fermi potential, and dTPB is a wavelength-shifting dye which shifts the wavelength of the ^4He scintillation light from 80 nm to a more manageable 400 nm for detection by the silicon photomultipliers (SiPMs) [Ahm+19]. Here deuterated materials are necessary, as hydrogen has a high cross section for UCN upscattering, leading to unwanted losses. Minimizing UCN losses is doubly important for nEDM@SNS, as a high loss rate limits the




Figure 3.3: Full-scale prototype cell illuminated under UV light. Image from [Leu25].

UCN density that can be accumulated at the start of a measurement run, as well as causes UCN to be lost over the course of the ~ 1000 s measurement time. For these reasons it is imperative that these losses be characterized and measured. To this end, storage time measurements have been performed with a full-scale cell prototype, shown in figure 3.3, to measure the UCN loss time with these materials. The most recent value for storage lifetime was found to be (including losses due to beta decay) $\tau_{\beta+\text{cell}} = 570 \pm 20$ s, which is near the design goal of $\tau_{\beta+\text{cell}} = 610$ s [Leu25].

Once inside the nEDM@SNS cryostat, the pair of measurement cells will be filled with superfluid liquid ^4He and cooled to a temperature $\sim 0.4\text{K}$, a choice which is the result of attempting to optimize for several competing factors. For one, the phonon density in the superfluid increases sharply with temperature, meaning that lower temperatures will reduce UCN losses due to two-phonon scattering [Wil67]. On the other hand, a key systematic effect for this experiment is the false EDM caused by the presence of a magnetic field gradient coupled with an electric field. This effect depends on the temperature of the superfluid through the mean free path of the ^3He atoms between phonon collisions, but for reasonable parameters there exists a temperature at which the false EDM is zero [A L06]. It should be noted though, that because the exact behavior of this shift relies on unknown parameters — such as the surface conditions of cell, it is likely that studies will need to be conducted on the completed cell in order to determine the optimal temperature.

3.4.4 Signal Detection

As noted in section 3.4.1, the rate of capture, and thus the rate of scintillation photons following the capture reaction, is proportional to $(1 - \cos(\theta_{n3}))$, where θ_{n3} is the angle between the neutron and ^3He spin vectors. We thus need a detection scheme to measure the instantaneous rate of this reaction. The capture reaction of neutrons on ^3He releases 764 keV of energy into the superfluid bath ($n + ^3\text{He} \rightarrow p + ^3\text{H} + 764 \text{ keV}$), producing several thousand photons with wavelength approximately 80nm [McK+04]. This is in the extreme ultraviolet (EUV) range, and would thus be absorbed by almost any window material. Thus, the dTPB coating's primary purpose is to act as a wavelength-shifting dye, downconverting the 80nm scintillation light to 400nm blue light [Ahm+19]. The light emitted is then captured by fiber optic cables positioned on the sides of the measurement cells, from which the photons are detected by silicon photomultipliers (SiPMs).

3.4.5 Electric Field System

Because sensitivity to an nEDM is proportional to the strength of the electric field, it is beneficial to have the largest electric field magnitude possible. Here again, the properties of liquid helium work in our favor, as the breakdown voltage of helium is expected to be much higher than that of vacuum. With liquid helium, techniques to maintain electric fields up to 100 kV/cm are being studied in various smaller-scale apparatuses [Pha23]. Even if breakdown is prevented, producing such a field in the first place also poses challenges. If high voltage is produced outside the cryogenic system, a feedthrough will be needed. However, such a feedthrough would need to be electrically insulated from the surrounding apparatus, while also being in thermal contact with it. To circumvent this, nEDM@SNS will use a more modest voltage outside the cryogenic system, and use a mechanical voltage multiplication system called Cavallo's multiplier [Cla+18] inside the cryogenic system to raise the voltage to the desired value. Cavallo's original design dates back to 1795 [Cav95], and uses a moving plate to transfer charge between a (relatively) low voltage source to a fixed pair of capacitor plates. This mechanism, in the proposed experimental design for nEDM measurement, is shown in figure 3.4. The ground electrode surrounds the two cells and, by repeated action of the moving electrode, charge is accumulated on the central high-voltage electrode.




Figure 3.4: 3D COMSOL simulation of the Cavallo voltage multiplier in the nEDM@SNS central volume, in which the central electrode is charged to 650 kV. Figure copied from [Cla+18].




Figure 3.5: Schematic of the inner magnet volume (IMV). G10 frame not shown.

3.4.6 Magnetic Field System

A diagram of the magnetic field system, comprised of the B_0 coil, dressing coil, shielding, and associated cryogenic vessel, is shown in figure 3.5. Within the inner magnet volume (IMV), the outermost layer is the superconducting lead shield, which is comprised of the side lead shield, the lead end caps, and the lead end plate. Combined, these form a near-hermetic barrier which insulates the inner layers from external magnetic field fluctuations. The lead end caps have an additional role,




Figure 3.6: Partially wound B_0 coil with G10 frame. Image from [Sun22].

as these act as magnetic mirrors, helping to maintain magnetic field uniformity; essentially, this coil configuration mimics a coil of infinite length[Ahm+19]. The B_0 magnet itself is designed as an optimized $\cos(\theta)$ coil, with wires running vertically. This coil must also be superconducting, as this ensures that its current is near-constant in time when operating in persistent current mode. This coil provides a uniform ~ 30 mG field that is orthogonal to the neutron beam direction, and which serves as a holding field in which the neutrons and ^3He precess. This magnetic field strength corresponds to a precession frequency of around 100 Hz for ^3He , or around 90 Hz for neutrons, values which were chosen to mitigate systematic frequency shifts such as those studied in [Pig+15; Pig19], which depend on the Larmor frequency. Just outside of the B_0 coil is the Metglas flux return, which serves to further improve uniformity. The B_0 frame, which holds the wires for the magnet in place, must be non-magnetic and capable of withstanding multiple cycles of cooling and warming. We use G10 (also known as G-10 or garolite), a type of fiberglass, for this purpose due to its low thermal contraction and known behavior at cryogenic temperatures. Due to limitations on the size of G10 pieces that could be cut by the manufacturer, each of the horizontal hoops of the frame needed to be constructed as twelve circular arcs, which were then glued together with epoxy over the course of several months. For more information on the construction process (and how it was done at Caltech during a global pandemic), see [Sun22]. A picture of the completed B_0 frame, with wire partially wound, is shown in figure 3.6. Interior to B_0 , and rotated at a 90-degree angle so that its field points along the beam direction, is the dressing coil. This




Figure 3.7: COMSOL model of the dressing coil used to simulate Joule heating due to the dressing field.

coil serves two purposes: first, it provides the $\pi/2$ pulse which tips the neutron and ^3He spins into the plane of precession. Second, it supplies an oscillating 1 kHz sine wave field with amplitude approximately 300 mG. The amplitude and frequency are chosen so as to cause the neutrons and ^3He precess at the same frequency [Swa+18]. This is known as the critical dressing condition. A section of a COMSOL model of this coil is shown in figure 3.7. The presence of the dressing field imposes a variety of constraints on other systems in the apparatus — notably, the oscillating field induces eddy currents in conductive materials, leading to Joule heating. To help mitigate this, an eddy current shield made of copper-clad G10 is included between the B_0 and dressing coils to reduce heating in the Metglas flux return. The Joule heating has been modeled extensively in COMSOL by C.M. Swank, and a sheet of copper of thickness $34\text{ }\mu\text{m}$ was found to be sufficient to reduce the heat below the heat budget for the cryogenic system.

Fabrication of the dressing frame and coil is similar to that of the B_0 coil, except that whereas the B_0 coil used to springs to keep the wires under tension, in the dressing coil each vertical wire sits in a channel cut into a piece of G10. The construction and winding of the dressing coil has recently been completed, with field uniformity tests underway.

Chapter 4

SPIN DRESSING THEORY

Given the importance of critical spin dressing for the sensitivity of the nEDM@SNS experiment, it is vital to understand and address the effect that magnetic field noise and field inhomogeneities have on this measurement mode. This chapter is split into three sections. The first section (section 4.1) is adapted from a published paper [TS22] studying the effects of noise on the critical dressing mode of the nEDM@SNS experiment, as well as developing ways to mitigate these effects. Beginning in section 4.2, we develop an additional technique termed “magic dressing” which can be used to measure an nEDM in a way that is insensitive to fluctuations in the dressing field amplitude. In section 4.3, we conclude with a discussion of data analysis techniques that may be employed to analyze spin dressing data in the presence of magnetic field noise.

4.1 Quantum Control of Two Critically-Dressed Spin-1/2 Species

When a particle’s spin precessing in a static magnetic field is exposed to an off-resonant oscillating orthogonal magnetic field, the effective gyromagnetic ratio of the particle — and thus its precession frequency — is modified. This phenomenon is known as spin dressing and was first studied theoretically in reference [PC65]. Spin dressing is found to have applications across multiple sub-fields in quantum information and fundamental physics, including increasing coherence time of free induction decay (FID) in relatively large magnetic field gradients [Bev+19; Swa+18]. The use of multiple dressing frequencies tuned to extend the coherence of atomic clocks in field gradients was first reported in reference [KS15], whereas the authors of reference [BIS18] used this technique to mitigate Stark shifts. With multiple dressing field directions, enhanced spin manipulation for quantum information is achievable, and can even accelerate the effective Larmor precession; this is discussed in references [Bev+20; Bev+22]. In reference [GL94], the authors found that critical spin dressing (CSD), the simultaneous dressing of two spin species to the same Larmor frequency, can be applied for a very sensitive measurement of the neutron electric dipole moment (nEDM). A detailed investigation in reference [Ahm+19] found that this technique would halve the statistical uncertainty of an nEDM measurement compared to an analogous measurement using FID. Additionally, by modulating the

parameters of CSD as described in references [GL94; Swa+18], systematic uncertainty from phase fluctuations not associated with noise in the spin dressing field (e.g., external magnetic field drifts) is significantly reduced even without correction from a comagnetometer. For example, with modulation at the angular frequency ω_m phase accumulation from slow field drifts in the apparatus will be modified by a factor $\sim 1/\sqrt{\omega_m}$ [Pap65].

In the nEDM@SNS experiment ultracold neutrons (UCN) are confined to a pair of acrylic (PMMA) measurement cells with dimensions $L_x \times L_y \times L_z = 40 \times 10.2 \times 7.6$ cm in solution with polarized ^3He in superfluid ^4He . A large electrode sits between the two cells which are sandwiched by ground electrodes on either side and provide a large electric field of order 7.5 MV/m between each cell in the $\hat{z}$ direction so the electric field strength has opposite sign for each cell. A superconducting persistent ‘ $\cos \theta$ ’ coil provides a static uniform magnetic field of magnitude $B_0 \geq 3 \mu\text{T}$ along the $\hat{z}$ direction and is the same in each cell. Inside the static field is a separate audio frequency ‘ $\cos \theta$ ’ coil which provides a dressing field with amplitude of up to $100 \mu\text{T}$ oriented along the $\hat{x}$ direction. In between the static uniform field and the dressing field is a copper shield, which limits interaction between the dressing coil and the static field coil and acts as a flux return for the dressing field. To determine the relative phase between the neutron and ^3He spins, nEDM@SNS detects the spin-dependent capture rate of neutrons on ^3He through the reaction $n + ^3\text{He} \rightarrow ^3\text{H} + p + 764 \text{ keV}$. The critical dressing mode of the experiment uses a strong oscillating magnetic field to dress gyromagnetic ratios of neutrons and ^3He to the same value. This improves sensitivity to the nEDM by allowing continuous operation at the most sensitive relative phase between the neutron and ^3He [Ahm+19] for a given ^3He concentration. However, this mode of measurement introduces additional noise, as current fluctuations in the spin dressing coil will lead to fluctuating magnetic fields. These fluctuating magnetic fields can cause the spins of neutrons and ^3He to drift apart over time.

Here, we use second-order time-dependent perturbation theory to assess the implications of magnetic field fluctuations for precision measurements performed on dressed spin systems, with particular emphasis on applying these results to mitigate uncertainty in the critical dressing mode of the nEDM@SNS experiment. In section 4.1.1, we introduce the relevant Hamiltonian for spin dressing for a spin- $1/2$ system. In section 4.1.2, we apply time-dependent perturbation theory to this Hamiltonian to compute the phase variance, frequency shift, and run-to-run variance of the capture

rate signal induced by a fluctuating dressing field. We express these results in terms of the power spectrum of these fluctuations. In this section we also discuss the absence of run-to-run signal variance due to magnetic field gradients in either the static uniform field or the dressing field. In section 4.1.3, we confirm these calculations by numerical integration of the Bloch equations, and discuss the implications of these results for the nEDM@SNS experiment. Section 4.1.4 proposes several strategies to mitigate dressing field fluctuations for the nEDM@SNS experiment. Finally, in section 4.1.5, we introduce a new modulation scheme which can be used to combat the dressing field noise addressed in the previous section, as well as enabling precise control of dressed spins in the presence of magnetic field gradients.

4.1.1 Spin Dressing Hamiltonian

A spin- $1/2$ system with gyromagnetic ratio γ precessing in a static uniform field $B_0 \hat{z}$ and a strong oscillating magnetic dressing field $B_1 \cos(\omega t) \hat{x}$ (with $B_1 > B_0$) is described by the Hamiltonian

$$H = \omega a^\dagger a + \frac{\Omega}{2} \sigma_x (a + a^\dagger) + \frac{\omega_0}{2} \sigma_z, \quad (4.1)$$

where σ_x , σ_y , and σ_z are the usual Pauli spin matrices, $a^\dagger$ and a are the raising and lowering operators for the photon field, $\omega_0 = \gamma B_0$ is the undressed frequency of a spin in a static uniform field, defined along $\hat{z}$, and Ω is a coupling constant given by

$$\Omega = \frac{\gamma B_1}{2\lambda^{1/2}}, \quad (4.2)$$

where $\lambda = \langle n \rangle$ is the average photon number. We are interested in how applying a small fluctuating magnetic field $\delta B(t)$ with a known power spectrum would affect the time evolution of an initial state $|\psi_0\rangle$ under this Hamiltonian.

4.1.2 Theory

In the following section, we first outline a general framework to compute the expectation value and variance of operators under small, time-dependent perturbations to the Hamiltonian of section 4.1.1. We then use this framework to evaluate two observables to second order for a magnetic fluctuation $\delta B(t)$ which occurs parallel to the dressing field. These observables are the variance in phase for each spin species, and the frequency shift for each spin species. For an experiment like nEDM@SNS, these results can be used to estimate the sensitivity of dressed systems in the presence of field fluctuations. In deriving these formulae, we make no assumption about the source of the fluctuations, and therefore these results apply

equally whether the random magnetic field arises from current fluctuations in the dressing coil or static magnetic field gradients coupled with the stochastic motion of the spin species throughout the measurement cell. In the latter case, the power spectrum of $\delta B(t)$ may be calculated through the position-position autocorrelation function, as in [SPG16]. In addition to these two observables, we also compute the variance of the phase between two species that are exposed to the same fluctuating magnetic field, as would be the case for current fluctuations. This is useful for nEDM@SNS, as the capture rate of neutrons on helium-3 depends linearly on this phase difference [Ahm+19].

4.1.2.1 Time-Dependent Perturbation Theory

Suppose that at time $t = 0$, a system is in the state $|\psi_0\rangle$, and we wish to compute the expectation value of some operator A_0 at some future time t . Given a time-dependent perturbation $V(t)$ to a Hamiltonian H , The time-evolution operator $T(t)$ to second order in time-dependent perturbation theory is given by

$$T(t) = 1 - i \int_0^t dt' V_I(t') - \int_0^t \int_0^{t'} dt' dt'' V_I(t') V_I(t'') + O(V^3), \quad (4.3)$$

where $V_I(t) \equiv e^{iHt} V(t) e^{-iHt}$ is the perturbation evaluated in the interaction picture. Thus, we need to compute $\langle \psi(t) | A_0 | \psi(t) \rangle = \langle \psi_0 | T(t)^\dagger A T(t) | \psi_0 \rangle$, where A is likewise evaluated in the interaction picture, i.e., $A(t) \equiv e^{iHt} A_0 e^{-iHt}$.

If $V(t)$ is proportional to $\delta B(t)$, then we can decompose $V_I(t)$ as a sum of complex exponentials as follows:

$$V_I(t) = \left(\sum_j Q_j e^{-i\omega_j t} + Q_j^\dagger e^{i\omega_j t} \right) \delta B(t), \quad (4.4)$$

where Q_j are time-independent operators. Provided that $\delta B(t)$ has zero mean, the terms of $T^\dagger A T$ which are linear in $\delta B(t)$ will vanish when we compute an expectation value over functions $\delta B(t)$. Therefore we need only consider the quadratic terms, which can be decomposed as

$$\sum_{j,k} f(Q_j, Q_k) \int_0^t dt' \int_0^{t'} dt'' e^{-i\omega_j t'} e^{i\omega_k t''} \delta B(t') \delta B(t''), \quad (4.5)$$

where $f(Q_j, Q_k)$ is some operator which is a function of only Q_j , Q_k , and their Hermitian conjugates. We now evaluate the expectation value of the integral over functions $\delta B(t)$.

$$\int_0^t dt' \int_0^{t'} dt'' e^{-i\omega_j t'} e^{i\omega_k t''} \langle \delta B(t') \delta B(t'') \rangle, \quad (4.6)$$

$$= \int_0^t dt' \int_0^{t'} dt'' e^{-i(\omega_j - \omega_k)(t' + t'')/2} e^{-i(\omega_j + \omega_k)(t' - t'')/2} \langle \delta B(t') \delta B(t'') \rangle, \quad (4.7)$$

$$= \frac{1}{2} \left(\int_0^{2t} d\bar{t} e^{-i\Delta\omega\bar{t}} \right) \left(\int_0^{t-|\bar{t}|} d\Delta t e^{-i\bar{\omega}\Delta t} \left\langle \delta B\left(\frac{\Delta t}{2} + \bar{t}\right) \delta B\left(-\frac{\Delta t}{2} + \bar{t}\right) \right\rangle \right), \quad (4.8)$$

$$\approx \frac{1}{2} \left(\int_0^{2t} d\bar{t} e^{-i\Delta\omega\bar{t}} \right) \left(\int_0^\infty d\Delta t e^{-i\bar{\omega}\Delta t} R_{\delta B}(\Delta t) \right), \quad (4.9)$$

where in equation (4.8) we have made the variable substitutions

$$\bar{\omega} \equiv \frac{\omega_1 + \omega_2}{2}, \quad (4.10)$$

$$\Delta\omega \equiv \omega_1 - \omega_2, \quad (4.11)$$

$$\bar{t} = \frac{t' + t''}{2} \quad (\text{integration variable}), \quad (4.12)$$

$$\Delta t = t' - t'' \quad (\text{integration variable}), \quad (4.13)$$

$$R_{\delta B}(\Delta t) = \langle \delta B(\Delta t) \delta B(0) \rangle. \quad (4.14)$$

In equation (4.9), we assume that $\delta B(t)$ is stationary and has a short correlation time compared to t , and so it is valid to replace the limit of the second integral with infinity. The first integral in equation (4.9) only grows with time if $\Delta\omega \neq 0$, and therefore any term where $\Delta\omega \neq 0$ can be neglected. From this, we conclude that for the purpose of noise analysis, it is sufficient to consider the individual frequency components of $V_I(t)$ independently. Therefore, without loss of generality, we can write

$$V_I(t) = (Qe^{-i\omega t} + Q^\dagger e^{i\omega t})\delta B(t), \quad (4.15)$$

$T^\dagger AT$ can now be expressed in terms of Q and A to second order in V_I :

$$T^\dagger AT = A + \int_0^t \int_0^t dt' dt'' V_I(t'') A V_I(t') \quad (4.16)$$

$$\begin{aligned} & - \int_0^t dt' \int_0^{t'} dt'' V_I(t'') V_I(t') A + A V_I(t') V_I(t''), \\ & = A + uu^* (QAQ^\dagger + Q^\dagger AQ) - v(Q^\dagger QA + AQQ^\dagger) \\ & \quad - v^*(QQ^\dagger A + AQ^\dagger Q), \end{aligned} \quad (4.17)$$

where we have defined the integrals

$$u(t; \omega) \equiv \int_0^t dt' e^{-i\omega t'} \delta B(t'), \quad (4.18)$$

$$v(t; \omega) \equiv \int_0^t dt' \int_0^{t'} dt'' e^{-i\omega(t'-t'')} \delta B(t') \delta B(t''). \quad (4.19)$$

We can further simplify equation (4.17) by noting that

$$v + v^* = \int_0^t dt' \int_0^{t'} dt'' \left(e^{-i\omega\Delta t} + e^{+i\omega\Delta t} \right) \delta B(t') \delta B(t''), \quad (4.20)$$

$$= \int_0^t \int_0^t dt' dt'' e^{-i\omega\Delta t} \delta B(t') \delta B(t''), \quad (4.21)$$

$$= uu^*. \quad (4.22)$$

If we now collect the terms proportional to v and v^* , we get

$$T^\dagger AT = A + [Q^\dagger, [A, Q]]v + [Q, [A, Q^\dagger]]v^*. \quad (4.23)$$

We can alternatively choose to collect terms proportional to $\text{Re}(v)$ and $\text{Im}(v)$, in which case we get

$$\begin{aligned} T^\dagger AT &= A + \left([Q, [A, Q^\dagger]] + [Q^\dagger, [A, Q]] \right) \text{Re}(v) \\ &\quad + i[A, [Q^\dagger, Q]] \text{Im}(v). \end{aligned} \quad (4.24)$$

This allows us to calculate the expectation value of an operator A at time t in the interaction picture. In other cases, it may also be useful to calculate the variance of A . In the nEDM@SNS experiment for example, variance in the capture rate signal, which depends linearly on the dot product between neutron and ^3He spins, could increase uncertainties in the nEDM measurement. In this particular case, we are concerned not with the quantum mechanical variance of $\vec{\sigma}_1 \cdot \vec{\sigma}_2$ (which is nonzero even if the magnetic field has no fluctuations), but rather in how the magnetic field perturbs the spin vectors classically. In other words, we wish to find the variance of the quantum mechanical expectation value of $\vec{\sigma}_1 \cdot \vec{\sigma}_2$. We thus define the classical variance of an operator A as

$$\text{Var}_{cl}(A) \equiv \langle |\langle A \rangle|^2 \rangle_{\delta B} - |\langle \langle A \rangle \rangle_{\delta B}|^2, \quad (4.25)$$

where $\langle \cdot \rangle_{\delta B}$ denotes an average over random functions $\delta B(t)$, while $\langle \cdot \rangle$ denotes a quantum mechanical expectation value, i.e., $\langle A \rangle \equiv \langle \psi | A | \psi \rangle$. A similar analysis as was used to derive equation (4.17) yields, in the case that A is Hermitian,

$$\text{Var}_{cl}(T^\dagger AT) = 4 \langle [Q^\dagger, A] \rangle \langle [A, Q] \rangle \text{Re}(v). \quad (4.26)$$

These results can also be applied to a pair of non-interacting spin-1/2 systems through use of the tensor product. Suppose that in the interaction picture the time-dependent perturbations for each of the two spins are given by

$$V_1(t) = (Q_1 e^{-i\omega t} + Q_1^\dagger e^{i\omega t}) \delta B(t), \quad (4.27)$$

$$V_2(t) = (Q_2 e^{-i\omega t} + Q_2^\dagger e^{i\omega t}) \delta B(t). \quad (4.28)$$

If we wish to compute the expectation value $\langle A \rangle$ from an initial state $|\psi\rangle_0 = |\psi_1\rangle_0 \otimes |\psi_2\rangle_0$, then we can apply equation (4.24) or (4.26) by substituting

$$Q \rightarrow Q_1 \otimes I_2 + I_1 \otimes Q_2, \quad (4.29)$$

where I_1 and I_2 are identity operators.

Lastly, we express v in terms of the power spectrum $S(\omega)$ of $\delta B(t)$. The power spectrum is defined as

$$S(\omega) \equiv \int_{-\infty}^{\infty} dt e^{-i\omega t} R_{\delta B}(t). \quad (4.30)$$

From equation (4.9), we have

$$v(t; \omega) = \frac{1}{2} \left(\int_0^{2t} d\bar{t} \right) \left(\int_0^{\infty} d\tau e^{-i\omega\tau} R_{\delta B}(\tau) \right), \quad (4.31)$$

$$= \frac{t}{2\pi} \int_0^{\infty} d\tau e^{-i\omega\tau} \int_{-\infty}^{\infty} d\omega' S(\omega') e^{i\omega'\tau}, \quad (4.32)$$

$$= \frac{t}{2\pi} \int_{-\infty}^{\infty} d\omega' S(\omega') \int_0^{\infty} d\tau e^{i(\omega' - \omega)\tau}, \quad (4.33)$$

$$= \frac{t}{2\pi} \int_{-\infty}^{\infty} d\omega' S(\omega') \left(\pi \delta(\omega' - \omega) - \frac{i}{\omega' - \omega} \right), \quad (4.34)$$

$$= t \left(\frac{1}{2} S(\omega) - \frac{i}{2\pi} \int_{-\infty}^{\infty} d\omega' \frac{S(\omega')}{\omega' - \omega} \right). \quad (4.35)$$

4.1.2.2 Initial State

An oscillating magnetic field corresponds quantum mechanically to a coherent state, denoted by $|\alpha\rangle$. A coherent state is an eigenstate of the lowering operator with eigenvalue α . The average photon number in such a state is given by $|\alpha|^2 = \lambda$, and so we separate α into its magnitude and complex phase as $\alpha \equiv \sqrt{\lambda} e^{i\theta}$. In the appendix, it is shown that a complex phase of α is equivalent to a rotation of the neutron and ^3He spins. The initial state in the following is taken to be

$$|\psi_0\rangle = |\alpha\rangle |s\rangle, \quad (4.36)$$

where $|s\rangle$ is the initial spin state, or if we are considering both the neutron and ^3He spins,

$$|\psi_0\rangle = |\alpha\rangle |s_1\rangle |s_2\rangle. \quad (4.37)$$

4.1.2.3 Computation of Matrix Elements

We now compute the Q operators (and thus $V_I(t)$) for the spin dressing Hamiltonian in the case where the time-dependent fluctuating magnetic field is parallel to $\hat{x}$. This calculation will proceed in two steps. We first apply a carefully chosen time-independent unitary transform to the Hamiltonian in order to compute its eigenstates, and then compute the matrix elements of the operator σ_x with respect to these states. We then transform these matrix elements into the interaction picture by multiplying each matrix element by the appropriate time-dependent phase factor. In the limit where $\lambda \gg 1$ and $\omega \gg \omega_0$, the eigenstates and eigenvalues of the spin dressing Hamiltonian can be found approximately by expressing H as $H = H_0 + H_z$, where

$$H_0 = \omega a^\dagger a + \frac{\Omega}{2} \sigma_x (a + a^\dagger), \quad (4.38)$$

$$H_z = \frac{\omega_0}{2} \sigma_z. \quad (4.39)$$

As shown in [PC65], H_0 can be diagonalized by considering each of the two eigensubspaces of σ_x individually. The Hamiltonians of the two subspaces are given by

$$H_\epsilon = \omega a^\dagger a + \frac{\Omega}{2} \epsilon (a + a^\dagger), \quad (4.40)$$

where $\epsilon \in \{1, -1\}$ corresponds to the eigenvalues of σ_x . The part of the Hamiltonian contained in H_ϵ is diagonalized by applying the displacement operator $D(\epsilon\Omega/2\omega)$, which is defined by its action on the raising and lower operators:

$$D^\dagger(\eta) a D(\eta) \equiv a + \eta, \quad (4.41)$$

$$D^\dagger(\eta) = D(-\eta). \quad (4.42)$$

Explicitly, the displacement operator is given by

$$D(\eta) \equiv e^{\eta a^\dagger - \eta^* a}. \quad (4.43)$$

Note that while the definition of the displacement operator $D(\eta)$ allows for complex η , here we restrict our attention to the case where η is real. Applying this operator to H_ϵ and abbreviating $D(\epsilon\Omega/2\omega)$ as D_ϵ yields

$$D_\epsilon H_\epsilon D_\epsilon^\dagger = \omega D_\epsilon a^\dagger a D_\epsilon^\dagger + \frac{\epsilon\Omega}{2} D_\epsilon (a + a^\dagger) D_\epsilon^\dagger \quad (4.44)$$

$$= \omega \left(a^\dagger - \frac{\epsilon\Omega}{2\omega} \right) \left(a - \frac{\epsilon\Omega}{2\omega} \right) + \frac{\epsilon\Omega}{2} \left(a + a^\dagger - \frac{\epsilon\Omega}{\omega} \right) \quad (4.45)$$

$$= \omega a^\dagger a - \frac{\Omega^2}{4\omega}. \quad (4.46)$$

From D_ϵ we can construct a unitary operator U , which applies D_ϵ to the appropriate eigensubspace of σ_x . Explicitly,

$$U \equiv D \left(\frac{\Omega}{2\omega} \right) |+_x\rangle \langle+_x| + D^\dagger \left(\frac{\Omega}{2\omega} \right) |-_x\rangle \langle-_x|, \quad (4.47)$$

where $|\pm_x\rangle$ are the eigenvectors of σ_x . Applying U to H_0 yields,

$$U H_0 U^\dagger = \omega a^\dagger a - \frac{\Omega^2}{4\omega}. \quad (4.48)$$

We thus see that U diagonalizes H_0 , and that the eigenstates of $U H_0 U^\dagger$ are $|n\rangle |\pm_z\rangle$. Meanwhile, H_z in this basis becomes

$$U H_z U^\dagger = \frac{\omega_0}{2} \left[D^\dagger \left(\frac{\Omega}{\omega} \right) |+_x\rangle \langle-_x| + D \left(\frac{\Omega}{\omega} \right) |-_x\rangle \langle+_x| \right]. \quad (4.49)$$

For large photon number ($\lambda \gg 1$), the displacement operator can be approximated in terms of Bessel functions. Reference [PC65] gives the matrix elements of $D(\eta)$ as

$$\langle n | D(\eta) | n - q \rangle = \langle n | e^{\eta(a^\dagger - a)} | n - q \rangle = J_q(2\eta\sqrt{\lambda}), \quad (4.50)$$

where J_q is a Bessel function of the first kind with order q . From this we derive

$$D(\eta) \approx \sum_{n,q} J_q(2\eta\sqrt{\lambda}) |n + q\rangle \langle n|. \quad (4.51)$$

With this approximation, the full Hamiltonian H in the displaced basis is

$$\begin{aligned}
UHU^\dagger &= \omega a^\dagger a \\
&+ \frac{\omega_0}{2} J_0(x) \sigma_z \\
&- \frac{\omega_0}{2} J_1(x) \left(\sum_n |n\rangle \langle n+1| - |n+1\rangle \langle n| \right) i\sigma_y \\
&+ \frac{\omega_0}{2} J_2(x) \left(\sum_n |n\rangle \langle n+2| + |n+2\rangle \langle n| \right) \sigma_z \\
&+ O(q \geq 3),
\end{aligned} \tag{4.52}$$

where we have defined the spin dressing parameter $x \equiv 2\Omega\sqrt{\lambda}/\omega = \gamma B_1/\omega$. In reference [Swa+18], the matrix elements of various perturbations are calculated in time-dependent perturbation theory. Here, we extend those results to the case where ω_0 may not be small compared to ω by incorporating the first-order correction to the eigenstates of $UH_0U^\dagger$. We first consider the energy correction. Let the eigenstates of $UHU^\dagger$ be denoted by $\widetilde{|n, \pm_z\rangle}$. Treating H_z as a perturbation on H_0 , the energies of these states can be calculated to first order in ω_0/ω using perturbation theory, as in reference [PC65]. The energy shift is simply

$$E_{n,\pm} = E_{n,\pm}^{(0)} + \langle n | \langle \pm_z | H_z | \pm_z \rangle | n \rangle \tag{4.53}$$

$$= n\omega \pm \frac{\omega_0}{2} J_0(x). \tag{4.54}$$

H_z thus lifts the degeneracy of $|n\rangle |\pm_z\rangle$ and leads to an effective Larmor frequency $\omega'_0 \equiv \omega_0 J_0(x)$. We similarly derive the eigenstates of $UHU^\dagger$ to first order in ω_0/ω :

$$\begin{aligned}
\widetilde{|n, \pm_z\rangle} &= |n\rangle |\pm_z\rangle \\
&\pm \frac{J_1(x)\omega_0}{2\omega} (|n-1\rangle |\mp_z\rangle + |n+1\rangle |\mp_z\rangle) \\
&\pm \frac{J_2(x)\omega_0}{4\omega} (|n-2\rangle |\pm_z\rangle - |n+2\rangle |\pm_z\rangle) \\
&+ O(q \geq 3).
\end{aligned} \tag{4.55}$$

We can now calculate the matrix elements of various time-dependent perturbations with respect to these eigenstates. For dressing field noise generated from fluctuations in the current provided by the power supply, the only relevant operator is σ_x , provided

the linearly polarized dressing field is homogeneous. Fortunately, σ_x is unchanged by U , as $U\sigma_x U^\dagger = \sigma_x$. We obtain

$$\begin{aligned} \overline{\langle n, \pm_z | \sigma_x | n', \pm_z \rangle} &= \pm \frac{J_1(x)\omega_0}{\omega} \delta_{1,|n-n'|} \\ &+ O(q \geq 3), \end{aligned} \quad (4.56)$$

$$\begin{aligned} \overline{\langle n, \pm_z | \sigma_x | n', \mp_z \rangle} &= \delta_{n,n'} \pm \frac{J_2(x)\omega_0}{2\omega} \delta_{2,|n-n'|} \\ &+ O(q \geq 3). \end{aligned} \quad (4.57)$$

We have truncated terms containing Bessel functions of order 3 or greater. From these matrix elements, we can now evaluate the perturbation in the interaction picture by computing the individual frequency components of $e^{iUHU^\dagger t} \sigma_x e^{-iUHU^\dagger t}$ using

$$e^{iUHU^\dagger t} = \sum_{n,s} \overline{|n, s\rangle} \langle n, s| e^{i\omega_{n,s} t}, \quad (4.58)$$

where $\hbar\omega_{n,s}$ is the energy of the state $|n, s\rangle$. Expanding $e^{iUHU^\dagger t} \sigma_x e^{-iUHU^\dagger t}$, we get

$$\begin{aligned} e^{iUHU^\dagger t} \sigma_x e^{-iUHU^\dagger t} &= \\ &\sum_{n,n',s,s'} \overline{|n, s\rangle} \langle n, s| \sigma_x |n', s'\rangle \overline{\langle n', s'|} e^{i\Delta\omega t}, \end{aligned} \quad (4.59)$$

where $\Delta\omega$ is the frequency difference between $|n, s\rangle$ and $|n', s'\rangle$. Here we only consider the leading-order term of each frequency component, so we make the approximation that $\overline{|n, s\rangle} \langle n', s'|} \approx |n\rangle \langle s| \langle n'| \langle s'|$. With this approximation, we can now write

$$e^{iUHU^\dagger t} \sigma_x e^{-iUHU^\dagger t} = W + W^\dagger, \quad (4.60)$$

where

$$\begin{aligned} W &= -\sigma_- e^{-i\omega'_0 t} \\ &+ \sigma_z \sum_n |n-1\rangle \langle n| \frac{J_1(x)\omega_0}{\omega} e^{-i\omega t} \\ &- \sigma_- \sum_n |n-2\rangle \langle n| \frac{J_2(x)\omega_0}{2\omega} \left(e^{-i(2\omega+\omega'_0)t} - e^{-i(2\omega-\omega'_0)t} \right) \\ &+ O(q \geq 3). \end{aligned} \quad (4.61)$$

We can simplify this expression further. In the classical limit, we can make the approximation that

$$a \approx \sqrt{\lambda} \sum_n |n-1\rangle \langle n|, \quad (4.62)$$

provided that λ is large. Thus when applying operators of the form $\sum_n |n - q\rangle \langle n|$ to a coherent state, we get

$$\left(\sum_n |n - q\rangle \langle n| \right) |\alpha\rangle \approx \frac{a^q}{\sqrt{\lambda^q}} |\alpha\rangle, \quad (4.63)$$

$$= e^{iq\theta} |\alpha\rangle. \quad (4.64)$$

This allows us to replace the terms operating on the photon field in W with simple phase factors. Thus,

$$\begin{aligned} W = & -\sigma_- e^{-i\omega'_0 t} \\ & + \sigma_z e^{i\theta} \frac{J_1(x)\omega_0}{\omega} e^{-i\omega t} \\ & - \sigma_- e^{2i\theta} \frac{J_2(x)\omega_0}{2\omega} (e^{-i(2\omega+\omega'_0)t} - e^{-i(2\omega-\omega'_0)t}) \\ & + O(q \geq 3). \end{aligned} \quad (4.65)$$

Each independent frequency term of W corresponds to a Q operator. In particular for a perturbation $\delta H(t) = \sigma_x \delta B(t)$, the corresponding Q operators are

$$Q_{\omega'_0} = -\sigma_-, \quad (4.66)$$

$$Q_\omega = \frac{J_1(x)\omega_0}{\omega} \sigma_z, \quad (4.67)$$

$$Q_{2\omega+\omega'_0} = -\frac{J_2(x)\omega_0}{2\omega} \sigma_-, \quad (4.68)$$

$$Q_{2\omega-\omega'_0} = \frac{J_2(x)\omega_0}{2\omega} \sigma_+. \quad (4.69)$$

In writing these operators we have omitted any phase factors, as in both equations (4.24) and (4.26), Q is always paired with $Q^\dagger$, and so any phase factor will be eliminated.

4.1.2.4 Explicit Calculation for Magnetic Field Fluctuations

Recalling that equation (4.24) allows us to calculate the expectation value of an operator at any time in the interaction picture, and combining equations (4.24) and (4.66-4.69), we can now explicitly compute the effects of fluctuations in the amplitude and phase of the dressing field. As in [Swa+18], we compute the phase variance and frequency shifts by taking the real and imaginary parts of $\langle \sigma_x + i\sigma_y \rangle$ in the interaction picture (rotating frame) for a spin which starts in the $+\hat{x}$ direction at time $t = 0$, i.e., $|\psi_0\rangle = (|+_z\rangle + |-_z\rangle)/\sqrt{2}$. We will denote the rate of accumulation

of phase variance as Γ_ϕ and the frequency shift as $\delta\omega$. Given a perturbation $\delta H(t) = \gamma\delta B(t)\sigma_x/2$, we obtain

$$\begin{aligned}\Gamma_\phi = & -\frac{d\langle\sigma_x\rangle}{dt} = \frac{\gamma^2}{4}S(\omega'_0) \\ & + \left(\frac{\gamma J_1(x)\omega_0}{\omega}\right)^2 S(\omega) \\ & + \frac{1}{4}\left(\frac{\gamma J_2(x)\omega_0}{2\omega}\right)^2 [S(2\omega - \omega'_0) + S(2\omega + \omega'_0)] \\ & + O(q \geq 3),\end{aligned}\tag{4.70}$$

$$\begin{aligned}\delta\omega = & \frac{d\langle\sigma_y\rangle}{dt} = -\frac{\gamma^2}{4\pi} \int_{-\infty}^{\infty} d\omega' \frac{S(\omega')}{\omega' - \omega'_0} \\ & - \frac{1}{4\pi} \left(\frac{\gamma J_2(x)\omega_0}{2\omega}\right)^2 \int_{-\infty}^{\infty} d\omega' S(\omega') \\ & \times \left(\frac{1}{\omega' - 2\omega + \omega'_0} + \frac{1}{\omega' - 2\omega - \omega'_0}\right) \\ & + O(q \geq 3).\end{aligned}\tag{4.71}$$

We note that in the case where $\delta B(t)$ is different for spins in the same measurement cell, as would be the case for magnetic field gradients, Γ_ϕ may be regarded as the transverse relaxation rate, i.e., $\Gamma_\phi = 1/T_2$. Conversely, in the case where all neutrons experience the same time-dependent magnetic field as would be the case for dressing field fluctuations arising due to current fluctuations in the dressing coil, Γ_ϕ represents the run-to-run phase variation. That is, $\Gamma_\phi = \frac{d}{dt} \text{Var}(\phi)$, where the variance is taken over multiple runs of the experiment. While Γ_ϕ and $\delta\omega$ are quantities that are of general importance in situations where precise control of the phase of dressed spins is required, for nEDM@SNS we are particularly concerned with effects which can perturb neutron- ^3He capture rate. This rate depends linearly on $\vec{\sigma}_1 \cdot \vec{\sigma}_2$, so we would like to compute the variance of this operator in order to find the run-to-run variance in the signal rate, this time under the assumption that the neutrons and ^3He atoms are exposed to the same magnetic field. However, the expectation values computed from time-dependent perturbation theory are in the interaction picture, so to convert these values to those actually observed in the lab frame we must find the variance of $U_1 U_2 e^{iHt} \sigma_1 \cdot \sigma_2 e^{-iHt} U_2^\dagger U_1^\dagger$. Details of this calculation are found in section 4.1.8.5. Generically, this is a time-dependent operator having terms which oscillate at frequencies that are sums of integer multiples of the dressing field frequency ω and the dressed Larmor precession frequency ω'_0 . However, because capture signal rates in the nEDM@SNS experiment are not sufficiently high to resolve oscillations

at these frequencies, we consider only the time-independent component of this operator. We can now calculate the variance in the capture rate due to a fluctuating magnetic field. Let $\vec{b}_1 = \langle \vec{\sigma}_1 \rangle (t = 0)$ and $\vec{b}_2 = \langle \vec{\sigma}_2 \rangle (t = 0)$ be the Bloch vectors of the neutron and ^3He atom at time $t = 0$. Then

$$\begin{aligned} \text{Var}_{cl}(\sigma_1 \cdot \sigma_2)(t) = & \frac{1}{2}(\gamma_1 - \gamma_2)^2 \left| \hat{z} \times (\vec{b}_1 \times \vec{b}_2) \right|^2 S(\omega'_0)t \\ & + 2 \left(\frac{\gamma_1 J_1(x_1)\omega_1 - \gamma_2 J_1(x_2)\omega_2}{\omega} \right)^2 \left| \hat{z} \cdot (\vec{b}_1 \times \vec{b}_2) \right|^2 S(\omega)t \\ & + \frac{1}{2} \left(\frac{\gamma_1 J_2(x_1)\omega_1 - \gamma_2 J_2(x_2)\omega_2}{2\omega} \right)^2 \left| \hat{z} \times (\vec{b}_1 \times \vec{b}_2) \right|^2 [S(2\omega - \omega'_0) + S(2\omega + \omega'_0)] t \\ & + O(q \geq 3). \end{aligned} \quad (4.72)$$

The first term in equation (4.72) represents the variance due to noise at the effective Larmor frequency. However, if both spins start in the plane of precession, this term will not contribute. The second term represents variance due to noise at the dressing frequency, which unlike the first term is maximized if the spins start in the plane of precession. The third term, like the first, is minimized when the spins start in the plane of the precession and is generated from noise at twice the dressing frequency shifted by the effective Larmor precession frequency.

Equation (4.72) may appear to suggest that if $\vec{b}_1 = \vec{b}_2$, one could altogether eliminate variance due to current fluctuations. However, under critical spin dressing the quantity proportional to the nEDM is the relative phase between neutron and ^3He spins, i.e., $\phi_2 - \phi_1 \propto d_n$. Meanwhile, the signal rate of nEDM@SNS is linear in $\sigma_1 \cdot \sigma_2$, which is in turn related to the relative phase by

$$\sigma_1 \cdot \sigma_2 = \cos(\phi_2 - \phi_1), \quad (4.73)$$

provided both spins start in the plane of precession. Thus, setting $\phi_2 = \phi_1$ would decrease sensitivity to an nEDM, as $\sigma_1 \cdot \sigma_2$ would no longer be sensitive to the relative phase in this configuration. Furthermore, the variance of $\phi_2 - \phi_1$ is independent of the starting phase of either spin. Again assuming both spins start in the plane of precession,

$$\hat{z} \cdot (\vec{b}_1 \times \vec{b}_2) = \sin(\phi_2 - \phi_1), \quad (4.74)$$

The variance can be propagated using the Taylor expansion. We have,

$$\text{Var}_{cl}(\sigma_1 \cdot \sigma_2) \approx |\sin(\phi_2 - \phi_1)|^2 \text{Var}_{cl}(\phi_2 - \phi_1). \quad (4.75)$$

Therefore, in the plane of precession the phase of the spin of one species with respect to the other spin species does not matter, and the variance in relative phase will accumulate according to

$$\text{Var}_{cl}(\phi_2 - \phi_1) \approx 2 \left(\frac{\gamma_1 J_1(x_1) \omega_1 - \gamma_2 J_1(x_2) \omega_2}{\omega} \right)^2 S(\omega) t. \quad (4.76)$$

In figure 4.1, we show the per-run frequency uncertainty due to noise in the dressing field with a perfect band-pass around the dressing frequency, predicted by equation (4.76). Explicitly, the frequency uncertainty is found by taking the square root of the variance at a given time in the measurement to find the standard deviation of the phase at that time, then dividing this phase uncertainty by the time in the measurement to find the frequency uncertainty. The per-run sensitivity of nEDM@SNS according to reference [Ahm+19] is also shown for comparison. If we define SNR to be

$$\text{SNR} = P_{\text{signal}}/P_{\text{noise}} = (A_{\text{signal}}/A_{\text{noise}})^2, \quad (4.77)$$

where P is the power of the signal or noise and A is the corresponding amplitude, we find that for 80 dB SNR, and a perfect band-pass filter around the dressing frequency, the phase uncertainty is comparable in magnitude to the statistical sensitivity of nEDM@SNS.

Lastly, we find an absence of run-to-run signal variance for neutrons and ^3He , $\text{Var}_{cl}(\sigma_n \cdot \sigma_3)$, from magnetic fluctuations arising from motion in a static magnetic gradient of any field along any direction, assuming that the gradient field does not change from run to run. This is due to two effects: first, for each run $\sigma_n \cdot \sigma_3$ is averaged over many particles as each neutron and ^3He samples a different magnetic field; and second, the neutron and ^3He trajectories are uncorrelated, and so equation (4.72) (which is derived under the assumption that the neutron and ^3He are exposed to the same field) does not apply. Instead, because these trajectories are uncorrelated, the decay rate of the signal during a run can be computed from the relaxation rate of each individual species using the formalism of reference [Swa+18]. The relaxation rate and frequency shifts for each species are given by the expectation value of σ_+ :

$$\text{Re}(\langle \sigma_+ \rangle) = 1 - \frac{t}{T_2}, \quad (4.78)$$

$$\text{Im}(\langle \sigma_+ \rangle) = \delta \omega t. \quad (4.79)$$




Figure 4.1: Run-to-run uncertainty in the measured frequency due to dressing field noise for 57 pT/√Hz and 5.7 pT/√Hz dressing field noise and a perfect band-pass filter around the dressing frequency, calculated from equation (4.76). The predicted per-run sensitivity for the nEDM@SNS, formulated from reference [Ahm+19], is shown for comparison. It is shown that 57 pT/√Hz dressing field noise will significantly reduce the ultimate sensitivity of the measured frequency, while 5.7 pT/√Hz dressing field noise is acceptable.

Therefore, keeping only terms to second order in magnetic gradients after computing the expectation value of σ_+ and σ_z with the perturbed wave function we find

$$\text{Re}(\langle \sigma_{+n} \rangle \langle \sigma_{-3} \rangle) = 1 - \frac{t}{T_{2n}} - \frac{t}{T_{23}}, \quad (4.80)$$

$$\langle \sigma_{zn} \rangle \langle \sigma_{z3} \rangle \approx 0. \quad (4.81)$$

Using the identity $\sigma_n \cdot \sigma_3 = \sigma_{zn} \sigma_{z3} + \text{Re}(\sigma_{+n} \sigma_{-3})$, we derive

$$\frac{d \langle \sigma_n \cdot \sigma_3 \rangle}{dt} \frac{1}{\cos(\phi_n - \phi_3)} = - \left(\frac{1}{T_{2n}} + \frac{1}{T_{23}} \right) \approx - \frac{1}{T_{23}}. \quad (4.82)$$

Transverse relaxation, T_2 , for both species can be calculated from reference [Swa+18]. Typically, because the trajectories of UCN are ballistic, and highly oscillatory in the measurement cell resulting in significant motional narrowing, the contribution

ϕ_{n3}	ϕ_n	θ_n	ϕ_3	θ_3
0	-0.99650	1.81020	-1.00759	1.85511
$\pi/4$	-0.59186	1.84162	-1.40374	1.77373
$\pi/2$	-0.18485	1.83092	-1.78485	1.66284

Table 4.1: Optimized parameters for the starting position on the Bloch sphere for neutrons and ^3He to achieve cosine dressing with a sine dressing field described in section 4.1.2.5. $\phi_{n3} = \phi_n - \phi_3$ is the relative azimuthal angle of the neutron and ^3He that is desired during critical dressing, $\phi_{n,3}$ and $\theta_{n,3}$ are the absolute azimuthal angle and absolute polar angle of the neutron and ^3He , respectively. When the neutron and ^3He start at these positions with sine dressing, the result is cosine dressing dynamics for the respective relative phase ϕ_{n3} . All units are in radians. Values are for the fiducial nEDM@SNS experiment parameters for critical dressing, where the dressing field amplitude is $40.2497 \mu\text{T}$ oscillating at 1 kHz with a static uniform field strength of $5.2 \mu\text{T}$ in the $\hat{z}$ direction. Optimization of the fiducial nEDM@SNS parameters are discussed in reference [Ahm+19].

of the transverse relaxation of the neutron to the signal decay is small. This result shows that there is no run-to-run phase fluctuation if gradients of all magnetic fields remain the same run-to-run. Gradients in magnetic fields diminish the statistical sensitivity by decreasing the coherence time of a single measurement, and the effect is predictable without stochastic variations, contrary to the case with dressing field amplitude fluctuations.

4.1.2.5 The feasibility of dressing with a cosine waveform

In order to minimize the variance in equation (4.72), both spins should start in the plane of precession at $t = 0$. However, in deriving equation 4.72 we assumed a magnetic field with a cosine dependence, i.e., $B_x(t) = B_1 \cos(\omega t)$. This poses a problem, as such a dressing pulse would require an instantaneous change in magnetic field at $t = 0$, which is impossible to achieve due to the inductance of the dressing coil and limited slew rates of power supplies. Instead, the same dynamics can be achieved with sine dressing, i.e., $B_x(t) = B_1 \sin(\omega t)$, by having the neutron and ^3He spins start outside the plane of precession. The polar angle $\theta_{n,3}$ and absolute phase $\phi_{n,3}$ of neutron and ^3He required at $t = 0$ to achieve this are specified in table 4.1.

4.1.3 Simulation

The calculations in the previous section are verified by numerically integrating the Bloch equations with an adaptive Runge-Kutta integrator. Ensembles of spins are

simulated using the fiducial critical spin dressing parameters for the nEDM@SNS experiment ($B_1 \approx 40 \mu\text{T}$, $\omega \approx 2\pi \times 1 \text{ kHz}$, $\omega'_0 \approx 2\pi \times 100 \text{ Hz}$), with the addition of a noisy magnetic field with amplitude spectral density of $57 \text{ pT}/\sqrt{\text{Hz}}$. This assumed spectral density corresponds to a 0.01% rms variation in B_1 over a bandwidth of 5 kHz. While this is a reasonable figure for the power supplies being considered for nEDM@SNS, it should be noted that the noise power spectrum for the dressing field will generally be peaked around the dressing field frequency, and therefore $S(\omega = 1 \text{ kHz})$ may be significantly larger than the assumed value of $57 \text{ pT}/\sqrt{\text{Hz}}$. The noise power spectrum is varied by changing the cutoff frequency of a simulated high-pass filter. In computing the theoretical values we assume that this high-pass filter is ideal, so that

$$S(\omega) = \begin{cases} \left(57 \text{ pT}/\sqrt{\text{Hz}}\right)^2 & |\omega| > \omega_{\text{cutoff}} \\ 0 & \text{otherwise.} \end{cases} \quad (4.83)$$

We substitute this power spectrum into equations (4.70) and (4.71) to obtain the theory line in figures 4.2 and 4.3. Figure 4.2 shows the simulated and theoretical relaxation times for different values of the high-pass filter cutoff. As evidenced by a sharp decrease in relaxation rate as the cutoff frequency is swept past $2\pi\omega'_0$ and $2\pi\omega$, T_2 depends only on the power spectral density at a set of discrete frequencies. Figure 4.3 compares the simulated and theoretical frequency shifts. The frequency shift decreases logarithmically as the high-pass filter cutoff is increased past the dressed Larmor frequency. Therefore, if a physical high-pass filter is employed to reduce the run-to-run phase variation due to current fluctuations, its cutoff frequency should be chosen to be well above the Larmor frequency in order to mitigate the frequency shift. It should be noted however that for the case of nEDM@SNS experiment, a frequency shift caused by current fluctuations can be corrected for by comparing the signal from the two measurement cells for the same run, or by comparing the signal after reversal of the electric field.

4.1.3.1 Discussion of Results

The results derived in section 4.1.2 and verified in section 4.1.3 apply generally to any time-dependent magnetic field perturbation parallel to the dressing field. For nEDM@SNS, these perturbations can arise primarily from two sources: current fluctuations in the dressing coil, or gradients in the magnetic field coupled with the random motion of neutrons and ^3He atoms in the measurement cells. We first consider the case of current fluctuations, in which all spins are exposed to the




Figure 4.2: Predicted and simulated values of Γ_ϕ for various high-pass filter cutoffs. The theory line is calculated from equations (4.70) and (4.83). Each simulated point represents 200 neutrons. Note that cutoff values beyond the dressing frequency (1000 Hz) are only achievable in simulation, and are shown here for comparison to the theory.

same fluctuating magnetic field. As shown in figures 4.2 and 4.3, the effect of these fluctuations can be separated into two components: the run-to-run variation in phase caused by the stochastic nature of the fluctuations, and a systematic frequency shift. The frequency shift is relatively harmless for nEDM@SNS, as this shift is eliminated by reversal of the electric field. However, the variation in phase over the course of multiple runs can pose a problem, as this contributes directly to the statistical uncertainty of the nEDM measurement, as evidenced by figure 4.1. Because the phase accumulated by this effect varies from run to run, electric field reversal will not nullify this effect. One could also attempt to correlate the scintillation signals from each of the two measurement cells under the assumption that the perturbing magnetic field is identical between the two cells. However, this approach also has limitations, as we discuss in the following section. Especially problematic is noise at the dressing field frequency, as the results of equation (4.72) indicate that the primary contribution to phase variance for spins precessing in the plane orthogonal to the static field comes from noise at this frequency. This is compounded with the results from tests with the power supply used in section 4.1.4.2, which indicate




Figure 4.3: Predicted and simulated frequency shifts, with magnetic field noise given by equation (4.83), for various high-pass filter cutoffs. The theory line is calculated from equations (4.71) and (4.83). Each simulated point represents 100,000 neutrons.

that the noise power spectrum is peaked around the dressing field frequency. Thus, to reduce the statistical uncertainty caused by fluctuations in the dressing field, the following section will explore several techniques to mitigate this type of noise.

4.1.4 Mitigation Strategies

The considerations in the previous sections show that fluctuations in the nEDM signal due to current fluctuations in the dressing coil would arise primarily due to the noise power spectrum at $S(\omega)$, which corresponds to amplitude fluctuations in the dressing field. We discuss two strategies to mitigate this effect. First, one can leverage the fact that the two measurement cells of nEDM@SNS are exposed to opposite electric fields in order to eliminate this effect during data postprocessing; however, as we show, this approach has limitations when there are static field gradients between the two cells. An alternative approach is to directly reduce $S(\omega)$ by applying feedback to the dressing coil in order to maintain the dressing field amplitude close to its target value.

4.1.4.1 Cross-Cell Correlation

In order to combat systematic effects, the nEDM@SNS experiment uses two measurement cells which will be exposed to opposite electric fields [Ahm+19]. Under the assumption that the magnetic field noise $\delta B(t)$ is identical in the two cells, it may be possible to correlate the effects of dressing field noise in the two cells and thus eliminate its effect on the measurement of the neutron EDM. One confounding factor for this strategy is the presence of static magnetic field gradients. Due to imperfections in the static and dressing field coils, the dressed Larmor frequencies between the two cells may differ by a small amount. In this case, the same noise field $\delta B(t)$ may affect the two cells differently, hampering our ability to correlate the two cells. In order to determine the correlations between the spins in different cells, we compute the expectation value $\langle \vec{\sigma}_1 \cdot \vec{\sigma}_2 \rangle$, where $\vec{\sigma}_1$ and $\vec{\sigma}_2$ represent two spins of the same species in opposite cells. Let ω_1 and ω_2 be the dressed Larmor frequencies in the two cells. In this case, calculating the expectation value is complicated by the fact that ω_1 and ω_2 may be very close in frequency, and thus we can no longer use the approximation that cross-frequency terms may be neglected. These terms are calculated in the appendix (section 4.1.8.1). The expectation value $\langle \vec{\sigma}_1 \cdot \vec{\sigma}_2 \rangle$ is then given by

$$\langle \vec{\sigma}_1 \cdot \vec{\sigma}_2 \rangle = 1 - \frac{\gamma^2}{2} S(\bar{\omega}) \left(t - \frac{\sin(2\Delta\omega t)}{2\Delta\omega} \right). \quad (4.84)$$

We see that decorrelation between the two cells may only be neglected if either $S(\bar{\omega})$ is small, or $\Delta\omega t$ remains small over the course of the experiment.

4.1.4.2 Feedback Control

In situations where decorrelation between the two cells is too large to apply the strategy of the previous section, we can instead employ feedback control on the dressing field. As we are primarily concerned with amplitude fluctuations in the dressing field, we regard the instantaneous dressing amplitude as a constant ideal value $B_{1,\text{ideal}}$, plus a small fluctuating $\delta B_1(t)$, so that the dressing field is given by $B(t) = (B_{1,\text{ideal}} + \delta B_1(t)) \cos(\omega t)$. Under critical dressing, the phase accumulated between the neutron and ^3He is, to first order in $\delta B_1(t)$, proportional to the integral of $\delta B_1(t)$ over time. Thus, one way to limit the impact of current fluctuations would be to employ feedback control to minimize the deviation of this integral from a target value. Using a Kepco four-quadrant 400 36-12 power supply and a DT9837A digital signal analyzer, we demonstrate a proof-of-concept for using feedback control in the nEDM@SNS experiment. The signal analyzer outputs a 1000 Hz sine wave,

which we feed into the voltage control input of the power supply. The power supply drives a cosine coil, at the center of which is a pickup coil. The voltage across this pickup coil is measured by the signal analyzer, and the output of the signal analyzer is adjusted based on this measurement using a digital PI loop to maintain the integral of the signal amplitude near its target value. We then use the measured fluctuating magnetic field to simulate the motion of neutron and ^3He spins, using the CODATA values for their gyromagnetic ratios [NIS]. The measured magnetic field is postprocessed with a high-pass filter to remove lower-frequency noise, and a single global scaling factor is applied to the field to replicate the CSD parameters of the nEDM@SNS experiment. Figures 4.4 and 4.5 show the resulting rms deviation of the relative phase shift between neutron and ^3He spins. In figure 4.5, we also plot the phase accumulation estimated by applying equation (4.76) to the measured noise power spectrum of the Kepco power supply. Assuming a $400 \text{ pT}/\sqrt{\text{Hz}}$ white noise spectrum which is then high-pass filtered at 500 Hz, we can see that this approach underestimates the actual phase accumulation. This indicates that, as expected, the noise power spectrum is peaked around the dressing field frequency. Nonetheless, these results show that introducing the feedback loop substantially reduces the run-to-run phase variation caused by amplitude fluctuations in the dressing field. There are a number of ways to implement this type of feedback loop in the nEDM@SNS experiment. A simple scheme would be to place a shunt resistor in series with the dressing coil to measure the current. Alternatively, pickup coils may be used, as was done in this section. The advantage of this approach for nEDM@SNS is that multiple pickup coils can be used to detect spatial inhomogeneities in the dressing field. This latter approach also enables us to detect changes in the dressing field that are not due to the current; for example, temperature fluctuations may affect the conductivity of the copper flux return, which would in turn alter the dressing field amplitude even if the current amplitude remains constant.

4.1.5 Robust Dressing

While feedback and cross-cell correlation can be employed to reduce the impact of magnetic field fluctuations which are homogeneous across the measurement cells, these strategies do not mitigate decoherence caused by magnetic field gradients in either the static or dressing field. Thus, we now turn our attention to strategies to amplify coherence times of systems undergoing fluctuating fields, particularly in the presence of magnetic field gradients. The dressing scheme presented in this section is a novel technique which, among other applications, can be used




Figure 4.4: The root mean squared deviation of the angle between the neutron and ^3He spins (ϕ_{n3}) as a function of time, with dressing field powered by a Kepco four-quadrant 400 36-12 power supply, for magnetic fields generated with (solid line) and without (dashed line) feedback control. The dataset without feedback consists of 28 runs, while the dataset with feedback consists of 50. Data are smoothed for clarity.

to mitigate gradient decoherence of the dressed spins. In this scheme, we apply modulation with angular frequency ω_m to the dressing parameter $x = \gamma B_1/\omega$, where ω_m is faster than the dressed Larmor precession ω'_0 . This differs from the modulated CSD presented in references [GL94; Ahm+19; Swa+18] in that our modulation frequency is much higher. We find that by tuning the rate and amplitude of modulation in conjunction with the dressing parameters B_1 and ω so that there is no phase accumulation in the lab frame, we can significantly decrease decoherence, similar to what is found in references [KS15; BIS18; Bev+22]. In a system with two spin species, these parameters can be selected such that both states on a short timescale are rapidly oscillating but on average are effectively frozen in time with respect to the lab frame. With this combination of fast oscillation and no phase accumulation in the lab frame, quantum states become especially robust to field fluctuations. On average, both spin species remain fixed on the Bloch sphere in the lab frame for arbitrary long timescales limited by the coherence. This technique




Figure 4.5: Phase standard deviation with feedback control applied from the field powered by a Kepco four-quadrant 400 36-12 power supply (solid line). The rms phase deviation rises for a period of less than five seconds, then appears to settle near a fixed value. For comparison, the theoretical estimate (dotted line) is computed assuming $400 \text{ pT}/\sqrt{\text{Hz}}$ noise high-pass-filtered at 500 Hz. Data are smoothed for clarity.

may prove useful in storing quantum information by keeping the state static and robust to electromagnetic fluctuations.

For the nEDM@SNS experiment, the extended coherence time afforded by robust dressing is a valuable tool to rapidly identify and compensate for large magnetic field gradients. The nEDM@SNS experiment is performed in a large cryogenic apparatus; as such, modifications to the apparatus require a long dead time to allow the apparatus to warm, and then cool to its operating temperature ($T < 500 \text{ mK}$). Thus, it is crucial to be able to mitigate gradients in the measurement cell without the need to make modifications to the apparatus. Aside from robust dressing, nEDM@SNS has two techniques available to measure static magnetic field gradients. First, for stray static field gradients in the precession volume, gradients can be determined by the cryogenic probe array described in reference [NP14a]. This array consists of 39 fluxgate magnetometers which will reconstruct the magnetic field by fitting the coefficients of a harmonic scalar potential expansion to the magnetic field values measured by the probes. However, if the ideal static uniform field is distorted by

magnetic materials within the volume of interest, the probe array described will not be sufficient to reconstruct the field due to the reconstruction condition that there be no magnetic sources within the volume of interest. Second, gradients can be determined by measuring transverse relaxation rates versus the applied field. However, during the initial phases of the experiment, there may be static field gradients large enough that the scintillation signal decays too quickly to be measured effectively. Such gradients may be caused, for example, by an unknown nickel flash (nickel coating) on an electric connection close to the measurement cell. In the following sections we first outline a form of precession-free dressing which we call robust dressing because it is particularly robust against relaxation and dressing field noise. This is due to the fact that a spin undergoing robust dressing does not accumulate phase in the lab frame, and remains fixed on the Bloch sphere. The longer relaxation time of the robustly dressed system allows us to measure larger gradients than would normally be possible. These results are similar to those obtained for multiple field dressing in references [Bev+20; Bev+22]. While these references consider dressing fields acting along multiple axes, here we limit ourselves to a single dressing field in the $\hat{x}$ direction, as is the case for nEDM@SNS experiment. We apply this technique to find precession-free dressing for both ^3He and neutron simultaneously, which we call robust critical dressing (RCD).

4.1.5.1 Description within the Bessel function approximation

The dressed system of a single species can be held at a value where no phase accumulates in the lab frame, around the zero crossing of the Bessel function, $J_0(\gamma B_1/\omega) \approx 0$. This is described in detail in reference [ZCA12] for a number of multiple spin systems to achieve a Zeeman-insensitive optical clock atomic transition. Alternatively, the dressing parameter can be modulated between values that achieve a system where no phase accumulation occurs in the lab frame. For example, the effective absolute phase of the neutron is approximately

$$\phi_n(t) = \int_0^t \gamma_n J_0 \left(\frac{\gamma_n B_1(t)}{\omega} \right) B_0 dt. \quad (4.85)$$

There are many possible choices for $B_1(t)$. Here we choose $B_1(t)$ to be a constant plus a cosine of frequency ω_m , so that

$$\phi_n(t) = \int_0^t \gamma_n J_0 \left\{ \frac{\gamma_n B_1}{\omega} \left[x_0 + \frac{x_1}{2} (1 + \cos \omega_m t) \right] \right\} B_0 dt, \quad (4.86)$$

for dimensionless parameters x_0 and x_1 . Robust dressing is achieved when there is no phase accumulation in the lab frame,

$$\langle \phi_n(t) \rangle = 0, \quad (4.87)$$

parameters for which can be found numerically. The parameters required to achieve robust dressing in simulations with cosine modulation can be estimated from equation (4.86) and (4.87) by integration of the modulation function over a modulation period; moderate deviations from the values predicted analytically are expected due to violation of the Landau-Zener approximation. This is described in the context of spin dressing in reference [Eck+12]. It is found that the shorter the modulation period, the further the RCD parameters deviate from the analytic prediction. In the simple analytic formulation it is found that if $x_0 = 1.2$ and $x_1 \approx 3.17$ then $\langle \phi_n(t) \rangle \approx 0$ (for both the neutron and ^3He .) Despite violation of the Landau-Zener approximation, detailed simulations of the Bloch equation find a continuum of values for x_0 and x_1 that satisfy the robust dressing condition, $\langle \phi_n(t) \rangle = 0$. Deviations away from the average phase can be made small when $\omega'_0 < \omega_m$. The average projection of the spin on the Bloch sphere remains fixed, with fluctuations on the order of the modulation period and the instantaneous precession rate.

4.1.5.2 Robustness

Robust dressing suppresses transverse relaxation due to field fluctuations in B_z and B_y . For B_z , this can be seen from the dressing approximation — the effective gyromagnetic ratio under precession-free dressing is zero, and therefore the local static field strength is irrelevant. Likewise, a variation in B_y can be viewed as a rotation of the static field, which in the case where the effective gyromagnetic ratio is zero has no effect on the spins' overall behavior. Thus, the only static gradients and field offsets which contribute to the transverse relaxation rate are those in B_x along any direction. To analyze these, it is useful to consider a simplified model of robust dressing.

4.1.5.3 An Intuitive Model of Robust Dressing

While we are not aware of any purely analytical solutions for the robust dressing field, robust dressing has several properties that allow for a convenient approximation. We have chosen the modulation frequency ω_m such that it evenly divides the dressing field frequency ω . Thus, the pulse is periodic with frequency ω_m , and its behavior is

entirely defined by its action on a spin during the interval $t = 0$ to $t = 2\pi/\omega_m$. This property allows us to analyze this pulse using Floquet theory, described in reference [BNN77]. In short, Floquet's theorem allows us to write the time evolution operator of robust dressing as

$$T(t) = M(t)e^{-i\Lambda t}, \quad (4.88)$$

where $M(t)$ is a unitary operator with period $2\pi/\omega_m$ and Λ is a time-independent Hermitian operator. The operator, Λ , can then be written in the form

$$\Lambda = -\frac{1}{2}\gamma\vec{B}_{\text{eff}} \cdot \vec{\sigma}. \quad (4.89)$$

The overall rotation under robust dressing can thus be treated as arising due to an effective magnetic field $\vec{B}_{\text{eff}}$. The effective magnetic field, $\vec{B}_{\text{eff}}$, can be computed by numerical integration of the Bloch equations over a single period, or approximated perturbatively as in references [Bev+20] and [Bev+22].

Numerically, we find that robust dressing corresponds to a weak ($|\gamma B_{\text{eff}}| \approx 1$ Hz) magnetic field whose primary component is in the $\hat{x}$ direction. The precise orientation of $\vec{B}_{\text{eff}}$ may be adjusted by careful selection of the robust dressing parameters.

In light of this, we analyze relaxation in the robust dressing scenario by treating $\vec{B}_{\text{eff}}$ as an effective static uniform field which lies along $\hat{x}$ and redefining T_1 and T_2 accordingly. We also define ϕ to be the angle between the spin and $\hat{x}$. A geometrical representation of the robust dressed system for a neutron is shown in figure 4.6.

4.1.5.4 Robust dressing applications

Before presenting an analysis of relaxation under robust dressing, we describe several applications of robust dressing for an nEDM search such as the nEDM@SNS experiment.

4.1.5.5 Gradient metrology through spin relaxation

Because the effects of gradients in B_y and B_z along any direction are suppressed by robust dressing, T_2 is always much shorter than T_1 for the nEDM@SNS experiment's expected operating and initial commissioning conditions. This will lead to a fast initial decay dominated by T_2 , followed by a slower decay whose rate is given by T_1 . In general, the polarization of spins decays as

$$P(t) = \sqrt{e^{-\frac{2t}{T_1}} \cos^2 \phi_0 + P_{T_2}^2(t) \sin^2 \phi_0}, \quad (4.90)$$




Figure 4.6: A classical visualization of the neutron spin trajectory undergoing robust dressing on the Bloch sphere in the laboratory frame, with arrows representing the static uniform field, B_0 , the dressing field, B_1 , and the robust dressing effective field B_{eff} .

where $P_{T_2}(t)$ is the component of the polarization transverse to $\hat{x}$ and ϕ_0 is the value of ϕ at the onset of robust dressing. The transition between fast and slow decay then occurs at polarization $P_0 \approx \cos(\phi_0)$. This behavior can be applied for gradient metrology, as a gradient in B_x will cause the spins to relax quickly to a known value, while robust dressing extends the coherence of the system from gradients in B_y or B_z . For the nEDM@SNS experiment, a static magnetic gradient of any field transverse to the static uniform field direction along the dressing axis represents the largest potential systematic effect by generating a frequency shift linear in the applied electric field, as evidenced by references [Com91; Pen+04a; LG05; AL06; Cla11; CM12; Pig+15; SPG16; Swa+18]. If this precise gradient measurement technique is used to feedback on magnetic gradient shimming, a powerful mitigation strategy for this linear-in-E frequency shift is achieved.

4.1.6 Relaxation under Robust Dressing

The longitudinal relaxation of the system can be decomposed into two contributions. One contribution arises from the Redfield-like longitudinal decay, which is generated from field imperfections in B_y and B_z , and which we compute by evaluating the system in the interaction picture, similar to the dual harmonic dressing treatment provided by reference [Bev+22], but with the oscillating field in only the $\hat{x}$ direction. A derivation is found in appendix 4.1.8.7. The model predicts relaxation to be extended beyond the T_1 time found in reference [Red57; McG90] for the corresponding static uniform field of the nEDM@SNS experiment's operating parameters, discussed in reference [Ahm+19]. This is because the integral over the field correlation function is shifted in frequency according to the harmonics of the dressing field and not the static uniform field as is the case in reference [Red57]. An expansion of the spin's reaction to the applied field into a Fourier series is used to formulate the longitudinal relaxation rate. In the diffusion limit we find

$$\frac{1}{T_1} \approx \gamma^2 (G_y^2 + G_z^2) D \sum_n \frac{|a_n|^2}{n^2 \omega_m^2}, \quad (4.91)$$

where a_n are the Fourier coefficients of the series and are found numerically. This result can be extended beyond the diffusion limit by using the spectrum of the trajectory autocorrelation function; this is described in appendix 4.1.8.7.

An additional source of longitudinal relaxation arises from the field imperfections in B_x , which generate longitudinal relaxation by rotating the local effective field off of the x direction, resulting in a contribution of the effective field along y and/or z . This in turn will cause relaxation according to

$$\frac{1}{T_1} \approx (G_{\perp,\text{eff}})^2 \frac{D}{|B_{0,\text{eff}}|^2}, \quad (4.92)$$

$$\approx \gamma^2 \left(\frac{\partial B_{\perp,\text{eff}}}{\partial B_x} \right)^2 G_x^2 \frac{D}{|B_{0,\text{eff}}|^2}, \quad (4.93)$$

where $G_{\perp,\text{eff}}$ and $B_{\perp,\text{eff}}$ refer to components of the effective gradient and magnetic field orthogonal to the effective static uniform field $B_{0,\text{eff}}$. Validity of this estimate depends on the magnitude of gradients in B_x in the cell. In particular, if L_x is the length of the cell along the x direction, then this approximation breaks down when $|G_x L_x / 2| \gtrsim |B_{\text{eff}}|$ because the scaling of $B_{\perp,\text{eff}}$ with B_x becomes non-linear. The longitudinal relaxation will not be proportional to G_x^2 , but will scale with a higher power of G_x due to the non-linearity in the effective field scaling. To find the exact

scaling in closed form is difficult. However, it is found from simulations that the robust dressing parameters can be tuned to mitigate this source of relaxation, even for large gradients in B_x . By tuning the robust critical dressing parameters to minimize the relaxation rate, we can mitigate the contribution from the B_x field imperfections to the level of the contribution from the other field directions. With this optimized tuning, the relaxation rate from this gradient returns to G_x^2 scaling, implying a return to the linear scaling of the effective field under optimized robust dressing. Due to the ability to mitigate or increase relaxation by tuning the robust dressing parameters, we find that this modulation technique is a useful form of quantum control.

In practice, because Maxwell's equations require a gradient contribution in at least two directions, this optimization procedure to minimize relaxation obtains diminishing returns when the relaxation from the gradient in B_x becomes comparable in magnitude to the relaxation contributed from gradients in B_y and/or B_z . Due to the nonlinear scaling, and the precision of the robust dressing parameters required, we expect that in the large gradient regime the relaxation from gradients in B_x can be mitigated to parity with the contribution from B_y and B_z . However, in the small gradient regime, where tuning need not be as precise, we may more easily find optimized robust dressing parameters where the contribution from the gradient in B_x can be ignored. The prediction is formulated as

$$\frac{1}{T_1} = C_{dB_x} G_x^2 + \gamma^2 (G_y^2 + G_z^2) D \sum_n \frac{|a_n|^2}{n^2 \omega_m^2}, \quad (4.94)$$

where C_{dB_x} is determined from specifics of the parameters, but in general, it can be estimated to be of magnitude

$$C_{dB_x} \leq \gamma^2 D \sum_n \frac{|a_n|^2}{n^2 \omega_m^2}. \quad (4.95)$$

The relaxation from simulations of a linear gradient, where $G_x = dB_x/dx = -dB_z/dz$, is shown in figure 4.7 for all regimes of robust dressing, as well as the theoretical prediction after optimization for both limits of C_{dB_x} . The smallness of B_{eff} compared to B_0 and the nonlinear dependence on G_x lead naturally to three different regimes depending on the magnitude of the G_x gradient relative to B_{eff} .

4.1.6.1 Small gradient regime

If $|G_x L_x/2| \ll |B_{\text{eff}}|$, the magnetic field gradient across the cell is smaller in magnitude than the effective static uniform field. The tuning of the parameters need

not be terribly precise ($\delta x_{0,1} \sim x_{0,1} \times 10^{-3}$) to achieve the T_1 relaxation time shown in equation (4.91). Additionally, the G_x contribution to T_1 may be found using the effective magnetic field formalism using equation (4.92).

Similar to traditional NMR transverse relaxation, the RCD analog can be predicted from references [Red57; McG90]; however, only gradients in B_x contribute significantly. Thus, for the transverse polarization we find

$$P_{T_2}(t) = e^{-\frac{t}{T_2}}, \quad (4.96)$$

where in the diffusion limit,

$$\frac{1}{T_2} = \gamma^2 G_x^2 \frac{L_x^4}{120D}. \quad (4.97)$$

Outside the diffusion limit, the transverse relaxation can be found from the spectrum of the trajectory autocorrelation function at zero frequency.

4.1.6.2 Intermediate gradient regime

If $|G_x L_x/2| \sim |B_{\text{eff}}|$, then the gradient magnetic field nearly cancels B_{eff} in a significant portion of the cell. For untuned robust dressing parameters, the effective magnetic field may have a nonzero $B_{z,\text{eff}}$ component in this region, leading the local magnetic field here to be misaligned with the x -axis. This drastically increases the longitudinal relaxation rate for this gradient regime. While no theoretical prediction exists for the dynamics of the spins in this regime, a thorough investigation is presented in reference [SHD91]. The strong dependence of T_1 on G_x in the vicinity of the intermediate gradient regime can be employed for gradient metrology. Specifically, a gradient large enough to induce this regime could be identified by measuring the T_1 relaxation of the spin species under study. Although the gradients in B_x required to achieve this regime are a couple orders of magnitude larger than the initial gradients expected at the onset of commissioning of the nEDM@SNS experiment, they are in the realm of possibility if there is an unaccounted for magnetized material — for example, a nickel flash (nickel coating) on an electrical connection, or other unknown magnetic material close to the measurement cell. Note that while precise ($\delta x_{0,1} \sim x_{0,1} \times 10^{-5}$) tuning of the robust dressing parameters can eliminate the $B_{z,\text{eff}}$ component in the misaligned region, thus recovering the T_1 of equation (4.91), this is not required nor useful for gradient metrology.

4.1.6.3 Large gradient regime

As $|G_x L_x/2|$ becomes large compared to $|B_{\text{eff}}|$, the region of low magnetic field shrinks in volume, and the effect of misalignment becomes less significant. The treatment of T_1 and T_2 is straightforward. Longitudinal relaxation can be computed from equation (4.91). Transverse relaxation under gradients of this magnitude behave according to the adiabatic regime described in reference [Pig+15], as the spins dephase on a shorter timescale than the diffusion time. This results in a relatively fast relaxation, given by

$$P_{T_2}(t) = \frac{2}{\gamma G_x L_x t} \sin(\gamma G_x L_x t/2), \quad (4.98)$$

$$= \text{sinc}\left(\frac{\gamma G_x L_x t}{2}\right), \quad (4.99)$$

where $\text{sinc}(x) = \frac{\sin(x)}{x}$.

4.1.6.4 Robust dressing and dressing field noise

From the simple model that describes robust dressing, we expect that robust dressing is also robust to fluctuations of the dressing field, because the integral of the total phase accumulated on average is zero. In simulations of noise we find that after an initial relaxation period on the timescale of the modulation frequency, the random phase accumulation nearly ceases, and dressing field fluctuations arising from power supply noise can typically be ignored with commercial linear amplifiers. The polarization of a robust dressing waveform generated according to equation 4.86, where $B_1 = 40.2497 \mu\text{T}$ and $\omega_m = 2\pi \times 200 \text{ rad/s}$ with amplitude noise ranging from 20 to 80 dB SNR is shown in figure 4.8. In that figure it is shown that at 20 dB SNR, where the amplitude of the field noise is 10% of the amplitude of the pulse, polarization can be observed for a long time, and the rate in equation (4.91) is still achievable after an initial loss.

4.1.6.5 Robust critical dressing

There exists a non-trivial solution that satisfies the robust dressing conditions for both the neutron and ^3He at the same time. Robust critical dressing is achieved when

$$\langle \phi_n \rangle \approx \langle \phi_3 \rangle \approx 0. \quad (4.100)$$




Figure 4.7: Relaxation after optimization of robust dressing parameters for 200 Hz robust dressing and a wide range of gradient magnitudes. The gradients in this simulation arise from the divergence theorem, $\frac{dB_z}{dz} = -\frac{dB_x}{dx}$.

Table 4.2 shows optimized robust dressing parameters for the CODATA values of the magnetic moments of the neutron and ^3He . We find that robust critical dressing significantly reduces relaxation from field fluctuations in the static uniform field and dressing fields.

Figure 4.9 shows a simulation of the robust critical dressing signal for $\omega_m = 2\pi \times 200$ Hz in a large gradient, along with the theoretical estimation for robust critical dressing based on

$$\langle \sigma_n \cdot \sigma_3 \rangle \approx P_0^2 \exp \left(\gamma_3^2 G_z^2 D t \sum_n \frac{|a_n|^2}{n \omega_m} \right). \quad (4.101)$$

We have ignored neutron relaxation in this prediction, as it is expected to be small because neutrons exhibit strong motional narrowing. The applied gradient satisfies Maxwell's equations with $dB_z/dz = -dB_x/dx = 2 \times 10^{-2} \text{B}_0/\text{cm}$. For comparison,




Figure 4.8: 1- Polarization for robust dressing modulated according to equation 4.86, where $B_1 = 40.2497 \mu\text{T}$ and $\omega_m = 200 \text{ Hz}$, with simulated noise power varied from 20 dB SNR to 80 dB SNR

unmodulated critical dressing is also simulated in the same gradient, and the signal is found to decay rapidly compared to the robust case.

4.1.6.6 Mitigation of linear-in-E effects

Perhaps the most troublesome systematic effects of EDM experiments is the linear-in-E frequency shift. The linear-in-E frequency shift arises from correlated fluctuations due to the motion of the spin in magnetic field gradients perpendicular to the applied static field, coupled with the $\mathbf{v} \times \mathbf{E}$ motional magnetic field arising from the motion of the spin within the measurement cell. The effect can be described from second order perturbation theory in terms of the spectrum of the spin's position-velocity correlation function. An abundance of detail concerning this shift is presented by references [Com91; Pen+04a; LG05; AL06; Cla11; CM12; Pig+15; SPG16; Swa+18].

$\omega_m/2\pi$ (Hz)	x_0	x_1	$\sqrt{\langle \Delta\theta_{n3}^2 \rangle}$ (rad)
200	1.0470	3.1224	0.68
300	1.1005	3.6645	0.41
500	1.3290	3.4435	0.25

Table 4.2: Optimized parameters for robust critical dressing. All values are calculated for a static uniform field of 5.2 μT and amplitude modulation of a nominal critical dressing field of 40.2497 μT oscillating at 1000 Hz.

Nominally for nEDM@SNS, the electric field is applied parallel to the static uniform field. In this case, robust dressing cannot be used to detect an nEDM. However, if the electric field is instead applied along $\hat{x}$, robust critical dressing can be employed to measure an nEDM and simultaneously mitigate linear-in-E frequency shifts arising from the $\mathbf{v} \times \mathbf{E}$ motional magnetic field via the same mechanism that mitigates the longitudinal relaxation; specifically, the response of the spin precession system to magnetic field fluctuations is determined by the power spectrum of the noise evaluated at integer multiples of the modulation frequency rather than at the Larmor frequency. Due to the $1/\omega^2$ dependence of the trajectory correlation functions in this regime, described in reference [SPG16], the contribution to the linear-in-E frequency shift is significantly mitigated.

4.1.7 Conclusion

In this work, we developed a model to evaluate the effect of magnetic field noise on a critically dressed system of neutrons and ^3He atoms in the case where the Larmor frequency is a non-negligible fraction of the dressing field frequency. Applying this model to the case of magnetic field fluctuations parallel to the dressing field, we find that the relaxation time is given by an infinite sum over the noise power spectrum evaluated at discrete frequencies, with higher frequency terms suppressed by higher-order Bessel functions. For the purpose of calculating the impact of current fluctuations in the dressing field coil, these results demonstrate that the primary concern is amplitude fluctuations in the dressing field. These results were verified by numerical integration of the Bloch equations.

With these results as motivation, we propose several strategies to increase coherence and sensitivity for systems of dressed spins. Feedback control of the dressing amplitude, for example, can substantially reduce the phase uncertainty in the presence of current fluctuations in the dressing field coil, while correlating the signal rate from opposing cells allows for rejection of noise sources which are common




Figure 4.9: Simulated signal for highly tuned RCD modulated at 200 Hz (red) and CSD (green), with theoretical estimation for RCD (blue). The starting phase for ^3He is $\phi_3 = \pi/8$, and it is the same for neutrons, $\phi_n = \pi/8$. Signal is observable for a long time in the RCD scenario, while the signal is indistinguishable from noise in a traditional critical dressing scenario.

to both cells. We emphasize that while these strategies are proposed in the context of the nEDM@SNS experiment, these techniques are generally applicable to systems where precise control of dressed spins is required. Of particular interest is robust dressing, which sets the effective Larmor frequency of each species to zero by modulating the dressing field. This substantially reduces relaxation due to magnetic field fluctuations, and effectively freezes a state in time which is useful for extending the coherence of quantum information. Measuring the relaxation of a robustly dressed system allows for precise determination and feedback on gradients parallel to the dressing field, and thus can be applied as gradient metrology for fundamental physics measurements which require a uniform magnetic field. For the nEDM@SNS experiment, this effect can be leveraged to measure or shim relatively large spatial gradients in a critically dressed system of neutrons and ^3He . Furthermore, robust dressing allows for mitigation of linear-in-E frequency shifts in EDM

experiments where the electric field is chosen to be parallel to the dressing field.

4.1.8 Appendix to [TS22]

The calculations presented in the following sections are the appendix to the paper [TS22], in which we work through the quantum mechanics formalism necessary to derive the results in the article.

4.1.8.1 Near-Frequency Terms

Suppose $V_I(t)$ contains a pair of terms with frequencies ω_j and ω_k which are very close in frequency, such that $|\omega_j - \omega_k| \sim 1/t$. Then the cross-frequency term can no longer be neglected. Define the integral

$$v_{jk} \equiv \int_0^t dt' \int_0^{t'} dt'' e^{-i\omega_j t'} e^{i\omega_k t''} \langle \delta B(t') \delta B(t'') \rangle. \quad (4.102)$$

Using equation (4.9), v_{jk} can be written in terms of the noise power spectrum as

$$v_{jk} = \frac{e^{-2i\Delta\omega t} - 1}{-2i\Delta\omega} \left(\frac{1}{2} S(\bar{\omega}) - \frac{i}{2\pi} \int_{-\infty}^{\infty} d\omega' \frac{S(\omega')}{\omega' - \bar{\omega}} \right).$$

Then the cross terms of $T^\dagger AT$ are

$$\begin{aligned} & \int_0^t \int_0^{t'} dt' dt'' (Q_j e^{-i\omega_j t''} + Q_j^\dagger e^{i\omega_j t''}) A (Q_k e^{-i\omega_k t'} + Q_k^\dagger e^{i\omega_k t'}) \\ & - \int_0^t dt' \int_0^{t'} dt'' (Q_j e^{-i\omega_j t''} + Q_j^\dagger e^{i\omega_j t''}) (Q_k e^{-i\omega_k t'} + Q_k^\dagger e^{i\omega_k t'}) A \\ & - \int_0^t dt' \int_0^{t'} dt'' A (Q_k e^{-i\omega_k t'} + Q_k^\dagger e^{i\omega_k t'}) (Q_j e^{-i\omega_j t''} + Q_j^\dagger e^{i\omega_j t''}) \\ & + (j \leftrightarrow k) \\ & = Q_j A Q_k^\dagger (v_{jk} + v_{kj}^*) + Q_j^\dagger A Q_k (v_{kj} + v_{jk}^*) \\ & - Q_j Q_k^\dagger A v_{kj}^* - Q_j^\dagger Q_k A v_{kj} \\ & - A Q_k Q_j^\dagger v_{kj} - A Q_k^\dagger Q_j v_{kj}^* \\ & + (j \leftrightarrow k), \\ & = v_{jk} (Q_j A Q_k^\dagger + Q_k^\dagger A Q_j - Q_k^\dagger Q_j A - A Q_j Q_k^\dagger) \\ & + v_{jk}^* (Q_k A Q_j^\dagger + Q_j^\dagger A Q_k - Q_k Q_j^\dagger A - A Q_j^\dagger Q_k) \\ & + (j \leftrightarrow k), \end{aligned}$$

$$\begin{aligned}
&= v_{jk} [Q_k^\dagger, [A, Q_j]] \\
&+ v_{jk}^* [Q_k, [A, Q_j^\dagger]] \\
&+ (j \leftrightarrow k).
\end{aligned}$$

4.1.8.2 Explicit Calculation for Case of Static Gradients

We first consider noise near the dressed Larmor frequency ω'_0 . Using $Q_j = \frac{\gamma}{2}\sigma_{+,1}$ and $Q_k = \frac{\gamma}{2}\sigma_{+,2}$, we get

$$\begin{aligned}
\langle \vec{\sigma}_1 \cdot \vec{\sigma}_2 \rangle &= 1 - \frac{\gamma^2}{4}(v_1 + v_1^* + v_2 + v_2^*) - \frac{\gamma^2}{4}(v_{kj}^* + v_{kj} + v_{jk} + v_{jk}^*), \\
&= 1 - \frac{\gamma^2}{2}tS(\bar{\omega}) - \frac{\gamma^2}{4} \frac{\sin(2\Delta\omega t)}{\Delta\omega} S(\bar{\omega}), \\
&= 1 - \frac{\gamma^2}{2}S(\bar{\omega}) \left(t - \frac{\sin(2\Delta\omega t)}{2\Delta\omega} \right).
\end{aligned}$$

Next, we consider the contribution of noise near ω , the dressing field frequency. In this case, static gradients do not affect the time-dependence of the noise Hamiltonian — that is, the frequency of the complex exponential $e^{i\omega t}$ is unaffected by a change in ω_0 or B_1 (see equation (4.1.2.3)). Therefore, there are no cross-frequency terms to consider, and so noise near ω which applies equally to both cells cannot cause the spins between the two cells to decorrelate.

4.1.8.3 Transformation of the Signal Rate into the Interaction Picture

Computing the variance of $\sigma_1 \cdot \sigma_2$ in the lab frame corresponds to computing the variance of $U_1 U_2 e^{iHt} \sigma_1 \cdot \sigma_2 e^{-iHt} U_2^\dagger U_1^\dagger$ in the interaction picture. We begin by calculating $U_1 U_2 \sigma_1 \cdot \sigma_2 U_2^\dagger U_1^\dagger$.

From the definition of U_1 ,

$$U_1 = D(\eta_1/2) |+_x\rangle \langle+_x| + D^\dagger(\eta_1/2) |-_x\rangle \langle-_x|, \quad (4.103)$$

$$= \frac{1}{2}D(\eta_1/2)(1 + \sigma_x) + \frac{1}{2}D^\dagger(\eta_1/2)(1 - \sigma_x). \quad (4.104)$$

The Pauli spin operators then transform as

$$U_1 \sigma_x U_1^\dagger = \sigma_x, \quad (4.105)$$

$$U_1 \sigma_y U_1^\dagger = \frac{1}{2}(D(\eta_1) + D^\dagger(\eta_1))\sigma_y + \frac{1}{2}i(D(\eta_1) - D^\dagger(\eta_1))\sigma_z, \quad (4.106)$$

$$U_1 \sigma_z U_1^\dagger = \frac{1}{2}(D(\eta_1) + D^\dagger(\eta_1))\sigma_z - \frac{1}{2}i(D(\eta_1) - D^\dagger(\eta_1))\sigma_y. \quad (4.107)$$

So the observable $\sigma_1 \cdot \sigma_2$ transforms as

$$U_1 U_2 (\sigma_1 \cdot \sigma_2) U_2^\dagger U_1^\dagger = (U_1 \sigma_{x1} U_1^\dagger) (U_2 \sigma_{x2} U_2^\dagger) + (U_1 \sigma_{y1} U_1^\dagger) (U_2 \sigma_{y2} U_2^\dagger) + (U_1 \sigma_{z1} U_1^\dagger) (U_2 \sigma_{z2} U_2^\dagger) \quad (4.108)$$

$$= \sigma_{x1} \sigma_{x2} \quad (4.109)$$

$$+ \frac{1}{2} (D(\Delta\eta) + D^\dagger(\Delta\eta)) (\sigma_{y1} \sigma_{y2} + \sigma_{z1} \sigma_{z2}) \quad (4.110)$$

$$- \frac{i}{2} (D(\Delta\eta) - D^\dagger(\Delta\eta)) (\sigma_{y1} \sigma_{z2} - \sigma_{z1} \sigma_{y2}). \quad (4.111)$$

where $\Delta\eta = \eta_1 - \eta_2$.

Applying the approximation in equation (4.51), we have that

$$D(\Delta\eta) \approx \sum_{n,q} J_q(x_1 - x_2) |n+q\rangle \langle n|. \quad (4.112)$$

For the fiducial nEDM@SNS experimental parameters, the dressing parameters of the neutrons and ^3He atoms are, respectively,

$$x_n \approx -1.184,$$

$$x_3 \approx -1.317.$$

Because $|x_n - x_3|$ is significantly smaller than 1, the magnitude of $J_q(x_n - x_3)$ is small for $|q| > 0$. For example,

$$J_0(x_n - x_3) = 0.9956,$$

$$J_1(x_n - x_3) = 0.0662,$$

$$J_2(x_n - x_3) = 0.0022.$$

We therefore make the approximation that

$$U_1 U_2 (\sigma_1 \cdot \sigma_2) U_2^\dagger U_1^\dagger \approx \sigma_1 \cdot \sigma_2. \quad (4.113)$$

Next, we need to compute $e^{iH_U t} \sigma_1 \cdot \sigma_2 e^{-iH_U t}$, where we define $H_U = U_2 U_1 H U_1^\dagger U_2^\dagger$.

We make use of the identity

$$e^{iH_U t} \sigma_1 \cdot \sigma_2 e^{-iH_U t} = \sum_{n,n',s_1,s'_1,s_2,s'_2} \overline{|n, s_1, s_2\rangle} \langle n', s'_1, s'_2| \overline{\langle n_1, s_1, s_2| \sigma_1 \cdot \sigma_2 | n', s'_1, s'_2\rangle} e^{i\Delta\omega t}, \quad (4.114)$$

where as before $\Delta\omega$ is the frequency difference between the states $\overline{|n, s_1, s_2\rangle}$ and $\overline{|n', s'_1, s'_2\rangle}$. We make two simplifications to this expression. First, we are concerned

with the time-averaged signal, and therefore we ignore any terms where $\Delta\omega \neq 0$. Second, we consider only the leading-order terms in ω'_0/ω , and therefore make the approximation that $\overline{|n, s_1, s_2\rangle} \approx |n\rangle |s_1\rangle |s_2\rangle$. The expression in (4.114) then becomes

$$\begin{aligned} \left(e^{iH_U t} \sigma_1 \cdot \sigma_2 e^{-iH_U t} \right)_{\Delta\omega=0} &= \sum_{n, n', s_1, s'_1, s_2, s'_2} |n\rangle |s_1\rangle |s_2\rangle \langle n'| \langle s'_1| \langle s'_2| \langle n_1| \langle s_1| \langle s_2| \sigma_1 \cdot \sigma_2 |n'\rangle |s'_1\rangle |s'_2\rangle \\ &\quad \times \delta_{n, n'} \delta_{s_1+s_2, s'_1+s'_2}. \end{aligned} \quad (4.115)$$

It turns out that $\langle s_1| \langle s_2| \langle n| \sigma_1 \cdot \sigma_2 |n'\rangle |s'_1\rangle |s'_2\rangle = 0$ for $n \neq n'$, and for $s_1+s_2 \neq s'_1+s'_2$. Therefore,

$$\left(U_1 U_2 e^{iHt} \sigma_1 \cdot \sigma_2 e^{-iHt} U_2^\dagger U_1^\dagger \right)_{\Delta\omega=0} = \sigma_1 \cdot \sigma_2. \quad (4.116)$$

4.1.8.4 Calculation of Variance

If a random variable X can be written as a Taylor expansion in some small parameter λ , i.e.,

$$X = \sum_{n=1}^{\infty} \lambda^n Y_n. \quad (4.117)$$

where Y_n are random variables, then to second order in λ the variance of X is simply $\text{Var}(X) = \lambda^2 \text{Var}(Y)$. Therefore, to compute the variance of $T^\dagger AT$ we need only compute the variance of the terms of $T^\dagger AT$ which are linear in $\delta B(t)$. We get

$$\text{Var}_{cl}(T^\dagger AT) = \text{Var}_{cl} \left(i \int_0^t dt' V_I(t') A - A V_I(t') \right), \quad (4.118)$$

$$= \text{Var}_{cl} \left(i \int_0^t dt' (Q e^{-i\omega t} + Q^\dagger e^{i\omega t}) A - A (Q e^{-i\omega t} + Q^\dagger e^{i\omega t}) \right) \delta B(t'), \quad (4.119)$$

$$= \text{Var}_{cl}(i([Q, A]u + [Q^\dagger, A]u^*)). \quad (4.120)$$

$$(4.121)$$

Using the fact that $\langle u \rangle_{\delta B} = 0$, we re-write this expression as

$$\text{Var}_{cl}(T^\dagger AT) = \left\langle \left| [Q, A]u + [Q^\dagger, A]u^* \right|^2 \right\rangle_{\delta B}, \quad (4.122)$$

$$= \left\langle \left(\langle [Q, A] \rangle \langle [A^\dagger, Q^\dagger] \rangle + \langle [Q^\dagger, A] \rangle \langle [A^\dagger, Q] \rangle \right) uu^* \right\rangle_{\delta B}, \quad (4.123)$$

$$= 2 \left(\langle [Q, A] \rangle \langle [A^\dagger, Q^\dagger] \rangle + \langle [Q^\dagger, A] \rangle \langle [A^\dagger, Q] \rangle \right) \text{Re}(v). \quad (4.124)$$

If A is Hermitian, then this can be further simplified to

$$\text{Var}_{cl}(T^\dagger AT) = 4 \langle [Q^\dagger, A] \rangle \langle [A, Q] \rangle \text{Re}(v). \quad (4.125)$$

4.1.8.5 Signal Variance Calculation

For $Q_{\omega'} = \frac{\gamma_1}{2}\sigma_{-,1} + \frac{\gamma_2}{2}\sigma_{-,2}$, we need to calculate

$$[\vec{\sigma}_1 \cdot \vec{\sigma}_2, \sigma_{-,1}] = \sigma_{z1}\sigma_{x2} - \sigma_{x1}\sigma_{z2} + i(\sigma_{y1}\sigma_{z2} - \sigma_{z1}\sigma_{y2}). \quad (4.126)$$

Then

$$[\vec{\sigma}_1 \cdot \vec{\sigma}_2, \frac{\gamma_1}{2}\sigma_{-,1} + \frac{\gamma_2}{2}\sigma_{-,2}] = \frac{1}{2}(\gamma_1 - \gamma_2)(\sigma_{z1}\sigma_{x2} - \sigma_{x1}\sigma_{z2} + i(\sigma_{y1}\sigma_{z2} - \sigma_{z1}\sigma_{y2})). \quad (4.127)$$

Since $\vec{\sigma}_1 \cdot \vec{\sigma}_2$ is Hermitian, we compute the magnitude squared of the above quantity to find $\text{Var}_{cl}(T^\dagger \vec{\sigma}_1 \cdot \vec{\sigma}_2 T)$:

$$\text{Var}_{cl}(T^\dagger \vec{\sigma}_1 \cdot \vec{\sigma}_2 T)_{\omega'_0} = 4 \left\langle [Q_{\omega'_0}^\dagger, A] \right\rangle \left\langle [A, Q_{\omega'_0}] \right\rangle \text{Re}(v), \quad (4.128)$$

$$= \frac{1}{2}(\gamma_1 - \gamma_2)^2 \left(\langle \sigma_{z1}\sigma_{x2} - \sigma_{x1}\sigma_{z2} \rangle^2 + \langle \sigma_{y1}\sigma_{z2} - \sigma_{z1}\sigma_{y2} \rangle^2 \right) S(\omega'_0)t, \quad (4.129)$$

$$= \frac{1}{2}(\gamma_1 - \gamma_2)^2 |\hat{z} \times \{ \langle \vec{\sigma}_1 \rangle \times \langle \vec{\sigma}_2 \rangle \}|^2 S(\omega'_0)t. \quad (4.130)$$

For $Q_\omega = \frac{\gamma_1 J_1(x_1)\omega_1}{2\omega}\sigma_{z1} + \frac{\gamma_2 J_1(x_2)\omega_2}{2\omega}\sigma_{z2}$, we get

$$[\vec{\sigma}_1 \cdot \vec{\sigma}_2, \sigma_{z1}] = 2i(\sigma_{x1}\sigma_{y2} - \sigma_{y1}\sigma_{x2}). \quad (4.131)$$

Then

$$\left[\vec{\sigma}_1 \cdot \vec{\sigma}_2, \frac{\gamma_1 J_1(x_1)\omega_1}{2\omega}\sigma_{z1} + \frac{\gamma_2 J_1(x_2)\omega_2}{2\omega}\sigma_{z2} \right] = i \left(\frac{\gamma_1 J_1(x_1)\omega_1 - \gamma_2 J_1(x_2)\omega_2}{\omega} \right) (\sigma_{x1}\sigma_{y2} - \sigma_{y1}\sigma_{x2}). \quad (4.132)$$

Thus

$$\text{Var}_{cl}(T^\dagger \vec{\sigma}_1 \cdot \vec{\sigma}_2 T)_\omega = 4 \left\langle [Q_\omega^\dagger, A] \right\rangle \left\langle [A, Q_\omega] \right\rangle \text{Re}(v), \quad (4.133)$$

$$= 2 \left(\frac{\gamma_1 J_1(x_1)\omega_1 - \gamma_2 J_1(x_2)\omega_2}{\omega} \right)^2 \langle \sigma_{x1}\sigma_{y2} - \sigma_{y1}\sigma_{x2} \rangle^2 S(\omega)t, \quad (4.134)$$

$$= 2 \left(\frac{\gamma_1 J_1(x_1)\omega_1 - \gamma_2 J_1(x_2)\omega_2}{\omega} \right)^2 |\hat{z} \cdot \{ \langle \vec{\sigma}_1 \rangle \times \langle \vec{\sigma}_2 \rangle \}|^2 S(\omega)t. \quad (4.135)$$

4.1.8.6 Initial State of Magnetic Field

In this analysis, we assume that the magnetic field oscillates sinusoidally with amplitude B_1 . However, this leaves some freedom in choosing the phase of magnetic

field at time $t = 0$. A magnetic field $B(t) = B_1 \cos(\omega t - \phi)$ corresponds to the coherent state $|\sqrt{\lambda} e^{i\phi}\rangle$ where λ is the average photon number. Because calculations in this work take place in the displaced basis (which is obtained by applying the unitary transformation U to operators in the lab frame), we must also apply U to the initial state $|\alpha\rangle |s\rangle$. We first show that for $|\eta| \ll 1$ and $|\alpha| \gg 1$,

$$D(\eta) |\alpha\rangle \approx e^{\eta\alpha^* - \eta^*\alpha} |\alpha\rangle. \quad (4.136)$$

We use the identity that for any complex numbers α and η ,

$$D(\eta)D(\alpha) = e^{\eta\alpha^* - \eta^*\alpha} D(\alpha)D(\eta). \quad (4.137)$$

Therefore,

$$D(\eta) |\alpha\rangle = D(\eta)D(\alpha) |0\rangle, \quad (4.138)$$

$$= e^{\eta\alpha^* - \eta^*\alpha} D(\alpha)D(\eta) |0\rangle, \quad (4.139)$$

$$= e^{\eta\alpha^* - \eta^*\alpha} D(\alpha) |\eta\rangle. \quad (4.140)$$

Coherent states are given explicitly in terms of the Fock states by

$$|\eta\rangle = e^{-|\eta|^2/2} \sum_{k=0}^{\infty} \frac{\eta^k}{\sqrt{k!}} |k\rangle. \quad (4.141)$$

For $\eta \ll 1$ all but the $k = 0$ terms are small, and so we make the approximation $|\eta\rangle \approx |0\rangle$. From this we conclude that

$$D(\eta) |\alpha\rangle \approx e^{\eta\alpha^* - \eta^*\alpha} D(\alpha) |0\rangle, \quad (4.142)$$

$$\approx e^{\eta\alpha^* - \eta^*\alpha} |\alpha\rangle. \quad (4.143)$$

Now we examine what happens when we apply

$$U = D(\eta) |+_x\rangle \langle+_x| + D(\eta)^\dagger |-_x\rangle \langle-_x|, \quad (4.144)$$

where $\eta = \Omega/2\omega$ to the state $|\alpha\rangle |s\rangle$. We get

$$U |\alpha\rangle |s\rangle = D(\eta) |\alpha\rangle |+_x\rangle \langle+_x|s\rangle + D(\eta)^\dagger |\alpha\rangle |-_x\rangle \langle-_x|s\rangle, \quad (4.145)$$

$$= |\alpha\rangle \left(e^{i\theta} |+_x\rangle \langle+_x|s\rangle + e^{-i\theta} |-_x\rangle \langle-_x|s\rangle \right), \quad (4.146)$$

$$= |\alpha\rangle \left(e^{i\theta} |+_x\rangle \langle+_x| + e^{-i\theta} |-_x\rangle \langle-_x| \right) |s\rangle, \quad (4.147)$$

$$= |\alpha\rangle e^{i\theta\sigma_x} |s\rangle, \quad (4.148)$$

where $\theta = \text{Im}(2\eta^*\alpha)$. We thus see that an initial phase ϕ corresponds to a rotation by an angle θ about the x axis.

4.1.8.7 Robust dressing relaxation due to a relatively strong gradient field

The dynamics of the system can be described by

$$\dot{U}_I = e^{i\varphi_x \sigma_x / 2} \omega_0 \cdot \sigma e^{-i\varphi_x \sigma_x / 2} U_I, \quad (4.149)$$

where

$$\varphi_x = \gamma_i \int_0^t B_1(t') dt'. \quad (4.150)$$

Within this interaction picture we can write longitudinal relaxation as the decay rate of $\langle \sigma_x \rangle$, which is given by time-dependent perturbation theory as

$$\langle \sigma_x \rangle = 1 - \gamma^2 \text{Re} \left\{ \int_0^t dt' \int_0^{t'} dt'' e^{-i\varphi_x(t')} e^{i\varphi_x(t'')} \langle \delta B(t') \delta B(t'') \rangle \right\}, \quad (4.151)$$

where $\delta B(t)$ in this case represents fluctuations in either B_z or B_y . In writing this approximation, we have taken the static uniform field to be small, which allows us to neglect the contribution to T_1 relaxation from B_x . The integral can be evaluated in terms of the Fourier transform of $e^{-i\varphi_x(t)}$, given by

$$e^{-i\varphi_x(t)} = \sum_{n=-\infty}^{\infty} a_n e^{in\omega_m t}, \quad (4.152)$$

where ω_m is the modulation frequency. By the same reasoning as in section 4.1.2.1, we need only consider terms that are the same frequency. Thus we get

$$\langle \sigma_x \rangle = 1 - \gamma^2 \text{Re} \left\{ \int_0^t dt' \int_0^{t'} dt'' \sum_n |a_n|^2 e^{in\omega_m(t'-t'')} \langle \delta B(t') \delta B(t'') \rangle \right\}, \quad (4.153)$$

$$= 1 - 2t \sum_n |a_n|^2 S(n\omega_m). \quad (4.154)$$

Therefore we get

$$\frac{1}{T_1} = \frac{\gamma^2}{2} \sum_n |a_n|^2 S(n\omega_m), \quad (4.155)$$

$$\approx \gamma^2 (G_z^2 + G_y^2) D \sum_n \frac{|a_n|^2}{(\omega_m n)^2}, \quad (4.156)$$

where in the last step we took the diffusion approximation to evaluate $S(n\omega_m)$. If the system is not in the diffusion limit a more exact formulation can be implemented, for example, the spectrum of the correlation function presented in reference [SPG16] would allow accurate predictions from the ballistic through the diffusive regimes.

4.2 Magic Dressing

We also investigate another approach to modulated dressing, termed “magic dressing.” The concept is similar to that of the “magic wavelength” used to mitigate the AC Stark shift for atoms in an optical lattice [Kat+03]. Unlike robust dressing, where the objective is to maintain a polarized sample even in the presence of large gradients, magic dressing aims to simultaneously address fluctuations in the dressing and holding field.

4.2.1 Magic Dressing for a Single Species

For a single species (i.e., UCN or ^3He alone), This is accomplished by maintaining a dressing amplitude and frequency such that the dressing parameter $x = \gamma B_1/\omega$ is near a local extremum of the zeroth-order Bessel function. In this case, the dressed Larmor frequency is given by $\gamma B_0 J_0(x)$. Thus by operating near an extremum of J_0 (or equivalently, a zero of J_1), the precession frequency can be made insensitive to small changes in x . The single-species magic dressing technique works both in the case where the dressing field amplitude varies in space, as is the case for inhomogeneities caused by inaccuracies in dressing coil construction, or time as could be caused by fluctuations from the power supply controlling the current. This technique was studied using the spin-exchange optical pumping (SEOP) apparatus discussed in Chapter 6. As a cell containing both UCN and helium-3 was not available at the time of writing, only single-species magic dressing is investigated experimentally here.

4.2.2 Magic Dressing for Two Species

As described, magic dressing can only be used on one species at a time, as the zero of the J_1 Bessel function will differ for two species with different gyromagnetic ratios. Instead, one can apply a dressing field such that the difference between the derivatives of the two Bessel functions is zero. In other words, we want

$$\frac{d\omega_n}{dB_1} = \frac{d\omega_3}{dB_1}, \quad (4.157)$$

where ω_n and ω_3 are the dressed Larmor frequencies of the neutron and ^3He , respectively, given by

$$\omega_n = \gamma_n B_0 J_0(\gamma_n B_1/\omega) \quad (4.158)$$

$$\omega_3 = \gamma_3 B_0 J_0(\gamma_3 B_1/\omega). \quad (4.159)$$




Figure 4.10: Relative precession frequency for neutron and ^3He as a function of dressing amplitude (B_1). Dressing amplitudes where the frequency shifts due to changes in dressing amplitude are equal for neutron and ^3He are shown as dashed and dotted lines.

The first two solutions (including the trivial solution, $B_1 = 0$), are shown in figure 4.10. Neither solution is useful as a dressing scheme on its own, as neither solution corresponds to critical dressing, and so we cannot take advantage of the sensitivity gains afforded by critical dressing.

However, if we modulate between the two solutions shown in figure 4.10, then we can adjust the amount of time we spend at each solution to achieve both critical and magic dressing. Because the Bessel function model is only an approximate description of the dressed spin system [C C69], it is necessary to numerically optimize a parametrized modulation function. We proceed in two steps: first, we find a pulse shape for which the magic dressing condition is achieved (ignoring the critical dressing condition). Then, we adjust the time interval between pulses in order to achieve critical dressing. Because the B_1 field will be off between pulses (and because $x = 0$ trivially satisfies the 4.157), this second step will not disturb the magic dressing achieved in the first step. The modulation envelope is a smoothed square wave, given by

$$B_{\text{pulse}}(t) = \frac{A}{2} \sin(\omega t + \phi) \left\{ 1 + \tanh \left[\frac{3\tau}{2\pi\alpha} \left(\sin \left(\frac{2\pi}{3\tau} (t - 4\alpha) + \frac{\pi}{6} \right) - \frac{1}{2} \right) \right] \right\}, \quad (4.160)$$

where the parameters A , ϕ , τ , and α are optimized for. A is the average amplitude of the B_1 field, ϕ controls the relative phase between the modulation and the carrier

Table 4.3: Optimized magic dressing parameters for neutron and ^3He , assuming a dressing frequency of $\omega = 1$ kHz and a holding field strength of $B_0 = 50$ mG

Parameter	Value	Units
A	1.1176	Gauss
ϕ	1.3941	
τ	7.3122	ms
α	1.4754	ms
t_{rest}	11.149	ms

wave, τ controls the duration of the pulse, and α controls the time taken to go between the maximum and minimum amplitude. To optimize these parameters, we must define an objective function. The objective function has three terms: first, the deviation of the neutron spin vector out of the plane of precession; second, the deviation of the ^3He spin vector out of the plane of precession; and third, a term proportional to the derivative of the phase difference between the neutron and ^3He with respect to dressing amplitude. The first two terms are computed by integrating the Bloch equations and taking the z -component of the spin vectors at the end of the simulation time. The final term is computed by running two additional simulations, with slightly increased and slightly decreased dressing amplitude. In summary, the objective function is given by

$$\text{Obj}(A, \phi, \tau, \alpha) = (\hat{z} \cdot \vec{\sigma}_n)^2 + (\hat{z} \cdot \vec{\sigma}_3)^2 + \|U_{n+}U_{3+}^T - U_{n-}U_{3-}^T\|^2, \quad (4.161)$$

where $\vec{\sigma}_n$ and $\vec{\sigma}_3$ are the neutron and ^3He spin vectors at the end of the simulation time, and the matrices U_{n+} (U_{n-}) and U_{3+} (U_{3-}) are the propagators for the neutron and ^3He , simulated with a 1% larger (smaller) value for A . The operator $\|\cdot\|^2$ is the Frobenius norm. The optimization was conducted using the SAMIN() simulated annealing algorithm available in the Optim.jl package in Julia, and the Bloch equations were integrated with a custom-made Julia package. The optimized parameter values are given in table 4.3.

With the pulse shape optimized, we next optimize the time between pulses t_{rest} , using binary search to ensure critical dressing. This is given by the parameter t_{rest} . The final, periodic pulse sequence then has the form

$$B_x(t) = B_{\text{pulse}}(t \% (t_{\text{rest}} + \tau)) \quad (4.162)$$

where $a \% b$ means to take the remainder of a when divided by b . The magnetic field strength as a function of time is shown for 100 ms in figure 4.11. To confirm that this pulse train reduces depolarization due to magnetic field amplitude fluctuations,




Figure 4.11: Two-species modulated magic dressing pulse with optimized parameters.




Figure 4.12: Heatmap of the frequency difference between neutron and ^3He as holding field strength and dressing amplitude are varied about the magic dressing parameters.

we simulate neutrons and ^3He spin vectors over time for various perturbations in B_0 and B_1 about the magic dressing parameters above. A heatmap of the frequency difference for these perturbations is shown in figure 4.12. It is worth noting that near the nominal magic dressing parameters (the center of the heatmap), the variation with respect to B_1 is quadratic, indicating that magic dressing has been achieved.

We can also compare the results of magic dressing with that of unmodulated critical dressing. This comparison is shown in figure 4.13. Implementing magic dressing




Figure 4.13: Comparison of the response of magic dressing and unmodulated critical dressing to changes in the dressing field amplitude.

decreases the amplitude dependence of the frequency difference by several orders of magnitude.

Unlike for robust dressing, the dressed Larmor frequency (and thus the effective gyromagnetic ratio) for spins undergoing magic dressing is nonzero. Thus, two-species magic dressing is suitable for an nEDM measurement when a comagnetometer species is present, as is the case for nEDM@SNS [Ahm+19]. This allows one to reap the benefits of critical dressing, while loosening the amplitude stability requirement imposed on the dressing field. This can make it easier to procure a power supply that can provide the necessary current. Should this technique be employed for measurement, then the capture rate signal, being related to the angle between the neutron and ^3He , will oscillate between two values as the dressing field oscillates between the two solutions shown in 4.10. We illustrate this in figure 4.14. The dot product makes a triangle wave in time, which is useful as this can be repurposed as a modulation scheme of the type discussed in 3. If instead of modulating about $\theta_{n3} = \pi/2$ we modulate around $\theta_{n3} = 0$, then a nonzero nEDM would manifest as signal asymmetry that grows linearly with time.

4.3 Considerations for Spin Dressing Data Analysis

As the datasets generated in the critical dressing mode differ substantially from those in the free precession mode, different data analysis techniques will need to be developed. In this section we explore some of the considerations in analyzing spin dressing data in the case where magnetic field fluctuations are present.




Figure 4.14: Dot product between the neutron and ^3He spin vectors undergoing two-species magic dressing. The product is evaluated at integer multiples of the magic dressing modulation period $\tau + t_{\text{rest}}$, and the initial angle between the two spin vectors is $\pi/2$.

4.3.1 Magnetic Field Fluctuations as a Random Walk

Unlike the Poisson statistics from photon counting, fluctuations in magnetic field amplitude affect the relative phase between the ^3He and neutron spin vectors, which in turn leads to uncertainties that are proportional to the time integral of the fluctuations. For example, adding a time-dependent magnetic field fluctuation $\Delta B(t)$ to equation 3.8 results in

$$\theta_{n3}(t) = \pm\phi_0 + \frac{2J_0(x_n)d_nEt}{\hbar} + \int_0^t dt' (\gamma_n J_0(x_n) - \gamma_3 J_0(x_3)) \Delta B(t'), \quad (4.163)$$

provided that $\Delta B(t)$ is slow enough that the Bessel function approximation is valid. If we further assume that the fluctuations have mean zero and are uncorrelated on time scales over which the asymmetry measurement takes place (~ 1 Hz [Ahm+19]), we can treat the phase as having a random walk term. We thus consider a discrete random walk $y_i, i = 0 \dots N$, such that

$$y_0 = 0 \quad (4.164)$$

$$y_i = y_{i-1} + d + \epsilon_i \quad (4.165)$$

where the random variables ϵ_i are independent and normally distributed with variance ϵ^2 , i.e., $\epsilon_i \sim \mathcal{N}(0, \epsilon^2)$. The variance of ϵ is proportional to the variance of the magnetic field fluctuations $\Delta B(t)$. Now suppose that at each time i , we perform a

measurement z_i of y_i , and that these measurements are subject to Gaussian noise, i.e.,

$$z_0 = 0 \quad (4.166)$$

$$z_i = y_i + \sigma_i \quad (4.167)$$

where $\sigma_i \sim \mathcal{N}(0, \sigma^2)$. For the nEDM@SNS experiment, where counting statistics dominate, σ^2 should be proportional to the variance of the Poisson distribution. Our task is to determine a weighting h_i such that the estimator

$$\hat{d} \equiv \sum_i h_i z_i \quad (4.168)$$

estimates the true drift d with minimum uncertainty.

4.3.2 Matrix Formulation

This problem is easier if we work in terms of the difference of consecutive measurements. Define

$$x_i \equiv z_i - z_{i-1} \quad , \quad i \in 1 \dots N. \quad (4.169)$$

Then $\hat{d}$ is given by

$$\hat{d} \equiv \sum_i a_i x_i \quad (4.170)$$

where a_i are a new set of coefficients which can be related to h_i via a linear transformation. Clearly, if $\hat{d}$ is to be an estimator, it must satisfy

$$\mathbb{E}[\hat{d}] = d. \quad (4.171)$$

Therefore

$$d = \mathbb{E} \sum_i a_i x_i \quad (4.172)$$

$$= \mathbb{E} \sum_i a_i (z_i - z_{i-1}) \quad (4.173)$$

$$= \mathbb{E} \sum_i a_i (d + \epsilon_i + \sigma_i - \sigma_{i-1}) \quad (4.174)$$

$$= \mathbb{E} \sum_i a_i d \quad (4.175)$$

$$1 = \sum_i a_i, \quad (4.176)$$

where we have used the fact that σ_i has zero mean. To minimize the uncertainty, we minimize the variance of $\hat{d}$.

$$E(\hat{d} - d)^2 = E\left(\sum_i a_i(\epsilon_i + \sigma_i - \sigma_{i-1})\right)^2 \quad (4.177)$$

$$= E\left(\sum_i a_i \epsilon_i + \sum_i (a_i - a_{i+1})\sigma_i\right)^2 \quad (4.178)$$

$$= \epsilon^2 \sum_i a_i^2 + \sigma^2 \sum_i (a_i - a_{i+1})^2, \quad (4.179)$$

where we use the fact that the ϵ_i and σ_i are independent of each other (so that $E[\epsilon_i \sigma_j] = 0 \quad \forall i, j$). The objective then is to minimize the objective 4.179 subject to the constraint 4.176. Equation 4.179 is a quadratic form, so we can express it as a matrix multiplication. Define the matrices

$$U = \begin{bmatrix} 1 & -1 & 0 & & \\ -1 & 2 & -1 & \dots & \\ 0 & -1 & 2 & & \\ & \vdots & & \ddots & \\ & & & & 2 & -1 \\ & & & & -1 & 2 \end{bmatrix}$$

$$c = \begin{bmatrix} 1 & 1 & \dots & 1 \end{bmatrix}^T$$

$$a = \begin{bmatrix} a_1 & \dots & a_n \end{bmatrix}^T.$$

With these definitions, we need to solve the optimization problem

$$\min_a a^T (\sigma^2 U + \epsilon^2 I_n) a \quad (4.180)$$

$$\text{s.t. } c^T a = 1, \quad (4.181)$$

where I_n is the $n \times n$ identity matrix. The optimal weighting can be found with the method of Lagrange multipliers. The minimizing solution and corresponding uncertainty are

$$a = \frac{(\sigma^2 U + \epsilon^2 I_n)^{-1} c}{c^T (\sigma^2 U + \epsilon^2 I_n)^{-1} c} \quad (4.182)$$

$$\text{Var}(\hat{d}) = \frac{1}{c^T (\sigma^2 U + \epsilon^2 I)^{-1} c}. \quad (4.183)$$




Figure 4.15: Optimal weights for selected values of α . The upper plot shows the weights in the x_i basis, while the lower shows the same weights in the z_i basis.

Expressing the optimal weighting in this form allows for a straightforward method to compute the nEDM sensitivity numerically for a given set of parameters. Further, this method can be generalized to incorporate time-varying measurement noise by allowing σ^2 to depend on i in equation (4.179). Computing these weights is equivalent to constructing a matched filter [Tur60]; the matrix $(\sigma^2 U + \epsilon^2 I_n)$ is just the covariance matrix in the x_i basis, and the vector c can be viewed as the signal template in this same basis.

Using equation 4.182, optimal weights are computed for some selected values of α , with $N = 1000$. These are shown in figure 4.15. From the optimal weights in the

z_i basis one can see that measurements later in the measurement period are heavily weighted. This is a consequence of the so-called “lever-arm” effect, where later measurements are more sensitive to the nEDM, owing to the fact that there will have been more time for a linear-in-E phase to accumulate.

4.3.3 Continuum Limit

Next, we consider some limiting cases of the above results to obtain a quantitative description of how magnetic field fluctuations affect nEDM spin dressing data analysis. This will also allow for a more straightforward comparison with the results of [Ahm+19]. While the inverse of $(\sigma^2 U + \epsilon^2 I_n)$ is not trivial to compute analytically, in the continuum limit (i.e., many asymmetry measurements) we can make some progress. We expect this limit to be an appropriate description of the nEDM@SNS critical spin dressing parameters, as the measurement time is nominally 1000 s, while the modulation period for forming the signal asymmetry is on the order of 1 s. Let T be the duration of the random walk, and define $\Delta t \equiv T/N$ as the time difference between consecutive measurements. In the limit of small Δt (that is, many measurements), we define the continuous version of our estimator

$$\hat{d} \equiv \frac{1}{\Delta t} \int_0^T a(t)x(t)dt, \quad (4.184)$$

where $x(t)$ is the difference between consecutive measurements as in the discrete case. We approximate the summation in (4.179) as

$$\text{Var}(\hat{d}) = \sigma^2 \sum_i (a_i - a_{i+1})^2 + \epsilon^2 \sum_i a_i^2 \quad (4.185)$$

$$= \left[\sigma^2 \Delta t \sum_i \left(\frac{a_i - a_{i+1}}{\Delta t} \right)^2 + \frac{\epsilon^2}{\Delta t} \sum_i a_i^2 \right] \Delta t \quad (4.186)$$

$$\approx \int_0^T \left[\sigma^2 \Delta t \left(\frac{da}{dt} \right)^2 + \frac{\epsilon^2}{\Delta t} a^2 \right] dt, \quad (4.187)$$

and likewise the constraint (4.176) becomes

$$1 = \sum_i a_i \quad (4.188)$$

$$\approx \int_0^T \frac{a}{\Delta t} dt. \quad (4.189)$$

However in the continuous case, the drift d between consecutive measurements becomes infinitesimal, and so it is more natural to consider an estimator for $d/\Delta t$

— effectively the slope of the random walk. In that case we should take the normalization condition to be

$$1 = \int_0^T a dt, \quad (4.190)$$

This gives us a new optimization problem,

$$\min_a \int_0^T \left[\sigma^2 \Delta t \left(\frac{da}{dt} \right)^2 + \frac{\epsilon^2}{\Delta t} a^2 \right] dt \quad (4.191)$$

$$\text{s.t. } \int_0^T a dt = 1 \quad (4.192)$$

We can solve this problem with calculus of variations. The corresponding Lagrangian is

$$\mathcal{L} = (\sigma^2 \Delta t) \dot{a}^2 + \frac{\epsilon^2}{\Delta t} a^2 + 2\lambda(a - 1), \quad (4.193)$$

where λ is a Lagrange multiplier. Applying the Euler-Lagrange equation yields a differential equation for $a(t)$:

$$\frac{\partial \mathcal{L}}{\partial a} = \frac{d}{dt} \frac{\partial \mathcal{L}}{\partial \dot{a}} \quad (4.194)$$

$$\frac{\epsilon^2}{\Delta t} a + \lambda = (\sigma^2 \Delta t) \ddot{a}. \quad (4.195)$$

Examining the first and last rows of the matrix $(\sigma^2 U + \epsilon^2 I)$, we see that the boundary conditions are

$$\dot{a}(0) = 0 \quad (4.196)$$

$$a(T) = 0. \quad (4.197)$$

Solving for a , we get

$$a(t) = \beta \left(1 - \frac{\cosh(\alpha t/T)}{\cosh(\alpha)} \right) \quad (4.198)$$

where

$$\alpha \equiv \frac{N\epsilon}{\sigma} \quad (4.199)$$

$$\beta \equiv \frac{1}{T(1 - \tanh(\alpha)/\alpha)}. \quad (4.200)$$

The variance in $\hat{d}$ is then given by

$$\text{Var}(\hat{d}) = \int_0^T dt (\sigma^2 \Delta t) \dot{a}^2 + \frac{\epsilon^2}{\Delta t} a^2 \quad (4.201)$$

$$= \beta^2 \int_0^T dt (\sigma^2 \Delta t) \left(\frac{\alpha \sinh(\alpha t/T)}{T \cosh(\alpha)} \right)^2 + \frac{\epsilon^2}{\Delta t} \left(1 - \frac{\cosh(\alpha t/T)}{\cosh(\alpha)} \right)^2 \quad (4.202)$$

$$= \beta^2 \frac{\epsilon^2}{\Delta t} \int_0^T dt \left(\frac{\sinh(\alpha t/T)}{\cosh(\alpha)} \right)^2 + \left(1 - \frac{\cosh(\alpha t/T)}{\cosh(\alpha)} \right)^2 \quad (4.203)$$

$$= \beta^2 \frac{\epsilon^2}{\Delta t} T \left(1 - \frac{\tanh(\alpha)}{\alpha} \right) \quad (4.204)$$

$$= \beta^2 \frac{\epsilon^2}{\Delta t}. \quad (4.205)$$

We now consider the cases where either the magnetic field fluctuations (ϵ) or the counting statistics (σ) dominate the other.

1. $\epsilon \rightarrow 0$.

Then $\alpha \rightarrow 0$, and so

$$\lim_{\epsilon \rightarrow 0} \beta = \frac{1}{T} \lim_{\epsilon \rightarrow 0} \frac{1}{1 - \tanh(\alpha)/\alpha} \quad (4.206)$$

$$= \frac{1}{T} \lim_{\epsilon \rightarrow 0} \frac{1}{1 - (\alpha - \alpha^3/3)/\alpha} \quad (4.207)$$

$$= \frac{3}{T\alpha^2} \quad (4.208)$$

$$= \frac{3\sigma^2}{TN^2\epsilon^2}. \quad (4.209)$$

And therefore

$$\lim_{\epsilon \rightarrow 0} \text{Var}(\hat{d}) = \frac{3\sigma^2 \Delta t}{T^3}. \quad (4.210)$$

2. $\sigma \rightarrow 0$.

Then $\alpha \rightarrow \infty$, and so

$$\lim_{\sigma \rightarrow 0} \beta = \frac{1}{T}, \quad (4.211)$$

$$\lim_{\sigma \rightarrow 0} \text{Var}(\hat{d}) = \frac{\epsilon^2}{T\Delta t}. \quad (4.212)$$

From this we see that in the $\epsilon \rightarrow 0$ limit, the squared uncertainty scales as $1/T^3$. This is again a consequence of the “lever-arm” effect; extending the measurement time does not only increase the number of measurements taken, but also later measurements are more sensitive to the nEDM. For the $\sigma \rightarrow 0$ limit, phase accumulation

hinders us instead, because the phase from the magnetic field fluctuations (recall equation 4.163) also accumulates over time. This leads to $1/T$ scaling. This suggests that for very long measurements, magnetic field noise will begin to dominate over counting statistics. It is unlikely that this effect will matter for B_0 fluctuations, as the B_0 current is near-constant below the superconducting transition when in persistent current mode [Ahm+19], but may be relevant for the B_1 spin dressing field if the power supply driving it causes drifts in amplitude over time.

In summary, the optimal estimator for d , the slope of the random walk process, is given by

$$\hat{d} = \frac{1}{\Delta t} \int_0^T a(t)x(t)dt \quad (4.213a)$$

$$a(t) = \beta \left(1 - \frac{\cosh(\alpha t/T)}{\cosh(\alpha)} \right) \quad (4.213b)$$

$$\alpha \equiv \frac{N\epsilon}{\sigma} \quad (4.213c)$$

$$\beta \equiv \frac{1}{T(1 - \tanh(\alpha)/\alpha)} \quad (4.213d)$$

$$\text{Var}(\hat{d}) = \beta \frac{\epsilon^2}{\Delta t}. \quad (4.213e)$$

4.3.4 Application to nEDM

Finally, we compare the results derived here with the estimates given by [Ahm+19]. We will work in the Chibane limit, where $\dot{N}_B \rightarrow 0$, $\epsilon_\beta \rightarrow 0$, $\Gamma \rightarrow 0$. We begin by adding a random walk term to equation 2.19 of [Ahm+19]

$$\theta_{n3}(t) = \pm\phi_0 + \left(\frac{2J_0(x)Ed_n t}{\hbar} + W(t) \right), \quad (4.214)$$

where $W(t)$ is a Brownian motion such that $W(t) - W(s) \sim \mathcal{N}(0, \gamma^2(t-s))$. The time-dependent asymmetry then becomes

$$A_{ds}(t_i) = \frac{\left(\frac{N_0\epsilon_3}{\tau_3} \right) P \sin \phi_0 \left[\frac{2J_0(x_c)d_n|\vec{E}|t_i}{\hbar} + W(t) \right] e^{-\Gamma t_i}}{\left(\frac{N_0\epsilon_\beta}{\tau_\beta} \right) e^{-\Gamma t_i} + \left(\frac{N_0\epsilon_3}{\tau_3} \right) (1 - P \cos \phi_0) e^{-\Gamma t_i} + \dot{N}_B} \quad (4.215)$$

$$\equiv g(t_i) \left(d_n t_i + \frac{\hbar}{2J_0 E} W(t) \right), \quad (4.216)$$

where I have defined

$$g(t_i) \equiv \frac{\left(\frac{N_0\epsilon_3}{\tau_3} \right) P \sin \phi_0 \left[\frac{2J_0(x_c)|\vec{E}|t_i}{\hbar} \right] e^{-\Gamma t_i}}{\left(\frac{N_0\epsilon_\beta}{\tau_\beta} \right) e^{-\Gamma t_i} + \left(\frac{N_0\epsilon_3}{\tau_3} \right) (1 - P \cos \phi_0) e^{-\Gamma t_i} + \dot{N}_B}, \quad (4.217)$$

and the uncertainty on A_{ds} is

$$\sigma_A^2 = \frac{1}{\Delta t \left[\left(\frac{N_0 \epsilon_\beta}{\tau_\beta} \right) e^{-\Gamma t_i} + \left(\frac{N_0 \epsilon_3}{\tau_3} \right) (1 - P \cos \phi_0) e^{-\Gamma t_i} + \dot{N}_B \right]}. \quad (4.218)$$

The parameter we want to measure is d_n , so to make it easier to apply our previous result, we define

$$B(t_i) \equiv \frac{A_{ds}(t_i)}{g(t_i)} = d_n t_i + \frac{\hbar}{2J_0 E} W(t). \quad (4.219)$$

The corresponding uncertainty on measurements of B is then

$$\sigma_B^2 = \frac{\sigma_A^2}{g^2(t_i)} \quad (4.220)$$

$$= \frac{\left(\frac{N_0 \epsilon_\beta}{\tau_\beta} \right) e^{-\Gamma t_i} + \left(\frac{N_0 \epsilon_3}{\tau_3} \right) (1 - P \cos \phi_0) e^{-\Gamma t_i} + \dot{N}_B}{\left(\frac{N_0 \epsilon_3}{\tau_3} \right)^2 (P \sin \phi_0)^2 \left[\frac{2J_0(x_c) |\vec{E}|}{\hbar} \right]^2 e^{-2\Gamma t_i} \Delta t}. \quad (4.221)$$

$$(4.222)$$

Note that this differs from equation 2.28 of [Ahm+19] by a factor of 2, owing to the fact that Δt here is taken to be the duration of a complete modulation cycle (one measurement at each of $\pm \phi_0$). In the Chibane limit, σ_B^2 becomes constant with time:

$$\sigma_B^2 = \frac{1 - P \cos \phi_0}{(P \sin \phi_0)^2} \left(\frac{\hbar}{2J_0(x_c) |\vec{E}|} \right)^2 \frac{1}{\left(\frac{N_0 \epsilon_3}{\tau_3} \right) \Delta t}. \quad (4.223)$$

Now with the shuffling of parameters out of the way, we can apply our previous result with the substitutions

$$\sigma^2 \rightarrow \sigma_B^2 \quad (4.224)$$

$$\epsilon^2 \rightarrow \gamma^2 \Delta t \left(\frac{\hbar}{2J_0(x_c) |\vec{E}|} \right)^2. \quad (4.225)$$

The uncertainty is then given by

$$\text{Var}(\hat{d}_n) = \left(\frac{\hbar}{2J_0 E} \right)^2 \frac{\gamma^2}{1 - \tanh(\alpha)/\alpha}, \quad (4.226)$$

where

$$\alpha = TP \sin \phi_0 \sqrt{\left(\frac{N_0 \epsilon_3}{\tau_3} \right) \frac{1}{1 - P \cos \phi_0}}. \quad (4.227)$$

We can also compare the variance predicted here to the “naive” variance, which




Figure 4.16: nEDM Uncertainty for a single cell, single cycle, versus phase noise due to power supply.




Figure 4.17: Plot of naive and exact predictions of nEDM uncertainty.




Figure 4.18: Ratio of predictions (naive/exact).

would just be the sum of the two limiting cases. That is,

$$\text{Var}(\hat{d}) \approx \frac{\epsilon^2}{T\Delta t} + \frac{3\sigma^2\Delta t}{T^3}. \quad (4.228)$$

These comparisons are shown in figures 4.17 and 4.18.

Figure 4.16 shows the noise power at which the magnetic fluctuation effect becomes comparable to the noise due to counting statistics, at around 10^{-4} radians/ $\sqrt{s}$. From figure 4.18, we see that the naive estimate of the variance, which simply adds the two noise sources in quadrature, provides a result that is within 10% of the true variance. From these results, it is clear that magnetic field fluctuations cannot be neglected in the nEDM uncertainty, which underscores the importance of methods that can mitigate this type of noise, particularly in regard to the dressing field.

Chapter 5

SPIN INTEGRATION - COMPUTATIONAL TECHNIQUES

In this Chapter we discuss techniques for improving the speed or accuracy of numerical integration techniques applied to solving for the time-dependence of a spin under a (potentially time-varying) magnetic field. We will primarily deal with the Bloch equation,

$$\frac{d\vec{\sigma}}{dt} = \gamma \vec{\sigma} \times B(t) \quad (5.1)$$

where $\vec{\sigma}$ is the 3-component Bloch vector corresponding to the spin and $B(t)$ is the magnetic field vector. In section 5.1, we discuss the use of the Magnus expansion to accelerate the integration of the Bloch equation for individual spins. In section 5.2 we take a broader view, and develop an algorithm that can rapidly determine the average behavior of ensembles of spins. This section is adapted from the article “The semi-classical Floquet-Markov master equation for Monte Carlo spin integration,” [Tat25].

5.1 Commutator-Free Exponential Time-Propagator (CFET) Method

The standard tool for integrating the Bloch equations are the Runge-Kutta methods. These are a family of numerical integration schemes, all of which can be derived from a table of coefficients with certain properties [WNH93]. The simplest of these is Euler’s method, which for a given time step Δt estimates the future vector as

$$x(t + \Delta t) \approx x'(t)\Delta t + x(t). \quad (5.2)$$

By repeatedly iterating this procedure, one acquires an estimate of the trajectory $x(t)$ over time. The time step Δt can be made smaller until the desired accuracy is reached. More sophisticated methods, such as the standard 4th order Runge-Kutta method, improve the convergence by evaluating the derivative $x'(t)$ at points inside the time interval (i.e., between t and $t + \Delta t$). This allows for larger steps to be taken at the cost of computation time per step. Thus depending on the accuracy needed, a method of suitable order should be chosen in order to maximize computational efficiency. Generally, when smaller error is required, higher order methods will yield better results. While the Runge-Kutta methods are an effective choice for a variety of ordinary differential equations, further optimizations can be made in

cases where the structure of the equations allows for additional knowledge about the trajectory. In particular in the case of the Bloch equations (ignoring for now the effect of relaxation), the time evolution of the spin vector $\vec{\sigma}$ is unitary, and therefore the length of the spin vector never changes. We therefore consider a method using the Magnus expansion [Bla+09], which intrinsically preserves the length of the spin vector.

5.1.1 The Magnus Expansion

Consider a linear ODE giving the time dependence of a vector $x(t)$:

$$x'(t) = A(t)x(t) \quad (5.3)$$

where A is some time-dependent matrix. The Magnus expansion gives the propagator $U(t)$ as

$$U(t) = e^{\sum \Omega_i(t)} \quad (5.4)$$

where the $\Omega_i(t)$ are time-ordered integrals of increasingly complicated commutators of A . Here are the first few terms [Bla+09]:

$$\Omega_1(t) = \int_0^t dt_1 A(t_1) \quad (5.5)$$

$$\Omega_2(t) = \int_0^t dt_1 \int_0^{t_1} dt_2 [A(t_1), A(t_2)] \quad (5.6)$$

$$\Omega_3(t) = \int_0^t dt_1 \int_0^{t_1} dt_2 \int_0^{t_2} dt_3 [A(t_1), [A(t_2), A(t_3)]] + [[A(t_1), A(t_2)], A(t_3)]. \quad (5.7)$$

To construct a numerical integration scheme from 5.5, we must first re-write the terms of Magnus expansion in some suitable polynomial basis, then find a procedure to numerically approximate the coefficients in this basis. The authors of [AF11] work in the basis of Legendre polynomials, which have the property

$$\int_0^1 P_m(x) P_n(x) dx = \frac{1}{2n+1} \delta_{mn}. \quad (5.8)$$

If we now expand $A(t)$ on a fixed interval $[0, \delta t]$ in terms of the Legendre polynomials, i.e.,

$$A(t) = \frac{1}{\delta t} \sum_{n=1}^N A_n P_{n-1} \left(\frac{t}{\delta t} \right) + O(\delta t^{N+1}). \quad (5.9)$$

It is important to note that the term A_n is of order δt^n , which will be important when we plug this into the Magnus expansion.

Now we can re-write each term of the Magnus expansion in terms of the coefficients A_i . For example,

$$\Omega_1(\delta t) = \int_0^{\delta t} dt_1 A(t_1) \quad (5.10)$$

$$= \frac{1}{\delta t} \int_0^{\delta t} dt_1 \sum_{n_1=1}^{\infty} A_{n_1} P_{n_1-1} \left(\frac{t_1}{\delta t} \right) \quad (5.11)$$

$$= \sum_{n_1=1}^{\infty} \left(\int_0^1 dx_1 P_{n_1-1}(x_1) \right) A_{n_1} \quad (5.12)$$

$$= \sum_{n_1=1}^{\infty} \delta_{n_1-1,0} A_{n_1} \quad (5.13)$$

$$= A_1 \quad (5.14)$$

$$\Omega_2(\delta t) = \int_0^{\delta t} dt_1 \int_0^{t_1} dt_2 [A(t_1), A(t_2)] \quad (5.15)$$

$$= \frac{1}{2\delta t^2} \int_0^{\delta t} dt_1 \int_0^{t_1} dt_2 \left[\sum_{n_1} A_{n_1} P_{n_1-1} \left(\frac{t_1}{\delta t} \right), \sum_{n_2} A_{n_2} P_{n_2-1} \left(\frac{t_2}{\delta t} \right) \right] \quad (5.16)$$

$$= \frac{1}{2} \sum_{n_1, n_2} \left(\int_0^1 dx_1 \int_0^{x_1} dx_2 P_{n_1-1}(x_1) P_{n_2-1}(x_2) \right) [A_{n_1}, A_{n_2}] \quad (5.17)$$

$$= -\frac{1}{6} [A_1, A_2] \quad (5.18)$$

and so on. The calculation of these terms was done in [AF11], the first few terms of which are

$$\begin{aligned} \Omega(\delta t) &= A_1 - \frac{1}{6} [A_1, A_2] \\ &+ \frac{1}{60} [A_1, [A_1, A_3]] - \frac{1}{60} [A_2, [A_1, A_2]] + \frac{1}{360} [A_1, [A_1, [A_1, A_2]]] - \frac{1}{30} [A_2, A_3] \\ &+ O(\delta t^7). \end{aligned} \quad (5.19)$$

One can easily determine the order of a given nested commutator expression by adding up the subscript indices on the A_i , e.g., $[A_1[A_1, A_3]]$ is of order δt^5 . You might notice that many commutators are missing. This is due to two facts:

1. Time-reversal symmetry (i.e., we demand that $U(\delta t, 0) = U(0, \delta t)^{-1}$) ensures that there are no even-order terms.

2. The orthogonality of the Legendre polynomials causes a bunch of other commutators to drop out. The condition for this to happen is if there is an index n_k which exceeds the sum of the others by 2 or more.

5.1.2 The Baker-Cambell-Hausdorff (BCH) Formula

The final ingredient in the Magnus CFET method is the BCH formula, which allows one to combine exponentials. For two exponentials, the formula is:

$$\exp(A) \exp(B) = \exp\left(A + B + \frac{1}{2}[A, B] + \frac{1}{12}[A, [A, B]] - \frac{1}{12}[B, [A, B]] + \dots\right). \quad (5.20)$$

A Magnus integrator then takes the form

$$U(\delta t) \approx \prod_i \exp\left(\sum_j f_{ij} A_j\right) \quad (5.21)$$

where f_{ij} are coefficients chosen so that when the BCH formula is applied, it reproduces the Magnus expansion to some order. The simplest case is the 2nd order CFET, given by

$$U_{CF2:1}(\delta t) = \exp(A_1) \quad (5.22)$$

which clearly reproduces the first term in the Magnus expansion, since $A_1 = \Omega_1(\delta t)$. The first non-trivial case is the 4th-order CFET with two exponentials.

$$U_{CF4:2}(\delta t) = \exp(f_{1,1}A_1 + f_{1,2}A_2) \exp(f_{1,1}A_1 - f_{1,2}A_2). \quad (5.23)$$

You might notice that there are four coefficients, but only two are independent. This is due to time-reversal symmetry. To solve for the two remaining coefficients, we invoke the BCH formula and compare with the Magnus expansion. We get:

$$U_{CF4:2}(\delta t) = \exp\left(2f_{1,1}A_1 + [f_{1,1}A_1 + f_{1,2}A_2, f_{1,1}A_1 - f_{1,2}A_2] + O(\delta t^5)\right) \quad (5.24)$$

$$= \exp\left(2f_{1,1}A_1 - 2f_{1,1}f_{1,2}[A_1, A_2] + O(\delta t^5)\right). \quad (5.25)$$

Comparing with the Magnus expansion (equation 5.19), we see that we must take

$$f_{1,1} = \frac{1}{2} \quad (5.26)$$

$$f_{1,2} = \frac{1}{3}. \quad (5.27)$$

This is a bit of a special case, since there is only one solution for the coefficients. Generally, there will be many more coefficients than constraints, and some optimization can be done to try to find the best set of coefficients.

5.1.3 Gauss Quadrature

Up until now, we have assumed that the terms in the Legendre polynomial expansion ($A_1, A_2, \dots$) were known. But in general if we are integrating numerically, these cannot be calculated exactly. Instead, we should estimate these terms using some Gaussian quadrature. Gaussian quadrature is a technique that estimates an integral of a function over an interval as a weighted sum over the function evaluated at a number of sample points:

$$\int_{-1}^1 dt f(t) \approx \sum_{i=1}^n w_i f(t_i) \quad (5.28)$$

for some carefully chosen w_i and t_i , the values for which you can find online in various tables. The matrices A_n are given by

$$A_n = (2n - 1)\delta t \int_0^1 A(x\delta t) P_{n-1}(x) dx \quad (5.29)$$

which is an integral we can approximate with Gauss quadrature. However, we must be careful to use a quadrature with high enough order (i.e., with enough sample points) so that we do not compromise the desired accuracy of our integrator. As it turns out, for an N -th order CFET, a quadrature with $N/2 + 1$ points is sufficient. Thus for an 8th order integrator, we need only evaluate the function $A(t)$ at 5 points.

Interestingly, there does not seem to be any reason that we could not use *more* than the minimum required number of quadrature points (other than of course having to do more computation). Thus, if we have already computed the 5-point quadrature for the 8-th order CFET, we can reuse the *same* 5 evaluations of $A(t)$ (and thus the same matrices A_i) in order to compute a 6-th order CFET over the same time step. This is the technique we use to construct the adaptive stepsize routine; we compare the output of a 6th- and 8th-order CFET to estimate the local error.

5.1.4 Performance

The performance of the 8th-order CFET integrator for free precession mode and critical spin dressing mode is shown in figures 5.1 and 5.2, respectively. For each of these plots 2^{14} ^3He spins were simulated for 10 seconds with an electric field of 75 kV cm^{-1} . The temperature of the simulation was varied from 0.3 K to 0.5 K to study the effect of the mean free path on the performance of the integrator. The simulations are benchmarked against DOP853 [WNH93], a standard 8-th order Runge-Kutta algorithm. All simulations were conducted on the Georgia Tech PACE cluster [PAC17] using a single Nvidia A100 GPU node. As expected, increasing




Figure 5.1: Performance of the 8th-order Magnus CFET integrator compared to DOP853 for integrating the spins under free precession.




Figure 5.2: Performance of the 8th-order Magnus CFET integrator compared to DOP853 for integrating the spins in critical spin dressing mode.

the temperature, which shortens the mean free path, worsens the performance of all integrators. Nonetheless, the Magnus CFET integrator outperforms DOP853 in all cases, especially in the case of free precession, where the magnetic field is near-constant. It should be noted that the speed of the Magnus CFET method depends critically on the existence of an efficient method to compute the matrix exponentials in 5.21. No such algorithm exists for matrices of arbitrary size, but in the 3×3 case Rodrigues' rotation formula allows us to perform these computations rapidly and exactly.

We also investigate the behavior of the Magnus CFET integrator for fixed step size (as opposed to the adaptive step size algorithm used for the previous plots). Here, we




Figure 5.3: Error comparison of Magnus CFET, MagnusGL8 (an integrator available in the DifferentialEquations.jl [RN17] library), and DOP853. The simulation parameters used are the critical dressing parameters, with no gradients or electric field applied.

compare three fixed step size integrators: the 8th-order Magnus CFET integrator, an 8th-order Magnus integrator available in the DifferentialEquations.jl library of the Julia programming language [RN17], and lastly DOP853 with fixed step size. The comparison is shown in 5.3, and shows the Magnus CFET method results in a significantly lower error at every step size.

5.2 The semi-classical Floquet-Markov Master Equation for Monte Carlo Spin Integration

The section is adapted from the article “Semi-classical Floquet-Markov master equation for Monte Carlo spin integration,” [Tat25] for which I am the sole author. We employ Floquet theory, which handles periodic differential equations, combined with time-dependent perturbation theory to design an integrator that vastly outperforms conventional Runge-Kutta integration for the case of time-periodic Hamiltonians.

5.2.1 Introduction

The search for physics beyond the Standard Model on the precision frontier necessitates an equally precise understanding of the uncertainties and systematics of the measurement. For an experiment to measure a neutron electric dipole moment (nEDM) for example, detailed knowledge of the systematic frequency shifts associated with the electric field and magnetic field gradients is required [GL94]. This is facilitated through extensive use of Monte Carlo simulations to investigate the

impact of these effects on the experiment; for an nEDM experiment, this may be done by simulating the Bloch equations for a large ensemble neutron spins and observing the frequency shift in simulation. The standard approach for this is to use a numerical integration algorithm such as Runge-Kutta [WNH93] to integrate the Bloch equations for each simulated neutron. Computationally, this is challenging; as the precision increases, so too does the number of simulated particles needed to constrain the statistical uncertainty to the level of the experiment. This difficulty is compounded when time-dependent (dressing) magnetic fields are applied, which have the effect of drastically reducing the step size taken by numerical integrators. Here we propose an alternative approach, which may provide significant gains in performance in cases where the effects of magnetic gradients and electric fields are relatively small. The strategy is to use the Floquet-Markov master equation (FMME) [Blü+91], an equation which predicts the time evolution of a system subject to a time-periodic Hamiltonian perturbed by a small time-dependent noise term with known power spectrum. Advances in GPU hardware allow this power spectrum to be estimated efficiently with the Monte Carlo method.

5.2.2 Spin Physics and Frequency Shifts

Nuclear magnetic resonance (NMR)-based EDM experiments require at least two fields: a static (time-independent) magnetic field, which causes the nuclear spin to precess, and a strong electric field which modifies the precession frequency of the nuclei when coupled to an EDM. There may additionally be a time-dependent spin dressing magnetic field as was proposed for the nEDM experiment at the Spallation Neutron Source (nEDM@SNS), which can be used to improve sensitivity [Ahm+19]. Because an EDM measurement is a measurement of frequency, it is necessary to exclude spurious sources of frequency shifts. These can arise due to magnetic field gradients, as well as from velocity-dependent effects due to the presence of the electric field. These shifts can be broadly divided into three categories — shifts which are quadratic in gradient magnitude, shifts which are quadratic in electric field strength, and shifts which are linear in gradient and electric field strength. The last shift, often referred to as the “geometric phase” or as the “ $\mathbf{v} \times \mathbf{E}$ effect,” poses a great challenge for EDM measurement, as being linear in electric field strength, it can mimic a real EDM. This was first noted in [Pen+04b] and its effect characterized for several important trap geometries in [LG05]. These shifts have since been studied extensively, for example in [McG90], [Pig+15], [SPG16], or [TS22]. However, several common scenarios still elude theoretical analysis —

for example, no analytic expression is known for the frequency shift for the case where the spins undergo diffuse reflections from the walls of their cell, nor is there a solution incorporating analytically the effects of gravity, which are relevant for ultra-cold neutrons (UCNs). The crux of the difficulty here is in determining the power spectrum of the fluctuating magnetic field, which depends on the correlation functions of the particle trajectories. These correlation functions can only be determined analytically in certain special cases. As a result, significant effort has been expended to study these shifts in simulation: for example, see [Pig19], [Swa+18], or [Cla11]. These works motivate our search for a more efficient approach to integrating large spin ensembles.

5.2.3 Floquet Theory and the Floquet-Markov Master Equation

Floquet theory describes the time evolution of a quantum system under a time-periodic Hamiltonian in terms of quasi-energies and quasi-eigenstates, which are the time-periodic analogues of the energies and eigenstates of a time-independent Hamiltonian. Specifically, given a Hamiltonian for which $H(t) = H(t+T)$, Floquet's theorem [Flo83; GH98] states that solutions to Schrödinger's equation have the form

$$\Psi(t) = \sum_{\alpha} c_{\alpha} \Psi_{\alpha}(t) \quad (5.30)$$

$$= \sum_{\alpha} c_{\alpha} \exp(-i\epsilon_{\alpha}t/\hbar) \Phi_{\alpha}(t), \quad (5.31)$$

where $\Psi_{\alpha}(t)$ are the Floquet states, and $\Phi_{\alpha}(t)$ are the Floquet modes, which are time-periodic, i.e., $\Phi_{\alpha}(t) = \Phi_{\alpha}(t+T)$. ϵ_{α} are the quasi-energies. Note that the quasi-energies are only defined up to multiples of $2\pi\hbar/T$. The Floquet modes can be computed numerically in several ways, but here we take the approach used in the QuTIP Python Library [Lam+24], which is to diagonalize the one-period propagator $U(T)$. The Floquet modes evaluated at $t = 0$ are the eigenstates of the one-period propagator. If one now adds a noise term to the Hamiltonian, the time evolution will no longer be given by equation 5.30, but by working within the basis of Floquet states, one can derive a master equation that describes the non-unitary evolution of the density matrix to second order in the size of the noise term. After applying the Born-Markov and rotating wave approximations as in Redfield theory [Red57; Sli90] one arrives at the Floquet-Markov master equation (FMME) [Blü+91], which has the form,

$$\frac{d}{dt} \widetilde{\rho_{ab}} = R_{abcd} \widetilde{\rho_{cd}}, \quad (5.32)$$

where $\widetilde{\rho}_{ab}$ is the density matrix in the basis of the Floquet modes, and R_{abcd} is a time-independent tensor whose elements depend on the power spectrum $S(\omega)$ of the noise term and the Floquet diagonalization of the unperturbed Hamiltonian. In particular, the relevant frequencies for calculating R_{abcd} are of the form $\Delta_{\alpha\beta k} = (\epsilon_\alpha - \epsilon_\beta)/\hbar + k\omega$, with k an integer and $\omega = 2\pi/T$. The set of relevant frequencies depends only on the unperturbed Hamiltonian, while the power spectrum depends only on the noise term. Thus, the FMME allows us to split the task of simulating an ensemble of spins into two parts — first to compute the Floquet diagonalization, and second to estimate the power spectrum of the noise term.

5.2.4 Floquet Diagonalization

The full derivation of the FMME formalism for this case is described in the appendix. Here we describe the Floquet diagonalization, wherein we find the Floquet modes and quasi-energies, and with these determine the Fourier transforms of the perturbation operators in the Floquet basis. Note that this need only be done once for a given dressing pulse, regardless of how many particles are simulated. As stated previously, the approach to finding the Floquet modes follows the method used by QuTIP [Lam+24]. This algorithm uses the fact that the Floquet modes obey the relation

$$U(T)\Phi_\alpha(0) = \exp(-i\epsilon_\alpha T/\hbar)\Phi_\alpha(0) \quad (5.33)$$

where $U(T)$ is the one-period propagator (assuming the period is T), $\Phi_\alpha(0)$ is a Floquet state at $t = 0$, and ϵ_α is the corresponding quasi-energy. To find the Floquet states, we begin by solving numerically the Schrödinger equation

$$i\hbar \frac{dU(t)}{dt} = H(t)U(t) \quad (5.34)$$

with initial condition $U(0) = 1$ to find $U(T)$. We then diagonalize $U(T)$, taking the eigenvectors to be $\Phi_\alpha(0)$ and the phase of the eigenvalues to determine ϵ_α . From there we compute the tensor $X_{\alpha\beta k}$, the Fourier transform of the perturbing operators, defined by

$$\langle \Phi_\alpha(t) | \sigma | \Phi_\beta(t) \rangle = \sum_{k=-\infty}^{\infty} e^{ik\omega t} X_{\alpha\beta k}, \quad (5.35)$$

where σ is the perturbing operator (for example, σ is the Pauli spin matrix σ_x for a noise magnetic field in the x direction). We can invert this relation to find $X_{\alpha\beta k}$,

$$X_{\alpha\beta k} = \int_0^T e^{-ik\omega t} \langle \Phi_\alpha(t) | \sigma | \Phi_\beta(t) \rangle dt/T. \quad (5.36)$$

To evaluate this expression, we must compute the Floquet modes evaluated at t , rather than zero. The time evolution of the Floquet modes over the first period is given by

$$\Phi_\alpha(t) = e^{i\epsilon_\alpha t/\hbar} U(t) \Phi_\alpha(0), \quad (5.37)$$

where $U(t)$ is found by again integrating the Schrödinger equation numerically. From this, we can evaluate the integral in equation 5.36 by dividing the time interval from 0 to T into evenly spaced sub-intervals and approximating the integral as a sum. Note that the Floquet diagonalization must be computed before the spectrum estimation step, as the relevant frequencies are determined by the quasi-energies; in particular, the relevant frequencies are of the form

$$\Delta_{\alpha\beta k} = (\epsilon_\alpha - \epsilon_\beta)/\hbar + k\omega. \quad (5.38)$$

Lastly, as the index k spans all integers, for the sake of computation a finite range of k must be chosen. For this work we assume that the range $-5 \leq k \leq 5$ is sufficient, as according to the analysis in [TS22], the contribution to frequency shifts and relaxation for higher values of k (higher frequencies) are suppressed by increasing orders of Bessel functions. For example, we would expect a suppression of the $k = 5$ term relative to the $k = 1$ term by a factor $(J_5(x_3)/J_1(x_3))^2$, where $x_3 = \gamma_3 B_1/\omega$ is the helium-3 spin dressing parameter. Here γ_3 is the helium-3 gyromagnetic ratio, and B_1 and ω are the dressing field amplitude and frequency, respectively. For nEDM@SNS x_3 is approximately 0.95 [Ahm+19], leading to a suppression factor of $\sim 2 \times 10^{-7}$. Provided that the range of k is large enough to ensure that further terms have a negligible impact on the simulated results, it is beneficial to choose as small a range of k as possible, as increasing the range of k increases the computational complexity of the method. However, depending on the application other ranges may be chosen, for example if noise is known to be concentrated around a specific frequency, or if the Larmor and dressing frequency differ greatly from the parameters used here.

5.2.5 Spectrum Estimation

The standard technique for estimating the power spectrum from evenly-spaced samples is to use the periodogram, which is the magnitude-squared of the discrete Fourier transform (DFT). In order to reduce variance of the estimate, rather than taking the periodogram of one long time series, the periodograms of many shorter windows of the time series are averaged. This is known as Bartlett's method [Bar48]. As the Floquet-Markov master equation only requires the power spectrum to be known at

specific discrete frequencies, we can tune the parameters of the periodogram so that the output frequencies match those needed by the Floquet-Markov master equation. We investigate two algorithms to compute the averaged periodograms — the first is the Goertzel algorithm [Goe58], which is a two-stage filter that computes the DFT at a given frequency. By operating several of these filters in parallel, one can compute as many of the relevant frequencies as is desired. The second approach is to use the fast Fourier transform (FFT). In this case, we simulate an ensemble of particles and collect samples of the noise magnetic field at evenly spaced points in time. We then, at regular intervals, take the FFT magnitude squared of the collected samples. Notably, for both approaches the minimum window length is determined by the desired frequency resolution, while the minimum sampling rate is determined by Nyquist-Shannon sampling theorem [Sha49], which states that the sampling rate must be at least twice the highest frequency component contained in the signal of interest.

5.2.5.1 Complex Spectrum

One subtlety involved in this approach to spin integration is that we require not the ordinary power spectrum, defined as the Fourier transform of the correlation function but rather the complex power spectrum, defined as the one-sided Fourier integral of the correlation function. In other words, whereas normally one works with

$$S(\omega) = \int_{-\infty}^{\infty} \langle H(t)H(t+\tau) \rangle e^{-i\omega\tau} d\tau \quad (5.39)$$

the FMME instead requires

$$J(\omega) = \int_0^{\infty} \langle H(t)H(t+\tau) \rangle e^{-i\omega\tau} d\tau. \quad (5.40)$$

Of note is the fact that although $S(\omega)$ is purely real, $J(\omega)$ may in general be complex. Without this modification, the FMME would not predict frequency shifts. Because the complex spectrum is needed, both the Goertzel algorithm and FFT approaches will need to be modified to estimate $J(\omega)$ instead of $S(\omega)$. For the Goertzel algorithm, this is accomplished by maintaining a partial sum of $J(\omega)$ in memory which is updated after every step. This change preserves algorithm's linear-in-time complexity, though at the cost of significantly increased computation at each step. To compute $J(\omega)$ with the FFT algorithm we first compute the correlation function $\langle H(t)H(t+\tau) \rangle$ via an inverse Fourier transform, apply a mask in the time domain to remove the negative-time component, then transform back to the the frequency

domain to recover $J(\omega)$. In the following sections we describe both algorithms and their corresponding modifications in further detail.

5.2.5.2 Goertzel Algorithm

The Goertzel algorithm computes the DFT at a given frequency ω in two stages. Given samples x_n separated by time interval Δt , the first stage computes

$$s_n = x_n + 2 \cos(\omega \Delta t) s_{n-1} - s_{n-2}, \quad (5.41)$$

and the second stage is

$$y_n = s_n - e^{-i\omega \Delta t} s_{n-1}. \quad (5.42)$$

Here s_n is an intermediate sequence, used only to calculate y_n . It can be shown that y_n is related to the DFT at frequency ω , namely

$$y_n = e^{i\omega n \Delta t} \sum_{k=0}^n x_k e^{-i\omega \Delta t k}. \quad (5.43)$$

Taking the magnitude squared of y_n results in an estimate of the ordinary power spectrum, which is given by the double summation

$$S(\omega) \approx \frac{\Delta t}{n} |y_n|^2 = \frac{\Delta t}{n} \sum_{j=0}^n \sum_{k=0}^n x_j x_k^* e^{-i\omega \Delta t (j-k)}, \quad (5.44)$$

whereas to obtain the complex spectrum we need to insert a Heaviside step function into the double sum:

$$J(\omega) \approx \frac{\Delta t}{n} \sum_{j=0}^n \sum_{k=0}^n x_j x_k^* e^{-i\omega \Delta t (j-k)} \Theta(j-k), \quad (5.45)$$

where

$$\Theta(x) = \begin{cases} 1 & x > 0 \\ 1/2 & x = 0 \\ 0 & x < 0. \end{cases}$$

Fortunately, the double sum in equation 5.44 can be computed efficiently by maintaining the partial sum y_m in memory for each m . More concretely, we can re-write the double-sum in equation 5.45. Defining $a_j = x_j e^{-i\omega\Delta t j}$, we get

$$J(\omega) \approx \frac{\Delta t}{n} \sum_{j=0}^n a_j \sum_{k=0}^n a_k^* \theta(j-k) \quad (5.46)$$

$$\approx \frac{\Delta t}{n} \sum_{j=0}^n a_j \left(-a_j^*/2 + \sum_{k=0}^j a_k^* \right) \quad (5.47)$$

$$\approx \frac{\Delta t}{n} \sum_{j=0}^n a_j \left(-a_j^*/2 + e^{i\omega\Delta t j} y_j^* \right) \quad (5.48)$$

$$\approx \frac{\Delta t}{n} \sum_{j=0}^n x_j y_j^* - x_j^2/2. \quad (5.49)$$

From this we can define the partial sum of $J(\omega)$ as $J_m(\omega)$, whose recurrence is given by

$$J_m(\omega) = J_{m-1}(\omega) + y_m x_m - \frac{x_m^2}{2} \quad (5.50)$$

$$J(\omega) = \frac{\Delta t}{n} J_n^*(\omega). \quad (5.51)$$

Combining the recurrence relations of equations 5.41, 5.42, and 5.50, we arrive at an algorithm which computes the imaginary spectrum $J(\omega)$ which is linear in time and constant in memory usage. This process must be repeated for each frequency ω , and for each generated time series.

5.2.5.3 Fast Fourier Transform (FFT)

Alternatively, one can collect samples of the noise in an array, then take the FFT of the array. Algorithmically, this corresponds to

$$S(\omega) = |\text{FT}[\vec{x}]|^2, \quad (5.52)$$

where $\vec{x}$ is the vector of noise samples $x_0, x_1 \dots x_n$, and FT is the Fourier transform. Modern computer hardware allows us to improve the speed of this method by computing multiple Fourier transforms in parallel. To do this, we collect an array of vectors of noise samples, with each vector corresponding to a single particle in the ensemble. Applying equation 5.52 then produces one spectrum estimate for each particle. We can take the average of these spectra to obtain a more accurate estimate of the noise power spectrum $S(\omega)$. Once $S(\omega)$ has been calculated, we can find the

complex spectrum using the fact that the power spectrum is the Fourier transform of the correlation function. We take the inverse Fourier transform, multiply by the Heaviside step function in the time domain, then Fourier transform back to the frequency domain. In other words, we compute

$$R(\tau) = \text{FT}^{-1}[S(\omega)] \quad (5.53)$$

$$J(\omega) = \text{FT}[\Theta(\tau)R(\tau)]. \quad (5.54)$$

These last two steps are performed only once per ensemble, after averaging over the noise spectra of individual particles. One final complication with the FFT-based method is that generally, the frequencies output by the Fourier transform will not correspond exactly to those needed by the FMME. To remedy this, we use nearest-neighbor interpolation to approximate intermediate values of the spectrum.

5.2.6 Acceleration with Graphics Processing Units (GPUs)

The next generation of nEDM experiments seek to measure the nEDM with better than 1×10^{-27} e-cm precision; nEDM@SNS, for example, is projected to achieve an ultimate statistical uncertainty of 3×10^{-28} e-cm [Ahm+19]. This precision is achieved in part by storing a large number of UCN for a time comparable to the neutron beta decay lifetime of 877.8 s [Gon+21]. Simulating enough particles to constrain frequency shifts to this level of precision requires on the order of millions of particles, each one tracked for a thousand seconds to replicate the conditions of the experiment under study. In the field of UCN, a variety of open source code packages have been developed which can integrate the Bloch equations along a particle's trajectory: for example starUCN [BD16], MCUCN [Zsi17], PENTrack [Sch+17], Geant4 [Ago+03], and Kassiopeia [Fur+17]. While these software packages have a variety of advantages in terms of simulating arbitrary geometries, material interactions, or sophisticated wall reflection models, for the specific purpose of simulating large ensembles of particles, generating particle trajectories serially will be too time-consuming. To expedite this process, we leverage general-purpose graphical processing unit (GPGPU) programming. Each GPU thread will be responsible for tracking the trajectory of a single particle, sampling the magnetic field at even time intervals along its path. This allows many particles to be simulated in parallel, reducing the computation time by orders of magnitude at the cost of generality of the simulation software. For the FFT approach, using GPUs offers an additional advantage, as in this case we can use existing CUDA libraries to efficiently perform the requisite FFT and tensor reductions on NVIDIA GPUs. These libraries are cuFFT,

which performs FFT; and cuTENSOR, which performs general tensor contractions and reductions. For GPGPU programming, some additional optimizations may be necessary to extract maximum performance out of GPU hardware; specifically, we ensure that the tensor which contains the noise samples is arranged so that samples collected consecutively in time are adjacent in memory, as the cuFFT library is optimized for this data layout [NVI25]. Further optimizations were required for the GPU-based particle trajectory tracking scheme, which are detailed in [Mor23].

5.2.7 Comparison with Theory

Once the complex spectrum $J(\omega)$ has been estimated, extracting frequency shifts from the FMME-based algorithms is straightforward; the FMME prescribes the time evolution of the density matrix (see equations 5.84 and 5.100 for further details), which can be numerically integrated to determine the density matrix at the end of the simulation time. From there, the frequency shift can be computed as the change in phase divided by the simulation time. To confirm that the output of the FMME-based algorithms agree with the theoretical frequency shifts, we compare with the analytical results found in [Pig+15] in the adiabatic limit for free precession. We also benchmark these simulations against a standard Runge-Kutta algorithm, Dormand-Prince 853 (DOP) [WNH93]. We show comparisons for B^2 (figure 5.4), E^2 (figure 5.5), and linear-in-E (figure 5.6) frequency shifts. Table 5.1 in Appendix 5.2.11 lists the parameters used for these simulations. These parameters model the motion of helium-3 atoms in a superfluid helium-4 bath including phonon collisions, as would be relevant for an experiment such as nEDM@SNS. We simulate phonon collisions using the same model as in [SPG16], where the helium-3 velocity is resampled from a Maxwell-Boltzmann distribution after each collision. Velocity is otherwise constant between collisions. We find excellent agreement with the analytical results for the B^2 and E^2 shifts, while we see slight discrepancies for the linear-in-E shift for all three algorithms. We attribute this to the fact that the linear-in-E shift is much smaller in magnitude than the B^2 and E^2 shifts for these parameters, and any simulation of the linear-in-E shift necessarily also includes all three types of shifts. Thus Monte Carlo simulations of this effect tend to have greater uncertainties than those of the B^2 and E^2 shifts.

Likewise, in figure 5.7 we compare the output of the FMME-based algorithms with the results given by the theory in [TS22] for the spin dressing case. For this, we apply a 1 kHz sinusoidal magnetic field along the $\hat{x}$ direction with amplitude




Figure 5.4: Comparison of DOP, FMME Goertzel, and FMME FFT methods with theoretical predictions for the B^2 frequency shift in the adiabatic limit. The data points for the three integrators overlap.




Figure 5.5: Comparison of DOP, FMME Goertzel, and FMME FFT methods with theoretical predictions for the E^2 frequency shift in the adiabatic limit. The data points for the three integrators overlap.




Figure 5.6: Comparison of DOP, FMME Goertzel, and FMME FFT methods with theoretical predictions for the linear-in-E frequency shift in the adiabatic limit. For these simulations the electric field is varied, while the gradient is fixed at $\frac{dB_z}{dz} = 3 \times 10^{-10} \text{ T m}^{-1}$. A relatively small gradient was chosen to minimize the effect of the B^2 shift in this calculation.

$3.0 \times 10^{-5} \text{ T}$. As the authors of [TS22] consider only perturbations along the dressing field direction, only the B^2 shift can be treated theoretically.

5.2.7.1 Effect of Window Length

In cases where the characteristic correlation time of the system is sufficiently long, the window length must be increased to fully capture the correlation function. This is particularly important for estimating the B^2 shift, as in this case the relevant correlation time is the time taken for the particles to diffuse across the cell, which for the parameters used for these simulations is on the order of one second for helium-3 atoms (this time strongly depends on temperature). We illustrate this phenomenon in figures 5.8 and 5.9, where we can see that decreasing the window length has the effect of increasing the bias for the Floquet-based methods. The bias is suppressed as the window length approaches the typical correlation time of the system. Notably, the bias for the two methods at short window length differ significantly. We attribute this to the fact that the two algorithms handle a spectrum with low frequency resolution differently; the FFT method uses nearest-neighbor interpolation to fill in missing




Figure 5.7: Comparison of DOP, FMME Goertzel, and FMME FFT methods with theoretical predictions for the B^2 frequency shift in the adiabatic limit for spin dressing. The data points for the three integrators overlap.

frequencies, while the Goertzel method is more akin to taking the average of nearby frequencies.

5.2.8 Performance

Figure 5.10 compares the performance of the Goertzel algorithm and FFT approach with that of an adaptive Runge-Kutta integrator, DOP853. The parameters used for performance testing are identical to those used for the theory comparison, except that the dressing amplitude is 3.0×10^{-5} T, the dressing frequency is 1 kHz, and the window length is shortened to 0.1 seconds. Note that whereas the Runge-Kutta integrator outputs the entire ensemble of spins, the FMME-based algorithms output only the density matrix. For similar number of particles, the Goertzel algorithm-based method outperforms the Runge-Kutta based approach, while the FFT-based approach outperforms both. For the algorithms using the FMME, we found that the main performance bottleneck was the collection of noise samples. Because of this, the performance of the Goertzel algorithm method suffers somewhat as several additional calculations must be made at each time step (see equation 5.50). Even so, the number of calculations and function evaluations per step for both methods is still far fewer than the number required for an eighth-order integrator such as




Figure 5.8: The B^2 shift estimated by the Floquet Goertzel method varies with window length. The data for 0.1 s and 0.01 s overlap.




Figure 5.9: The B^2 shift estimated by the Floquet FFT method varies with window length.

DOP853. It is also noteworthy that while the Floquet Goertzel method is slower than the Floquet FFT method, it uses much less memory; while the FFT method must maintain an array of all noise samples collected in a time window, the Goertzel method maintains only the variables y_m and J_m from equation 5.50 for each particle and frequency. It is possible that for larger ensembles, the FFT method will need to integrate in batches.




Figure 5.10: Comparison of time taken to simulate various numbers of particles with DOP853, Floquet Goertzel method, or Floquet FFT method.

5.2.9 Conclusion

In this work, we developed a new method to integrate the Bloch equation for large ensembles of spin-1/2 particles in a time-periodic magnetic field subject to noise. Our approach is to use the Floquet-Markov Master Equation to divide the problem of integrating a large number of differential equations into two subproblems: first, to compute the Floquet states and quasi-energies of the (noise-free) periodic Hamiltonian; and second, to compute the complex power spectrum of the noise. The former is computationally trivial, while the second can be computed efficiently using fast Fourier transform algorithms on modern hardware. To this end, we also adapt the

Floquet-Markov master equation to the semi-classical case, in which the environment noise is described by a classical correlation function rather than as a quantum mechanical bath. While this algorithm is applicable in general to any quantum system where both a periodic driving field and random fluctuations are present, it is particularly well-suited to simulating the frequency shifts and relaxations which are relevant to precision measurements experiments such as nEDM@SNS. We compare the results of the method when computing these frequency shifts both to a standard Runge-Kutta integrator and theoretical results based on the work of [Pig+15], and find good agreement with both. We additionally compare the runtime of the various algorithms, and find an order-of-magnitude improvement over the standard Runge-Kutta algorithm when using the FFT-based Floquet method, while the runtime savings with the Goertzel algorithm-based Floquet method are more modest.

5.2.10 Acknowledgments

The author would like to thank Brad Filippone, Robert Golub, David Davis, and the nEDM@SNS publications committee for their helpful input and review of the paper. The author would also like to thank Matthew Morano and David Mathews for their work in designing the GPU-based particle tracking software on which this work is based. This work was funded by the National Science Foundation (NSF) grants 2110898 and 1822515. This research was supported in part through research cyber-infrastructure resources and services provided by the Partnership for an Advanced Computing Environment (PACE) at the Georgia Institute of Technology, Atlanta, Georgia, USA. This research was supported in part through research infrastructure and services provided by the Rogues Gallery testbed hosted by the Center for Research into Novel Computing Hierarchies (CRNCH) at Georgia Tech. The Rogues Gallery testbed is primarily supported by the National Science Foundation (NSF) under NSF Award Number #2016701. Any opinions, findings and conclusions, or recommendations expressed in this material are those of the author, and do not necessarily reflect those of the NSF.

5.2.11 Simulation Parameters

Table 5.1 lists the parameters used for comparisons with the theoretical frequency shifts predicted by [Pig+15].

Parameter	Value	Unit
B0	5.0×10^{-6}	T
Dressing Amplitude	0	T
Dressing Frequency	N/A	Hz
Cell Dimensions	$1.0 \times 1.0 \times 1.0$	m
Particle Mass	1.20239×10^{-26}	kg
Simulation Time	1.0	s
Window Length	1.0	s
Gyromagnetic Ratio	-2.038×10^8	$\text{s}^{-1} \text{T}^{-1}$
Sampling Rate	10^4	Hz
Temperature	0.25	K
Wall diffusivity	0.0	
Gravity	0.0	m s^{-2}

Table 5.1: Parameters used for comparisons with [Pig+15].

5.2.12 Floquet-Markov Master Equation Derivation

In this section, we derive the semi-classical Floquet-Markov master equation including energy shifts. This derivation largely follows that in [Blü+91].

5.2.12.1 Von Neumann-Liouville Equation in the Floquet Basis

Consider a Hamiltonian

$$H(t) = H_0(t) + H_{\text{int}}(t), \quad (5.55)$$

where $H_0(t)$ is periodic and deterministic with period $T = 2\pi/\omega$, and $H_{\text{int}}(t)$ is a stochastic component given by

$$H_{\text{int}}(t) = f(t)\sigma \quad (5.56)$$

for some (quantum-mechanical) operator σ and some wide-sense stationary (WSS) scalar stochastic process of zero mean $f(t)$ whose autocorrelation is known.

We begin with the von Neumann-Liouville equation in the interaction picture:

$$i\hbar \frac{d\tilde{\rho}}{dt} = [\tilde{H}_{\text{int}}(t), \tilde{\rho}] \quad (5.57)$$

where $\tilde{\rho}$ and $\tilde{H}_{\text{int}}$ are the density matrix and Hamiltonian, each evaluated in the interaction picture:

$$\tilde{\rho}(t) = U^\dagger(t)\rho(t)U(t) \quad (5.58)$$

$$\tilde{H}_{\text{int}}(t) = U^\dagger(t)H_{\text{int}}(t)U(t). \quad (5.59)$$

The transformation to the interaction picture is given in the Floquet basis:

$$U(t) = \sum_{\alpha} e^{-i\mu_{\alpha}t} |\Phi_{\alpha}(t)\rangle \langle \Phi_{\alpha}(0)| \quad (5.60)$$

where the $|\Phi_{\alpha}(t)\rangle$ are the Floquet modes and the $\mu_{\alpha} = \epsilon_{\alpha}/\hbar$ are the frequencies associated with the corresponding quasi-energies. We can expand the interaction Hamiltonian in the Floquet basis, yielding

$$\tilde{H}_{\text{int}}(t) = U^{\dagger}(t) f(t) \sigma U(t) \quad (5.61)$$

$$= f(t) \sigma_{\text{int}}(t) \quad (5.62)$$

$$= f(t) \sum_{\alpha, \beta} e^{i(\mu_{\alpha} - \mu_{\beta})t} |\Phi_{\alpha}(0)\rangle \langle \Phi_{\alpha}(t)| \sigma |\Phi_{\beta}(t)\rangle \langle \Phi_{\beta}(0)|, \quad (5.63)$$

where $\sigma_{\text{int}}(t) = U^{\dagger}(t) \sigma U(t)$ is σ in the interaction picture.

Expanding the von Neumann-Liouville equation, we get:

$$\frac{d\tilde{\rho}}{dt} = \frac{1}{i\hbar} [\tilde{H}_{\text{int}}(t), \tilde{\rho}(t)] \quad (5.64)$$

$$= \frac{1}{i\hbar} \left[\tilde{H}_{\text{int}}(t), \tilde{\rho}(0) + \frac{1}{i\hbar} \int_0^t dt' [\tilde{H}_{\text{int}}(t'), \tilde{\rho}(t')] \right] \quad (5.65)$$

$$= \frac{1}{i\hbar} [\tilde{H}_{\text{int}}(t), \tilde{\rho}(0)] + \frac{1}{(i\hbar)^2} \int_0^t dt' [\tilde{H}_{\text{int}}(t), [\tilde{H}_{\text{int}}(t'), \tilde{\rho}(t')]] . \quad (5.66)$$

Substituting the expression in 5.62, this becomes

$$\frac{d\tilde{\rho}}{dt} = \frac{1}{i\hbar} f(t) [\sigma_{\text{int}}(t), \tilde{\rho}(0)] + \frac{1}{(i\hbar)^2} \int_0^t dt' f(t) f(t') [\sigma_{\text{int}}(t), [\sigma_{\text{int}}(t'), \tilde{\rho}(t')]] . \quad (5.67)$$

$$= \frac{1}{(i\hbar)^2} \int_0^t dt' R(t-t') [\sigma_{\text{int}}(t), [\sigma_{\text{int}}(t'), \tilde{\rho}(t')]] , \quad (5.68)$$

where in the second line we take the average over the functions $f(t)$ ($\langle f(t) \rangle = 0$, $\langle f(t)f(t') \rangle = R(t-t')$) and apply the Markov approximation, replacing $\tilde{\rho}(t')$ with $\tilde{\rho}(t)$. Next we introduce the Fourier transform $X_{\alpha\beta k}$ of the matrix element $\langle \Phi_{\alpha}(t) | \sigma | \Phi_{\beta}(t) \rangle$, defined by

$$\langle \Phi_{\alpha}(t) | \sigma | \Phi_{\beta}(t) \rangle = \sum_{k=-\infty}^{\infty} e^{ik\omega t} X_{\alpha\beta k} . \quad (5.69)$$

We can now write $\sigma_{\text{int}}(t)$ as

$$\sigma_{\text{int}}(t) = \sum_{\alpha, \beta, k} e^{i(\mu_{\alpha} - \mu_{\beta} + k\omega)t} |\Phi_{\alpha}(0)\rangle \langle \Phi_{\beta}(0)| X_{\alpha\beta k} \quad (5.70)$$

$$= \sum_{\alpha, \beta, k} e^{i\Delta_{\alpha\beta k}t} |\Phi_{\alpha}(0)\rangle \langle \Phi_{\beta}(0)| X_{\alpha\beta k}, \quad (5.71)$$

where we have defined $\Delta_{\alpha\beta k} = \mu_\alpha - \mu_\beta + k\omega$. Substituting into the von Neumann-Liouville equation we obtain

$$\begin{aligned} \frac{d\tilde{\rho}}{dt} = \frac{1}{(i\hbar)^2} \int_0^t dt' R(t-t') & \left[\sum_{\alpha,\beta,k} e^{i\Delta_{\alpha\beta k}t} |\Phi_\alpha(0)\rangle \langle \Phi_\beta(0)| X_{\alpha\beta k}, \right. \\ & \left. \left[\sum_{\alpha',\beta',k'} e^{i\Delta_{\alpha'\beta'k'}t'} |\Phi_{\alpha'}(0)\rangle \langle \Phi_{\beta'}(0)| X_{\alpha'\beta'k'}, \tilde{\rho}(t) \right] \right]. \quad (5.72) \end{aligned}$$

Rearranging terms, we get

$$\begin{aligned} \frac{d\tilde{\rho}}{dt} = \frac{1}{(i\hbar)^2} \int_0^t dt' R(t-t') \sum_{\alpha,\beta,k} \sum_{\alpha',\beta',k'} e^{i\Delta_{\alpha\beta k}t+i\Delta_{\alpha'\beta'k'}t'} X_{\alpha\beta k} X_{\alpha'\beta'k'} \\ [|\Phi_\alpha(0)\rangle \langle \Phi_\beta(0)|, [|\Phi_{\alpha'}(0)\rangle \langle \Phi_{\beta'}(0)|, \tilde{\rho}(t)]]]. \quad (5.73) \end{aligned}$$

We now make use of some useful symmetries of Δ and X :

$$\Delta_{\alpha\beta k} = -\Delta_{\beta\alpha -k} \quad (5.74)$$

$$X_{\alpha\beta k} = X_{\beta\alpha -k}^*. \quad (5.75)$$

Further note that since k' ranges from $-\infty$ to $+\infty$, we are free to replace k' with $-k'$ in the sum. And since α' and β' each index over the same set (the Floquet states), we can freely exchange α' and β' . Applying all of these changes, we get:

$$\frac{d\tilde{\rho}}{dt} = \frac{1}{(i\hbar)^2} \int_0^t dt' R(t-t') \sum_{\alpha,\beta,k} \sum_{\alpha',\beta',k'} e^{i\Delta_{\alpha\beta k}t-i\Delta_{\alpha'\beta'k'}t'} X_{\alpha\beta k} X_{\alpha'\beta'k'}^* \quad (5.76)$$

$$\begin{aligned} & [|\Phi_\alpha(0)\rangle \langle \Phi_\beta(0)|, [|\Phi_{\beta'}(0)\rangle \langle \Phi_{\alpha'}(0)|, \tilde{\rho}(t)]] \\ & = \frac{1}{(i\hbar)^2} \sum_{\alpha,\beta,k} \sum_{\alpha',\beta',k'} e^{i(\Delta_{\alpha,\beta,k}-\Delta_{\alpha'\beta'k'})t} \left(\int_0^t dt' R(t-t') e^{i\Delta_{\alpha'\beta'k'}(t-t')} \right) \\ & X_{\alpha\beta k} X_{\alpha'\beta'k'}^* [|\Phi_\alpha(0)\rangle \langle \Phi_\beta(0)|, [|\Phi_{\beta'}(0)\rangle \langle \Phi_{\alpha'}(0)|, \tilde{\rho}(t)]] \end{aligned} \quad (5.77)$$

We further apply the Markovian approximation, assuming that correlation times are short compared to the measurement time. For particles confined to a container, the typical timescales involved are the average time between collisions and the diffusion time across a container. For example for the nominal parameters of the nEDM experiment in [Ahm+19], the average collision time is on the order of milliseconds, while the diffusion time is on the order of a second — both of which are much

shorter than the proposed measurement time of 1000 s. In such cases $R(\tau)$ decays quickly enough that we can replace the lower limit of integration with $-\infty$. Thus,

$$\int_0^t dt' R(t-t') e^{i\Delta_{\alpha'\beta'k'}(t-t')} \approx \int_{-\infty}^t dt' R(t-t') e^{i\Delta_{\alpha'\beta'k'}(t-t')} \quad (5.78)$$

$$= - \int_{-\infty}^0 d\tau R(\tau) e^{i\Delta_{\alpha'\beta'k'}\tau} \quad (5.79)$$

$$= \int_0^{\infty} d\tau R(\tau) e^{i\Delta_{\alpha'\beta'k'}\tau} \quad (5.80)$$

$$= J(\Delta_{\alpha'\beta'k'}) \quad (5.81)$$

where we have defined $J(\omega) \equiv \int_0^{\infty} d\tau R(\tau) e^{i\omega\tau}$ as the one-sided Fourier transform of the autocorrelation function. At this point, we can observe that $J(\omega)$ is related to what we refer to simply as the spectrum, defined as $S(\omega) \equiv \int_{-\infty}^{\infty} d\tau R(\tau) e^{i\omega\tau}$. Notice the integral goes from $-\infty$ to $+\infty$. Let us note some properties of $J(\omega)$ and $S(\omega)$:

- $S(\omega)$ is real and even (because it is the Fourier transform of $R(\tau)$, which is real and even).
- $J(\omega) = J^*(-\omega)$.
- $S(\omega) = J(\omega) + J^*(\omega) = 2 \operatorname{Re}[J(\omega)]$.
- $J(\omega) = \frac{1}{2}S(\omega) - \frac{i}{2\pi} \mathcal{P} \int_{-\infty}^{\infty} \frac{S(\omega')}{\omega' - \omega} d\omega'$ ($\mathcal{P}$ denotes the Cauchy principal value).

The first three identities can be seen directly from the definitions of $J(\omega)$ and $S(\omega)$. The final identity can be derived either from the Kramers-Kronig relations, or by using the Fourier transform of the Heaviside step function (i.e., $J(\omega) = \int_{-\infty}^{\infty} d\tau \Theta(\tau) R(\tau) e^{i\omega\tau}$) along with the convolution theorem.

To simplify notation, we also introduce the operators

$$\Gamma_{\alpha\beta} \equiv |\Phi_{\alpha}(0)\rangle \langle \Phi_{\beta}(0)|. \quad (5.82)$$

With this, our master equation has now simplified to

$$\frac{d\tilde{\rho}}{dt} = \frac{1}{(i\hbar)^2} \sum_{\alpha,\beta,k} \sum_{\alpha',\beta',k'} J(\Delta_{\alpha'\beta'k'}) e^{i(\Delta_{\alpha,\beta,k} - \Delta_{\alpha'\beta'k'})t} X_{\alpha\beta k} X_{\alpha'\beta'k'}^* [\Gamma_{\alpha\beta}, [\Gamma_{\beta'\alpha'}, \tilde{\rho}(t)]] . \quad (5.83)$$

Next, we need to apply a rotating wave approximation. The argument is that the phase factor $e^{i(\Delta_{\alpha,\beta,k}-\Delta_{\alpha'\beta'k'})t}$ oscillates quickly, and thus can be neglected unless $\Delta_{\alpha,\beta,k} = \Delta_{\alpha'\beta'k'}$. From the definition of $\Delta_{\alpha\beta k} = \mu_\alpha - \mu_\beta + k\omega$ we can see that there are two cases where these two frequencies are equal:

1. $(\alpha, \beta, k) = (\alpha', \beta', k')$
2. $\alpha = \beta, \alpha' = \beta', k = k'$.

Some quantum systems may have additional degeneracies, but we will assume these are the only ones. We tackle case 1 first. Assuming $(\alpha, \beta, k) = (\alpha', \beta', k')$, the master equation further simplifies to:

$$\frac{d\tilde{\rho}}{dt} = \frac{1}{(i\hbar)^2} \sum_{\alpha,\beta,k} J(\Delta_{\alpha\beta k}) |X_{\alpha\beta k}|^2 [\Gamma_{\alpha\beta}, [\Gamma_{\beta\alpha}, \tilde{\rho}(t)]] \quad (5.84)$$

$$= - \sum_{\alpha,\beta} \zeta_{\alpha\beta} [\Gamma_{\alpha\beta}, [\Gamma_{\beta\alpha}, \tilde{\rho}(t)]] \quad (5.85)$$

$$= - \sum_{\alpha,\beta} \zeta_{\alpha\beta} (\Gamma_{\alpha\alpha} \tilde{\rho}(t) + \tilde{\rho}(t) \Gamma_{\beta\beta} - \Gamma_{\alpha\beta} \tilde{\rho}(t) \Gamma_{\beta\alpha} - \Gamma_{\beta\alpha} \tilde{\rho}(t) \Gamma_{\alpha\beta}) \quad (5.86)$$

where in the second line we perform the sum over k , and define

$$\zeta_{\alpha\beta} \equiv \frac{1}{\hbar^2} \sum_k J(\Delta_{\alpha\beta k}) |X_{\alpha\beta k}|^2. \quad (5.87)$$

Examining $\zeta_{\alpha\beta}$, we see that $\zeta_{\alpha\beta} = \zeta_{\beta\alpha}^*$, i.e., ζ is Hermitian. This matrix is related to the matrix $M_{\alpha\beta}$ given in [Blü+91]. For classical noise, we should take $n_{th}(\omega) = \frac{1}{2\pi\hbar^2} \gg 1$. We obtain:

$$M_{\alpha\beta} = \frac{1}{\hbar^2} \sum_{\alpha,\beta,k} \Theta(\Delta_{\alpha\beta k}) S(\Delta_{\alpha\beta k}) |X_{\alpha\beta k}|^2 + \Theta(\Delta_{\beta\alpha k}) S(\Delta_{\beta\alpha k}) |X_{\beta\alpha k}|^2 \quad (5.88)$$

$$= \frac{1}{\hbar^2} \sum_{\alpha,\beta,k} \Theta(\Delta_{\alpha\beta k}) S(\Delta_{\alpha\beta k}) |X_{\alpha\beta k}|^2 + \Theta(-\Delta_{\alpha\beta k}) S(\Delta_{\alpha\beta k}) |X_{\alpha\beta k}|^2 \quad (5.89)$$

$$= \frac{1}{\hbar^2} \sum_{\alpha,\beta,k} S(\Delta_{\alpha\beta k}) |X_{\alpha\beta k}|^2 \quad (5.90)$$

$$= \zeta_{\alpha,\beta} + \zeta_{\alpha,\beta}^* \quad (5.91)$$

$$= \zeta_{\alpha,\beta} + \zeta_{\beta,\alpha}. \quad (5.92)$$

Notably, if we assume $J(\omega)$ is purely real, which amounts to discarding the principal part of the integral as they do in [Blü+91], then $M_{\alpha\beta}$ and $\zeta_{\alpha\beta}$ are proportional. We now examine the matrix elements of $\tilde{\rho}$, which we will define as $\rho_{\nu\sigma} = \langle \Phi_\nu(0) | \tilde{\rho} | \Phi_\sigma(0) \rangle$. Taking the matrix elements of both sides, we get

$$\frac{d\rho_{\nu\sigma}}{dt} = - \sum_{\alpha,\beta} \zeta_{\alpha\beta} (\delta_{\nu\alpha} \rho_{\alpha\sigma} + \rho_{\nu\beta} \delta_{\beta\sigma} - \delta_{\nu\alpha} \delta_{\alpha\sigma} \rho_{\beta\beta} - \delta_{\nu\beta} \delta_{\beta\sigma} \rho_{\alpha\alpha}) \quad (5.93)$$

$$= - \sum_{\beta} \zeta_{\nu\beta} \rho_{\nu\sigma} - \sum_{\alpha} \zeta_{\alpha\sigma} \rho_{\nu\sigma} + \delta_{\nu\sigma} \sum_{\beta} \zeta_{\nu\beta} \rho_{\beta\beta} + \delta_{\nu\sigma} \sum_{\alpha} \zeta_{\alpha\sigma} \rho_{\alpha\alpha} \quad (5.94)$$

$$= - \left(\sum_{\alpha} \zeta_{\nu\alpha} + \zeta_{\alpha\sigma} \right) \rho_{\nu\sigma} + \delta_{\nu\sigma} \left(\sum_{\alpha} (\zeta_{\nu\alpha} + \zeta_{\alpha\sigma}) \rho_{\alpha\alpha} \right). \quad (5.95)$$

We can put this expression into perhaps a more intuitive form if we split this equation into the diagonal and off-diagonal cases: $\nu = \sigma$ and $\nu \neq \sigma$.

$$\frac{d\rho_{\nu\nu}}{dt} = \sum_{\alpha} [\zeta + \zeta^T]_{\nu\alpha} \rho_{\alpha\alpha} - \sum_{\alpha} [\zeta + \zeta^T]_{\nu\alpha} \rho_{\nu\nu} \quad (5.96)$$

$$\frac{d\rho_{\nu\sigma}}{dt} = - \left(\sum_{\alpha} \zeta_{\nu\alpha} + \zeta_{\alpha\sigma} \right) \rho_{\nu\sigma}, \quad \nu \neq \sigma \quad (5.97)$$

Finally, to match more closely with the notation in [Blü+91], we switch the labeling of the indices.

$$\frac{d\rho_{\alpha\alpha}}{dt} = \sum_{\nu} [\zeta + \zeta^T]_{\alpha\nu} \rho_{\nu\nu} - \sum_{\nu} [\zeta + \zeta^T]_{\alpha\nu} \rho_{\alpha\alpha} \quad (5.98)$$

$$\frac{d\rho_{\alpha\beta}}{dt} = - \left(\sum_{\nu} \zeta_{\alpha\nu} + \zeta_{\nu\beta} \right) \rho_{\alpha\beta}, \quad \alpha \neq \beta \quad (5.99)$$

In the case where $\alpha = \beta, \alpha' = \beta', k = k', \alpha \neq \alpha'$, we have

$$\frac{d\tilde{\rho}}{dt} = \frac{1}{(i\hbar)^2} \sum_{\alpha \neq \alpha', k} J(\Delta_{\alpha'\alpha'k}) X_{\alpha\alpha k} X_{\alpha'\alpha'k}^* [\Gamma_{\alpha\alpha}, [\Gamma_{\alpha'\alpha'}, \tilde{\rho}(t)]] . \quad (5.100)$$

It is worth noting that $\Delta_{\alpha\alpha k} = k\omega$ regardless of α , and that $X_{\alpha\alpha k} = X_{\alpha\alpha-k}^*$. Let us now define the matrix

$$\xi_{\alpha\alpha'} = \frac{1}{\hbar^2} \sum_k J(\Delta_{\alpha'\alpha'k}) X_{\alpha\alpha k} X_{\alpha'\alpha'k}^* \quad \alpha \neq \alpha' \quad (5.101)$$

$$= \frac{1}{\hbar^2} \sum_k J(k\omega) X_{\alpha\alpha k} X_{\alpha'\alpha'k}^* \quad \alpha \neq \alpha'. \quad (5.102)$$

Also, since we are summing only over terms with $\alpha \neq \alpha'$, we have that $\Gamma_{\alpha\alpha}\Gamma_{\alpha'\alpha'} = \Gamma_{\alpha'\alpha'}\Gamma_{\alpha\alpha} = 0$. Therefore we get

$$\frac{d\tilde{\rho}}{dt} = \sum_{\alpha \neq \alpha'} \xi_{\alpha\alpha'} (\Gamma_{\alpha\alpha}\tilde{\rho}\Gamma_{\alpha'\alpha'} + \Gamma_{\alpha'\alpha'}\tilde{\rho}\Gamma_{\alpha\alpha}) \quad (5.103)$$

$$\frac{d\rho_{\alpha\beta}}{dt} = \xi_{\alpha\beta}\rho_{\alpha\beta} + \xi_{\beta\alpha}\rho_{\beta\alpha} \quad \alpha \neq \beta. \quad (5.104)$$

5.2.12.2 Fields in Multiple Directions

Up until now we have considered only perturbations which apply to a single operator. What if there are multiple operators involved (e.g. in the case of spin, multiple magnetic fields)? Then there may be correlations between the perturbations for the various operators.

Suppose now that our perturbation Hamiltonian is given by

$$H_{\text{int}}(t) = \sum_i f^{(i)}(t)\sigma^{(i)}, \quad (5.105)$$

where $\sigma^{(i)}$ are operators, and the correlations between the $f^{(i)}(t)$ are given by a covariance matrix $R_{ij}(\tau)$

$$R_{ij}(\tau) \equiv \left\langle f^{(i)}(t)f^{(j)}(t-\tau) \right\rangle. \quad (5.106)$$

The perturbation Hamiltonian in the Floquet basis is then

$$\tilde{H}_{\text{int}} = \sum_{\alpha,\beta,k,i} e^{i\Delta_{\alpha\beta k}t} |\Phi_{\alpha}(0)\rangle \langle \Phi_{\beta}(0)| f^{(i)}(t) X_{\alpha\beta k}^{(i)}, \quad (5.107)$$

where

$$\langle \Phi_{\alpha}(t) | \sigma^{(i)} | \Phi_{\beta}(t) \rangle = \sum_{k=-\infty}^{\infty} e^{ik\omega t} X_{\alpha\beta k}^{(i)}. \quad (5.108)$$

Following the derivation for the single field case, the master equation can now be written as

$$\frac{d\tilde{\rho}}{dt} = \frac{1}{(i\hbar)^2} \sum_{\alpha,\beta,k,i} \sum_{\alpha',\beta',k',j} J_{ij}(\Delta_{\alpha'\beta'k'}) e^{i(\Delta_{\alpha\beta k} - \Delta_{\alpha'\beta'k'})t} X_{\alpha\beta k}^{(i)} X_{\alpha'\beta'k'}^{(j)*} [\Gamma_{\alpha\beta}, [\Gamma_{\beta'\alpha'}, \tilde{\rho}(t)]] , \quad (5.109)$$

where

$$J_{ij}(\omega) \equiv \int_0^{\infty} R_{ij}(\tau) e^{i\omega\tau} d\tau. \quad (5.110)$$

Once again, we specialize to two cases:

If $(\alpha, \beta, k) = (\alpha', \beta', k')$, then we can define the matrix

$$\zeta_{\alpha\beta} = \frac{1}{\hbar^2} \sum_{k,i,j} J_{ij}(\Delta_{\alpha\beta k}) X_{\alpha\beta k}^{(i)} X_{\alpha\beta k}^{(j)*} \quad (5.111)$$

and the rest of the derivation proceeds as in Section 5.2.12.1.

If $\alpha = \beta, \alpha' = \beta', k = k', \alpha \neq \alpha'$, then we can define the matrix

$$\xi_{\alpha\alpha'} = \frac{1}{\hbar^2} \sum_{k,i,j} J_{i,j}(k\omega) X_{\alpha\alpha k}^{(i)} X_{\alpha'\alpha' k}^{(j)*} \quad (5.112)$$

and the rest of the derivation proceeds as in Section 5.2.12.1.

5.2.13 Cross-Spectrum Properties

$$J_{xy}(\omega) \equiv \int_0^\infty \langle x(t)y(t+\tau) \rangle e^{-i\omega\tau} d\tau \quad (5.113)$$

$$J_{xy}(-\omega) = \int_0^\infty \langle x(t)y(t+\tau) \rangle e^{+i\omega\tau} d\tau = J_{xy}(\omega)^* \quad (5.114)$$

$$J_{xy}(\omega) + J_{yx}^*(\omega) = \int_0^\infty \langle x(t)y(t+\tau) \rangle e^{-i\omega\tau} + \langle y(t)x(t+\tau) \rangle e^{+i\omega\tau} d\tau \quad (5.115)$$

$$= \int_0^\infty \langle x(t)y(t+\tau) \rangle e^{-i\omega\tau} d\tau + \int_{-\infty}^0 \langle y(t)x(t-\tau) \rangle e^{-i\omega\tau} d\tau \quad (5.116)$$

$$= \int_{-\infty}^\infty \langle x(t)y(t+\tau) \rangle e^{-i\omega\tau} d\tau \quad (5.117)$$

$$J_{xy}(\omega) + J_{yx}^*(\omega) = S_{xy}(\omega). \quad (5.118)$$

Chapter 6

SEOP STUDIES OF SPIN DRESSING TECHNIQUES AT DUKE UNIVERSITY

In light of spin dressing's role in improving nEDM statistical sensitivity [Ahm+19], and calculations (see Chapter 4) showing that critical spin dressing can be modified to render the spin ensemble insensitive to fluctuations and inhomogeneities in dressing field amplitude, it is pertinent to conduct experiments to verify the utility of these new dressing techniques. Here we study the magic dressing concept outlined in Chapter 4. Magic dressing insulates the spins from the effects of magnetic field fluctuations or inhomogeneities, though in this case we restrict our attention to inhomogeneities in the dressing field. Normally, such gradients would result in a shorter transverse relaxation time T_2 , i.e. the spin polarization would decay while the dressing field is applied. In this context, magic dressing is achieved when the amplitude of the dressing field is tuned to a value where polarization can be maintained, even when gradients in the dressing field are large. To demonstrate magic dressing experimentally, we use the spin exchange optical pumping (SEOP) [WH97] apparatus at Duke to produce a sample of polarized ^3He . This chapter is split into three parts: first, we describe the theory of SEOP; second, we detail the measurement scheme and equipment used to perform magic dressing; and lastly, we analyze the effect of the magic dressing on spins as an artificially added dressing field gradient is applied.

6.1 Spin Exchange Optical Pumping (SEOP) Theory

SEOP is a technique for producing hyperpolarized spin samples, in this case ^3He . Because ^3He itself in its electronic ground state is not directly amenable to optical pumping, we instead polarize a species which is. Then the polarization can be transferred from the electrons of the optically pumped species onto the ^3He nucleus via collisions in the gas. For this application, we optically pump rubidium's D1 line ($5^2S_{1/2} \rightarrow 5^2P_{1/2}$) with circularly polarized 795 nm light. This process is shown schematically in figure 6.1. The relevant parts of the Hamiltonian for an alkali atom such as Rubidium will take the form

$$H = \alpha I \cdot S + g\mu_S S_z B_0 + \mu I_z B_z, \quad (6.1)$$

where the first term is the hyperfine interaction, the second term is the Zeeman splitting, and the last is the interaction with the nuclear spin. Generally speaking, the hyperfine interaction ($\nu \sim O(\text{GHz})$) will dominate over the Zeeman splitting in a weak field ($\nu \sim O(10\text{kHz})$). However, due to motional broadening, the hyperfine structure is not resolved [WH97].

With the rubidium optically pumped, its polarization is transferred to the nuclear spin of ^3He by collisions through a hyperfine-like interaction. This process is quite slow, taking 12 or more hours for maximum polarization. The laser and oven must stay on during this time, as the rubidium polarization must be replenished to continue polarizing the ^3He . Once polarized, however, the spin polarization of ^3He persists for quite some time, provided that a holding field of adequate strength and uniformity is maintained. A variety of sources contribute to relaxation of the accumulated ^3He polarization, including helium-helium collisions, wall collisions, and magnetic field gradients. Despite this, the longitudinal (T_1) relaxation time for our cell was estimated to be ~ 12 hours. While theoretically ^3He and rubidium are the only elements needed for SEOP, practically the performance of a cell is significantly improved by the inclusion of two additional elements, which will be described in the next sections.

6.1.1 Buffer Gas

Without a buffer gas, radiative decay of the rubidium in the $^2P_{1/2}$ state would produce unpolarized photons, resulting in radiation trapping which would cause relaxation of the excited state [WH97]. To mitigate this effect, a small amount of nitrogen (~ 0.1 bar) [Dip19] is added to the SEOP cell, allowing the decay from $^2P_{1/2}$ to occur without emission of a photon. Collisions with the buffer gas also mix the Zeeman sublevels of the P states, a process known as *collisional mixing*.

6.1.2 Hybrid SEOP

In the cells that we used, the spin-exchange efficiency is further improved through use of *hybrid* SEOP, wherein potassium is added to the cell. The spin-exchange efficiency of rubidium with ^3He is roughly 2%, whereas the efficiency between potassium and ^3He is 23% [Bab05]. Thus by introducing a sufficiently high pressure of potassium into the SEOP cell, the spin polarization will be first transferred to the potassium, and then transferred efficiently to the ^3He nuclear spins. In fact the spin-exchange cross section between K and Rb is so large — around $200\text{\AA}^2$ [Bab05] — that one can often assume that K and Rb are in spin equilibrium.





Figure 6.1: Level diagram of rubidium valence electron, used for optical pumping.



(a) The SEOP cell. Rubidium metal can be seen in the tip.



(b) Box containing the SEOP cell, with manufacturing information.

The SEOP cell used for these studies is shown in figure 6.2a, along with the box with its manufacturing information in figure 6.2b. The bulb of the cell measures approximately 5 cm in diameter. Note that there are two cells in the “KellyK” series, the first of which has a fill density of 1 amagat. We use only KellyK#2 here.

6.2 Description of SEOP Apparatus

We now describe the apparatus, and the various subsystems thereof, used to produce polarized ^3He at Duke. These subsystems include the laser optics, the hot air system, the B_0 system, the dressing field (B_1) system, the NMR electronics, and the SEOP cell itself.




Figure 6.2: Optics elements to produce two beams of circularly polarized light to optically pump the rubidium vapor. The orientation of the three quarter-wave plates was tuned by mounting a linear polarizing filter in the optics box in conjunction with a laser power meter downstream.

6.2.1 Laser Optics

The optics used to produce circularly polarized light 795 nm laser light is shown in figure 6.2 and is represented schematically in figure 6.3. The 60 W beam first hits a beam diffuser (which ensures that the beam is spread across the entire cell). The polarizing beamsplitter cube separates the two orthogonal components of polarization. The vertically polarized light passes through a quarter wave plate and is subsequently circularly polarized. The horizontally polarized beam first passes through a quarter wave plate twice (thus making it vertically polarized), and then passes through a quarter wave plate to make it circularly polarized with the same handedness as the other beam.

6.2.2 Heating System

For SEOP, a sufficient partial pressure of rubidium vapor must be maintained. The house air system provides continuous air flow, which we regulate to roughly 10 psi. The system has an interlock, which shuts off the air flow should the pressure drop below a predefined set point, around 5 psi. The air flow then passes through three OMEGA-brand AHP-7561 heaters which are mounted in series. Power to these three heaters is controlled by three heater controllers, which in conjunction with an RTD thermometer affixed to the inside of the oven, maintain an oven temperature of 195 degrees Celsius. These heaters had to be replaced, as it was discovered that




Figure 6.3: Schematic of optics used for the SEOP system. Having two beam paths that converge at the cell allows us to utilize all 60 W of laser power. It is important to ensure that both beams have the same handedness.

the system design had a flaw that caused the epoxy plug on these heaters to degrade over time. Specifically, the inlet of these heaters are only rated for temperatures of up to 121 degrees Celsius. Running the three heaters in series causes the epoxy plug on the downstream heaters to deteriorate, eventually leading to air leaking out of the heater tube. The new heaters are the same make and model as the old heaters, but following this incident, the heating system was modified in two ways: first, the oven was wrapped in insulating fabric, which reduced the heater power requirement; second, instead of powering all three heaters, we run only two. This is only possible because of the reduced heating requirement. Critically, in this configuration the downstream heater (the one closest to the oven) must be powered, as otherwise the inlet to this heater will be the hottest point in the system, and we run into the same problem as before.

6.2.3 Magnetic Field System

Figure 6.4 shows a side view of the magnetic field system used for SEOP. The large red coils provide a holding field parallel to the laser beam direction which produces the Zeeman splitting in Rubidium. These coils are maintained at a constant current of 3 A, corresponding to a field strength of approximately 6.3 G. Note that these coils also serve as the B_0 field for subsequent NMR experiments, for which the current is reduced to 0.91 A so that the ^3He Larmor frequency matches the pickup coil resonance frequency. This corresponds to a field strength of 1.9 G. The circular black coils shown in figure 6.4 are the shim coils, whose purpose is to cancel any




Figure 6.4: Side view of magnetic field components used for in-situ NMR with the SEOP system, with the oven shown in the center. The outer red coils are the B_0 , the circular black coils are shim coils, and the coils around the oven are the dressing field coils.

DC field component parallel to the dressing field (more on this in section 6.6.1). These coils are not active during SEOP, but are set to roughly 15 mA during spin dressing experiments.

6.2.3.1 Dressing Coil Design

For previous tests with this SEOP system, an ex situ NMR system was used to examine the ^3He polarization [Dip19]. After the cell was polarized, the oven would be shut off, the cell allowed to cool, after which the cell could be transported. In order to simplify this process, as well as to reduce the magnetic field gradients caused by the small RF coil used in the ex situ system, a new set of gradient and pickup coils were designed. Following the technique used by [Mag+10], a design consisting of a set of four square coils was chosen for the dressing coils. Enforcing symmetry on the model, the parameters varied were the side lengths of the coils, the ratio of current between the inner and outer coils, and the separation between the coils. With these parameters, we can numerically solve for the configuration that

Parameter	Value
Outer coil side length	14.26 cm
Outer coil separation	18.25 cm
Outer coil turns	19
Inner coil side length	21.29 cm
Inner coil separation	6.67 cm
Inner coil turns	27

Table 6.1: Optimized parameters used for dressing coil design. The actual design parameters will differ from these slightly owing to the thickness of the coil wires.




Figure 6.5: One octant of the SEOP dressing coil, simulated in COMSOL with arrows showing the magnetic field direction and magnitude. The sphere octant in the lower right corner represents the SEOP cell.

will set the 2nd, 4th, and 6th order gradient of the magnetic field at the center of the coils to be zero (the odd order gradients are zero by symmetry). The optimized coil parameters are shown in table 6.1. The coils were simulated in COMSOL to confirm that the gradients were within an acceptable range. The model is shown in figure 6.5, and the variation of the magnetic field along the primary field direction over the cell region was found to be less than 0.13%.

6.2.3.2 Dressing Coil Construction

As the spin dressing coil frame is exposed to AC magnetic fields, it must be non-conductive. The initial prototype for these coils, 3D-printed from polylactic acid (PLA), were found to deform when subjected to the high temperatures surrounding




Figure 6.6: Front view of the SEOP oven, showing the four square dressing coils around in the oven.

the oven. This problem persisted even when the oven was wrapped with insulating fabric. To combat this issue, the coils were reprinted using acrylonitrile butadiene styrene (ABS), a plastic with a much higher glass transition temperature of approximately 105 degrees Celsius [MS16]. The struts which hold these ABS coils in place are machined from acrylic (PMMA). The completed and wound dressing coils are shown in figure 6.6. The inductance of the coil was measured to be 1.514 mH and the resistance 1.39 Ω . A capacitor bank with capacitance 71.8 nF was placed in series with the coil to tune its resonance frequency to 15 kHz.

6.2.3.3 Pickup Coil

In order to maximize signal while having a relatively low resonance frequency (as the Larmor frequency for NMR for this system was intended to be on the order of 10 kHz), many turns were required for the pickup coil. We settled on using 2000 turns, hand-wound using the apparatus shown in figure 6.7b.

All components of the pickup coil must be able to withstand the oven's maximum temperature of 200 °C, so we used polyimide-coated magnet wire wound around a custom-made PTFE spool. The resonance frequency of the coil turned out to

(b) Coil winder used to wind the in-situ pickup coil for the SEOP system.

Figure 6.7: Figures showing the design and construction of the pickup coil.




Figure 6.8: Simplified diagram of electronics used for NMR in the SEOP Lab

be approximately 6.2 kHz, though this value varies slightly with temperature. All measurements here are taken with the oven near room temperature.

A diagram of the electronics used in these measurements is shown in figure 6.8.

Control of the dressing and tipping coils are handled by the Redstone NMR system, which is shown in figure 6.10. The Redstone system has many features, but for our purposes it is used as a combination arbitrary waveform generator and data acquisition module. We control the Redstone through the associated software interface (shown in figure 6.9). This software is called TNMR, and is distributed




Figure 6.9: TNMR software interface, showing a sample FID signal.



Figure 6.10: Electronics cabinet, including the Redstone NMR system used to perform data acquisition and arbitrary waveform generation. The LVC2016 amplifier and SR560 pre-amplifier are at the top of the cabinet.

by the same company, TecMag, that manufactures the Redstone system. Several scripts were written which can communicate with TNMR, allowing for the process of running and saving many NMR sequences to be automated. The Redstone can produce up to two arbitrary waveforms with voltage in the range ± 1.6 V. The dressing field waveform is amplified by the AE Techtron 8504 amplifier, a 4 kW amplifier specialized to driving reactive loads. To further improve the performance of the dressing coil system, a set of capacitors is placed in series with the dressing coil to tune the resonance frequency near 15 kHz. A 0.2Ω resistor is also in series with the dressing coil, and the voltage across this resistor, measured using a differential probe, is used as a monitor for the current. For the tipping pulse, the reactive power requirement is much lower because the operating frequency is lower (6.2 kHz vs 15 kHz), and because the amplitude is smaller. Thus, a LVC 2016 linear amplifier is used instead. The pickup coil signal first passes through a preamplifier, which applies a 0.3 kHz-30 kHz bandpass filter to reduce noise.

6.3 NMR Techniques

Two NMR techniques were used to study the ^3He spins in the SEOP lab. These are adiabatic fast passage (AFP) and free induction decay (FID). To understand AFP, it is most convenient to work within the rotating wave approximation, where an oscillating field of constant amplitude becomes a static field. Consider a field (in the lab frame) given by

$$\vec{B}_{\text{lab}}(t) = B_0 \hat{z} + B_1 \sin\left(\int_0^t \omega(t') dt'\right) \hat{x} \quad (6.2)$$

where we allow the frequency of the oscillating field $\omega(t)$ to vary with time. Then in the rotating frame, we get the effective field

$$\vec{B}_{\text{rot}}(t) = (B_0 - \omega(t)/\gamma) \hat{z} + B_1(t) \hat{x}. \quad (6.3)$$

By varying $\omega(t)$, we can change the direction of the effective field $\vec{B}_{\text{rot}}(t)$ from $+\hat{z}$ to $-\hat{z}$ and vice versa. If this change happens slowly enough (adiabatically), the spins will follow the time evolution of the effective magnetic field, and will flip their polarization. A demonstration of the AFP applied to the SEOP cell is shown in figure 6.11. In this test, we use a 0.5 s linear frequency ramp from 5 kHz to 8 kHz immediately followed by a 0.5 s ramp back down to 5 kHz. Note that the AFP pulse itself induces a response in the pickup coil, which has been subtracted in the bottom




Figure 6.11: Pickup coil response during frequency sweep AFP applied to the SEOP cell with polarized ^3He (top) and unpolarized ^3He (middle). The difference between the two (bottom) is also shown.

figure of 6.11 in order to recover the spin polarization signal. As expected, the signal is maximized when the AFP pulse is near the Larmor frequency of the spins.

In FID, we apply a $\pi/2$ pulse or “tipping pulse” to a spin-polarized sample, causing the spins to tip into the plane of precession. Then, using a pickup coil, the oscillation and decay of the magnetic field of the sample is measured. A full $\pi/2$ pulse is appropriate where the timescale for polarizing the sample is short, as is the case for water NMR. For the SEOP system, where polarization takes overnight, a $\pi/2$ pulse would allow us only one measurement per day, which is far too slow. Thus we instead use a small-angle tipping pulse, which tips the spins by a small angle, allowing polarization to be maintained over several FID measurements. The drawback to this is that the signal detected by the pickup coil will have a smaller amplitude. Thus the angle θ_{FID} for the tipping pulse is a tradeoff between two




Figure 6.12: Aluminum frame used for ambient field and B_0 monitoring. Some of the probe locations visible have been circled.

factors: the loss in polarization per FID, which is given by $1 - \cos(\theta_{\text{FID}})$, and the signal amplitude which is proportional to $\sin(\theta_{\text{FID}})$. Generally smaller angles are preferred, so long as the FID peak can be detected above the noise floor. This is because for smaller angles, the loss in polarization is proportional to the square of the angle, while the signal size is proportional to the angle; thus halving the tipping pulse amplitude doubles the number of measurements one can take for a given sample before the signal is below the noise floor. On the other hand, if one is in a situation where little to no signal can be detected, it may be advantageous to use a full $\pi/2$ pulse, if only to diagnose any problems with the SEOP system.

6.4 B_0 and Ambient Field Characterization

The T_2 time of the cell in the absence of dressing fields depends on several factors, including the diffusion time, geometry of the cell, and the static field gradient. We measure this with an 8020 frame built inside the B_0 coils. The field is monitored by 12 single-axis fluxgate magnetic field probes. The frame with probes installed is shown in figure 6.12. For the task of choosing where to place sensors, we take inspiration from [NP14b] and [Man+18]. The former derives a set of basis polynomials in Cartesian coordinates which can be used to decompose an arbitrary magnetic field in free space. The polynomials we use are given in table 6.2. The latter gives a procedure to pick a subset of sensor locations from a set of candidate sensor locations. Specifically, given a set of candidate locations $\vec{x}_i$, where each sensor measures the field along $\hat{e}_i$ and a set of basis functions $\vec{B}_j(\vec{r})$, we define the sensitivity matrix S_{ij}

$$S_{ij} = \vec{B}_j(\vec{x}_i) \cdot \hat{e}_i. \quad (6.4)$$

$B_x(\vec{x})$	$B_y(\vec{x})$	$B_z(\vec{x})$
0	0	1
-1	0	0
-x	-y	2z
-3z	0	-3x
6x	-6y	0
-3xz	-3yz	$-3(x^2 + y^2 - 2z^2)/2$
$3(3x^2 + y^2 - 4z^2)/2$	3xy	-12xz
30xz	-30yz	$15(x - y)(x + y)$
6y	6x	0

Table 6.2: Basis functions used for static field reconstruction.

	dB_x	dB_y	dB_z
dx	-18.5	-4.0	-16.0
dy	-4.0	16.7	0.0
dz	-16.0	0.0	1.8

Table 6.3: Ambient field gradient in $\mu\text{T m}^{-1}$.

	dB_x	dB_y	dB_z
dx	47.6	32.4	-173.9
dy	32.4	117.0	0.0
dz	-173.9	0.0	-164.6

Table 6.4: B_0 coil gradient in $\mu\text{T m}^{-1} \text{ A}^{-1}$.

The authors of [Man+18] prescribe that if n sensors are desired from the set of candidates, then one should take the first n pivot columns [Eng97] in the QR decomposition of the matrix SS^T . This approach turns out to be a greedy algorithm for maximizing the determinant of SS^T , which intuitively should correspond to better sensitivity to the basis function. At the least, maximizing the determinant ensures that there are no linear dependencies among the sensor locations.

With the sensors placed at these locations, we measure the field at each of the sensor locations for six different values of the B_0 current ranging from -0.3 A to 0.2 A . The current had to be limited to this range, as any larger would have saturated the fluxgate probes. We then use the approach of [NP14b] to determine the basis function coefficients and reconstruct the field at the SEOP cell. A linear fit to the field values with respect to current then gives the ambient field gradient and B_0 coil gradient. The final reconstructed gradient values are given in tables 6.3 and 6.4.




Figure 6.13: Frequency response of the pickup coil located in the SEOP oven.

6.5 Polarization Estimate

While the T_2 measurements in this chapter do not rely on the absolute polarization of the cell, it is nonetheless useful to estimate the cell polarization for diagnostic purposes. For this measurement, we mount a small calibration coil some distance underneath the oven. By passing an oscillating current at a frequency near the Larmor frequency ^3He (6.2 kHz, we can approximate the response of the pickup coil to an external magnetic field. First, we need to determine the coupling between the calibration coil and the pickup coil. As the pickup coil's response depends sharply on frequency, we sweep the driving frequency of the calibration coil, which gives us the coupling for a range of frequencies near the Larmor frequency. This measurement is shown in figure 6.13. We denote the response at a given temperature and frequency as $\beta(T, \omega)$, which has units of inverse amperes.

Next, since we generally use small-angle tipping pulses, the angle of the tipping pulse must also be measured. We accomplish this by performing a series of tipping pulses, one after another, on the same sample. After each tipping pulse, we take the amplitude of the FID signal. Suppose that the FID signal after the n th tipping pulse is A_n . Then A_n can be approximated as

$$A_n = A_0 \cos(\theta)^n \quad (6.5)$$

where θ is the tipping angle. This data is shown in figure 6.14. Fitting this data to an exponential decay, we can determine the maximum response assuming a $\pi/2$ pulse, given by $A_{\max} = A_0/\sin(\theta)$. Next, based on the relative positions of the




Figure 6.14: Amplitudes of a sequence of FID measurements, along with a fit to the exponential decay in polarization. The angle is given in radians.

calibration coil, SEOP cell, and pickup coil, we can determine the ratio between the pickup coil response induced by a given current in the calibration coil, and a given magnetization in the SEOP cell. Concretely, the ratio is

$$R(M_0, I_0) = \frac{\int_A B_{\text{cell}}(\vec{r}; M_0)}{\int_A B_{\text{cal}}(\vec{r}, I_0)}, \quad (6.6)$$

where M_0 is the magnetization of the SEOP cell, assumed to be uniform; I_0 is the amplitude of the current in the calibration coil; $B_{\text{cell}}(\vec{r}; M_0)$ is the magnetic field produced by the SEOP cell at location $\vec{r}$; B_{cal} is the magnetic field produced by the calibration coil at location $\vec{r}$; and the integral is taken over the area of the pickup coil. These quantities are simulated using Magpylib. We also need the density of ^3He atoms in the SEOP cell. Fortunately, this was recorded on the box containing the KellyK#2 cell when it was filled, and is 2 amagat (89.2 mol m^{-3}).

Putting all of these quantities together, the polarization of the cell is given by

$$P = \frac{M}{\rho \mu_3} \quad (6.7)$$

$$M = \frac{A_0 R(M_0, I_0)}{\sin(\theta) \beta(T, \omega)}, \quad (6.8)$$

where ρ is the number density of ^3He (2 amagat) and μ_3 is the magnetic moment of ^3He . This gives us a final polarization value of 30%.

6.6 Demonstration of Magic Dressing

In this section we describe the successful demonstration of single-species magic dressing at the Duke SEOP lab. To do this, we artificially apply a dressing field gradient to the SEOP cell and show that a dressing field amplitude can be found that, regardless of gradient magnitude, maximizes the T_2 relaxation time of the system. We conclude with a discussion of the role of magic dressing for a nEDM measurement.

6.6.1 Dressed Tipping

One of the most pernicious difficulties in performing this measurement was a phenomenon we term “dressed tipping.” This refers to the off-resonant excitation of the spins due to a dressing pulse which occurs when the dressing field and the holding field are not exactly orthogonal. This effect is most easily understood in terms of effective magnetic fields. We consider a magnetic field of the form

$$B(t) = B_0 \hat{z} + \delta B \hat{x} + B_1 \sin(\omega t) \hat{x} \quad (6.9)$$

where δB represents misalignment between the dressing and holding fields. Treating δB as a perturbation, we find that this corresponds to the effective magnetic field

$$B_{\text{eff}}(t) = J_0(\gamma B_1 / \omega) B_0 \hat{z} + \delta B \hat{x}. \quad (6.10)$$

The key point here is that as the J_0 term approaches zero, the effective magnetic field direction strays from $\hat{z}$. Thus, spins which are initially aligned along $\hat{z}$ will precess under the effective magnetic field and subsequently become depolarized if the dressing field is suddenly removed. Two techniques were employed to mitigate this phenomenon. First, dressing amplitudes that would put the system near the zero of the Bessel function were avoided; and second, a pair of shim coils were used to reduce the magnitude of the δB . The shim coils — the circular black coils shown in figure 6.4 — provide a field in the $\hat{x}$ direction, whose magnitude is adjusted at the start of each measurement cycle to minimize the size of the dressed tipping effect. For these calibration runs, no tipping pulse is used, as we want to isolate the effect of the dressed tipping phenomenon. Figure 6.15 shows a sample calibration figure. The necessary shim current was found to be relatively consistent from day to day, only varying around 5 mA from day to day. The current regulator used to adjust the shim current was fabricated by a Caltech postdoc, and are designed with current stability as a top priority [Mor23]. The device can be controlled using serial commands via a USB interface, which allowed for the calibration process to be automated.




Figure 6.15: Sample plot demonstrating the shim coil calibration process. The shim coil current is adjusted until the size of the signal caused by the dressed tipping effect is minimized.

6.6.2 Measurement Scheme

Even with the shim coils, the dressed tipping phenomenon discussed above still manifests at larger dressing amplitudes. We thus use the following measurement sequence to investigate magic dressing. The sequence consists of four subsequences:

1. Tipping Pulse, Rest, Acquire,
2. Rest, Dressing Pulse, Acquire,
3. Tipping Pulse, Rest, Acquire,
4. Tipping Pulse, Dressing Pulse, Acquire.

These four are repeated for each dressing amplitude under study. Subsequences (1) and (3) serve as normalization runs for subsequences (2) and (4), respectively, allowing us to account for the inevitable loss of polarization that occurs during the course of the experiment. Subsequence (2) is used to cancel the effect of dressed tipping; by subtracting the normalized signal acquired in step (2) from the normalized signal acquired in step (4), we can isolate the impact that the dressing pulse has on the spins that were tipped by the tipping pulse, and can eliminate the

signal arising from spins that are tipped by the dressing pulse. This process is summarized in equation 6.11.

$$P = \frac{A_4}{A_3} - \frac{A_2}{A_1} \quad (6.11)$$

$A_{\{1-4\}}$ denote the amplitude from subsequences 1-4, respectively, and P is the polarization remaining relative to the initially tipped sample (after the tipping pulse, but before the dressing pulse).

6.6.3 Analysis of Magic Dressing Data

Here we detail the process for transforming the raw pickup coil signal to the T_2 plots shown in figures 6.16, 6.17, 6.18, 6.19, and 6.20. Each of these plots represents a different dressing gradient applied to the SEOP cell. To apply an artificial gradient, we place a pair of coils in anti-Helmholtz configuration outside of the oven and in series with the dressing coil (see the red coils in figure 6.8). We can adjust the number of turns in the anti-Helmholtz coil with a set of switches, and we can also reverse its polarity, thus allowing us to reverse the sign of the dressing gradient applied to the cell. The “reverse orientation” mentioned in the captions for the T_2 plots refers to reversing the polarity of the gradient coils on both sides of the oven. For each gradient setting (defined by the number of turns and gradient coil polarity), we then generate a T_2 plot by performing the measurement scheme outlined in section 6.6.2 for each dressing field amplitude. Recall that a measurement sequence consists of four subsequences, and each of these subsequences results in a pickup coil signal. In all cases, the “amplitude” of a waveform is determined by taking the Fourier transform of the data, finding the peak, and taking the average in a small window around the peak. First, we take the signal from subsequence (2) and divide it by the amplitude of subsequence (1). Likewise, we take the signal of subsequence (4) and divide it by the amplitude of subsequence (3). We next subtract the normalized signal (2) from the normalized signal (4). The final step is to take the amplitude of the difference signal. This is what is plotted in the figures below, averaged over multiple runs at the same dressing amplitude.

For the global fit, Monte Carlo integration is used to approximate the longitudinal decay behavior of the ^3He spins. It is assumed that over the duration of the spin dressing pulse (10 ms), the motion of the ^3He is negligible. This is justified because the diffusion constant of ^3He at room temperature is approximately $1.759 \text{ cm}^2 \text{ s}^{-1}$ [Bur24], which over 10 ms corresponds to a rms change in position by only 0.13 cm. The interior of the cell is approximated to be a 4 cm sphere, accounting for the

thickness of the glass. We also account for the inhomogeneous response of the pickup coil. This is approximated using the fact that the sensitivity of a pickup coil to a dipole at a point in space is proportional to the strength of the magnetic field that would be produced by a current through the pickup coil at that same point in space. For a given current, we call this field $\vec{B}_{\text{pickup}}(\vec{r})$ and model it using Magpylib. We approximate the AC gradient near the center of the coil as being linear in position, and we assume that the static gradient is negligible, as the T_2 time of the SEOP system without dressing is on the order of 1 second. We use the Bessel function approximation to determine the phase shift caused by AC gradient. Putting this all together, the cell polarization is given by

$$P(t) = \frac{1}{A} \left| \int_V dV \left(\vec{B}_{\text{pickup}}(\vec{r}) \cdot \hat{y} \right) \exp \left[\frac{it\gamma^2 B_0}{\omega} J_1 \left(\frac{\gamma B_1}{\omega} \right) \vec{G}_x \cdot \vec{r} \right] \right| \quad (6.12)$$

$$A = \left| \int_V dV \vec{B}_{\text{pickup}}(\vec{r}) \cdot \hat{y} \right| \quad (6.13)$$

where the integral is taken over the SEOP cell volume V , and the origin ($\vec{r} = (0, 0, 0)$) is assumed to be the center of the cell. B_0 and B_1 are the nominal holding field and dressing field amplitudes at the center of the cell, the former determined by the Larmor frequency and the latter determined based on the current measured across the shim resistor in series with the dressing coil. $\vec{G}_x$ is the vector AC gradient $\left(\frac{\partial B_1}{\partial x}, \frac{\partial B_1}{\partial y}, \frac{\partial B_1}{\partial z} \right)$, which is modeled as a sum of two terms: a term independent of the number of gradient coils, $\vec{G}_{\text{base}}$, and a term proportional to the number of gradient coils n , $n\vec{G}_{\text{coil}}$. $\vec{G}_{\text{base}}$ is a fitted parameter, while $\vec{G}_{\text{coil}}$ is modeled in Magpylib based on the known coil dimensions. The reason for this is that $\vec{G}_{\text{base}}$ is intended to model imperfections in the dressing coil construction which lead to gradients; thus, its leading-order terms cannot be inferred from models of the dressing coil; meanwhile, the leading-order gradient terms for the gradient coils are due to their geometry.

For comparison, we also include a plot of a simpler T_2 model, where the polarization is given by

$$P(t) = \left| \text{sinc} \left(\frac{it\gamma^2 B_0}{\omega\pi} J_1 \left(\frac{\gamma B_1}{\omega} \right) |G_x| \cdot \Delta x \right) \right| \quad (6.14)$$

where $|G_x|$ is a fitted parameter representing the average gradient magnitude, Δx is the cell diameter, and $\text{sinc}(x)$ is the sinc function, defined as $\text{sinc}(x) = \sin(x)/x$. For this model, rather than performing a global fit for data across different numbers




Figure 6.16: Polarization of the ^3He spins after a 10-millisecond dressing pulse at 15 kHz. Four gradient coils are active on each side of the other, each in reverse orientation.




Figure 6.17: Polarization of the ^3He spins after a 10-millisecond dressing pulse at 15 kHz. Two gradient coils are active on each side of the oven, each in reverse orientation.

of gradient coil turns, we fit to each gradient coil setting individually. This gives us a rough approximation of the gradient magnitude for each setting.




Figure 6.18: Polarization of the ^3He spins after a 10-millisecond dressing pulse at 15 kHz. All gradient coils are disconnected.




Figure 6.19: Polarization of the ^3He spins after a 10-millisecond dressing pulse at 15 kHz. Two gradient coils are active on each side of the oven, each in standard orientation.




Figure 6.20: Polarization of the ^3He spins after a 10-millisecond dressing pulse at 15 kHz. Four gradient coils are active on each side of the oven, each in standard orientation.

6.6.4 Discussion of Results

From these plots, the effects of single-species magic dressing can be seen clearly. The first feature to note is that the polarization of the sample, and thus the T_2 time, decreases as the dressing amplitude increases, at least for dressing amplitudes less than around 2. This is consistent with the idea that variations in the effective gyromagnetic ratio caused by spatial inhomogeneities in the dressing field amplitude — inhomogeneities that become proportionally larger as the dressing field amplitude is increased — are responsible for transverse relaxation. Second, at a particular magnitude for dressing amplitude corresponding to the first root of the Bessel function of order one ($J_1(x)$), nearly all of the polarization is recovered, regardless of how many turns of the gradient coils are applied. This, in spite of the fact that the spatial inhomogeneities should be significant at this amplitude, indicates that the magic dressing condition has been achieved.

Lastly, we turn to the question of applying magic dressing for an nEDM measurement. Considering the difficulties in dressing the SEOP system at amplitudes beyond the first root of $J_0(x)$, especially the dressed tipping phenomenon and the impedance of a large spin dressing coil, it is unlikely that a scheme identical to this would be practical for a project such as nEDM@SNS. The modification proposed in Chapter 4 resolves this issue, as the average dressing field amplitude is much lower (less than the first root of $J_0(x)$). However, testing the modulation scheme of 4 would require

both ^3He and neutrons and so is outside the scope of this work. Nonetheless, from the experiments on single-species magic dressing, we can gain some insights into how this technique might be implemented for an EDM measurement with a comagnetometer species. A particular advantage of magic dressing is that the tuning of parameters need not be terribly precise. By design, magic dressing ensures that the behavior of the system is, to first order, independent of the amplitude of the dressing field; as shown in the plots above, polarization near unity can be maintained even when the amplitude differs from the “true” magic dressing parameters by a few percent. Thus for an experiment, it may be sufficient to tune the magic dressing parameters by artificially varying the amplitude of the dressing field, observing the parameters for which this variation has the least impact on the precession frequency between the two spin species under study.

CONCLUSION AND OUTLOOK

This dissertation presents an investigation and analysis of spin dressing, a technique where an oscillating magnetic field is used to modify the effective gyromagnetic ratio of a particle, with particular emphasis on the utility of this technique for neutron electric dipole moment (nEDM) experiments such as nEDM@SNS. In chapter one, we introduced the concepts of baryon asymmetry and CP violation, which are the key motivations for performing a precise measurement of the nEDM. In chapter two, we discussed nEDM experiments past, present, and future. We detailed some of the considerations that go into making a next-generation nEDM experiment, especially the use of ultra-cold neutrons (UCNs), which can be stored for hundreds of seconds in suitably designed vessels. We also gave an overview of the design and construction nEDM@SNS, which will use a combination of UCN and superfluid ^4He to achieve unprecedented sensitivity to the nEDM. In chapter three, we analyzed spin dressing from a theoretical standpoint, considering how noise arising from spatial inhomogeneities in the magnetic field or time-dependent field fluctuations could affect the nEDM@SNS experiment. We additionally developed two new modulation schemes, robust dressing and magic dressing, which can be used either to diagnose issues relating to magnetic field gradients or to improve statistical sensitivity of the measurement. In chapter four, we address the computational challenge of performing Monte Carlo simulations of the vast number of particles needed to assess systematics. The presence of a spin dressing field dramatically increases the runtime of these simulations, which motivates the development of two algorithms: the commutator-free exponential time-propagator (CFET) method, which uses the Magnus expansion to accurately approximate the time evolution of the spin vector over short time intervals; and the Floquet-Markov master equation (FMME) algorithm, which uses Floquet theory combined with time-dependent perturbation theory to approximate the density matrix evolution based on an estimate of the noise power spectrum. While these algorithms have seen use in other applications, here we introduce a variety of optimizations and modifications to optimize their performance for the task of computing frequency shifts and relaxation in spin dressing mode. Finally in chapter six, we use the spin-exchange optical pumping (SEOP) apparatus at Duke to produce a polarized sample of ^3He , and use this sample to experimentally verify the

magic dressing technique explored theoretically in chapter three. We successfully demonstrated that the magic dressing condition was achieved, as evidenced by the extended transverse relaxation time for particular values of the dressing field amplitude, even when a large dressing field gradient was applied.

While we aim for this to be a comprehensive analysis of spin dressing, science is never truly done. In the following sections we offer some thoughts on the prospects of the various projects discussed in the dissertation, and speculate about where future investigations might look.

7.1 Spin Dressing Theory: What's Next?

Given the analysis of noise in chapter 4 was primarily concerned with magnetic field fluctuations in the direction parallel to the dressing field, a natural question to ask is how these results may be extended to field fluctuations along all directions. The calculations presented here relied on a particular cancellation of terms that occurs when the perturbation is along the dressing field direction, and so the results will likely be messier, but we expect that the general structure should be the same. One avenue to tackle this problem may be to employ the Floquet-Markov master equation, which was used in chapter 5 to develop a new spin-integration algorithm. Some terms in the master equation, namely ζ and ξ (equations 5.87 and 5.101), may be associated with the frequency shift and relaxation rate of the spins under study. Invoking the Bessel function approximation in the Floquet basis diagonalization, one could then derive a result that holds for all noise directions. This approach could also be applied to more general dressing modulation schemes, provided that these modulation schemes are periodic in time and that a suitable approximation can be found for the diagonalization.

7.2 Bigger Computers, Faster Algorithms

Chapter 5 introduced two algorithms that can dramatically speed up Monte Carlo spin integration, especially for the case of spin dressing. The Magnus CFET method, which uses the Magnus expansion, is intended as a drop-in replacement to standard Runge-Kutta integrators. The main path we see to further optimize the Magnus CFET method is to adjust the adaptive step size subroutine. Currently, this subroutine is nearly identical to that used by DOP853; thus, it is reasonable to expect that by tuning the internal parameters we could obtain modest improvements in performances. One could also consider using different CFET coefficients; Alvermann et. al. [AF11] used numerical root-finding to obtain a set of coefficients which are

thought to have a small approximation error at eighth order, but due to the dimensionality of the coefficient space a complete search was not possible. Therefore, it is possible that a more favorable set of coefficients exists, but this is beyond the scope of this work. The Floquet-Markov master equation (FMME) method, meanwhile, is intended to be used to determine the frequency shifts and relaxation rates for large spin ensembles under time-periodic dressing fields. The FMME Goertzel algorithm could be further optimized by weighting the noise power spectrum based on how important a particular frequency component is to the final result. One could imagine a scheme where additional noise samples (i.e., more particle trajectories) are collected for the frequencies which have a large impact on the final result, thus saving time on the frequencies that matter less. This is only possible with FMME Goertzel because the Goertzel algorithm can compute the discrete Fourier transform at arbitrary frequencies. For the FFT-based algorithm, the main bottleneck is the collection of noise samples. The author is by no means of GPU programming, and it is certainly possible that significant gains in performance could be made by optimizing the memory use patterns during noise sample collection.

Finally, it is worth noting that both methods (Magnus CFET and FMME), and the spin-tracking software suite as a whole, could benefit from a more user-friendly interface. Currently parameters are passed to these methods using a text file, which is cumbersome and requires users to search for the names, permissible values, and default values for simulation parameters. Ideally a graphical user interface could be designed for this package. This could also help the significant work that has gone into this software to reach a wider audience.

7.3 Further SEOP Studies

In chapter 6, we described the SEOP apparatus and used it to demonstrate magic dressing for a single species. A natural next step would be to either attempt single-species robust dressing or two-species magic dressing. For single-species robust dressing on ^3He , the primary challenge will likely be the dressed tipping phenomenon, which will cause the sample to depolarize before robust dressing can be attempted. To mitigate this, a more complete set of shim coils may be helpful. While the SEOP lab at Duke has a set of three field compensation coils exterior to the B_0 coils, these have fallen into disrepair and would need to be rewound; all three coils are shorted to the aluminum frame around the SEOP system. Along with this, a larger set of dressing coils to improve uniformity could be useful. The size of the square dressing coils was primarily motivated by a desire to keep the

inductance of the dressing coil low, but a suitable capacitor in series could allow a larger dressing coil to be energized. This runs into another problem, however, as robust dressing requires fast modulation of the dressing amplitude and/or frequency (for example, modulating a 1 kHz sine wave at 200 Hz 4.2), which may be difficult for a sharply resonant dressing circuit. It remains to be seen whether the power supply in the SEOP lab is powerful enough to drive robust dressing at 15 kHz. For two-species magic dressing, one would need a second species to precess alongside ^3He , preferably neutrons. This is not possible with the current SEOP apparatus, but could potentially be tested at the Systematic and Operational Studies Apparatus at PULSTAR (SOS@PULSTAR) [Cia+25], which will have both neutrons and ^3He .

7.4 nEDMSF: Where We Stand

Despite the loss of funding from the Department of Energy, the nEDM@SNS project continues, though now under the name nEDMSF (SF for “super fluid”). Future efforts will be focused on transitioning the project to operate using neutrons from international sources. The directors of the Institut Laue Langevin (ILL) and the European Spallation Source (ESS), in a letter sent to the Department of Energy and the National Science Foundation, have proposed an international collaboration to examine the feasibility of achieving nEDMSF’s goals at either the neutron source currently at ILL in France, or the ESS, which is under construction in Sweden. Efforts have been underway to investigate the venues best suited to leverage nEDMSF’s unique measurement techniques. While the status of funding for nEDMSF is far from certain, it is clear that the international scientific community recognizes the value of this experiment and its potential to measure the nEDM with an unprecedented level of precision. This gives us hope that, in spite of the many setbacks, the tremendous effort and innovation that has already gone into realizing nEDMSF could continue into the future.

BIBLIOGRAPHY

- [A L06] A. L. Barabanov, R. Golub and S. K. Lamoreaux. “Electric dipole moment searches: effect of linear electric field frequency shifts induced in confined gases.” In: *Physical Review A* 74.5 (2006), 052115(11).
- [Aad+12] G. Aad et al. “Observation of a new particle in the search for the standard model higgs boson with the ATLAS detector at the LHC.” In: *Physics Letters B* 716.1 (Sept. 2012), pp. 1–29. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2012.08.020. URL: <http://dx.doi.org/10.1016/j.physletb.2012.08.020>.
- [Abe+20] C. Abel et al. “Measurement of the permanent electric dipole moment of the neutron.” In: *Physical Review Letters* 124 (8 Feb. 2020), p. 081803. DOI: 10.1103/PhysRevLett.124.081803. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.124.081803>.
- [AF11] A. Alvermann and H. Fehske. “High-order commutator-free exponential time-propagation of driven quantum systems.” In: *Journal of Computational Physics* 230.15 (2011), pp. 5930–5956. ISSN: 0021-9991. DOI: <https://doi.org/10.1016/j.jcp.2011.04.006>. URL: <https://www.sciencedirect.com/science/article/pii/S0021999111002300>.
- [Ago+03] S. Agostinelli et al. “Geant4—a simulation toolkit.” In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (2003), pp. 250–303. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8). URL: <https://www.sciencedirect.com/science/article/pii/S0168900203013688>.
- [Ahm+19] M. Ahmed et al. “A new cryogenic apparatus to search for the neutron electric dipole moment.” In: *Journal of Instrumentation* 14.11 (Nov. 2019), P11017–P11017. DOI: 10.1088/1748-0221/14/11/p11017. URL: <https://doi.org/10.1088/1748-0221/14/11/p11017>.
- [Ala+22] R. Alarcon et al. “Electric dipole moments and the search for new physics.” In: *Snowmass 2021*. Mar. 2022. arXiv: 2203.08103 [hep-ph].
- [Bab05] E. D. Babcock. “Spin-exchange optical pumping with alkali-metal vapors.” PhD thesis. University of Wisconsin-Madison, 2005.
- [Bar48] M. S. Bartlett. “Smoothing periodograms from time-series with continuous spectra.” In: *Nature* 161 (1948), pp. 686–687. URL: <https://api.semanticscholar.org/CorpusID:4068259>.
- [BD16] C. B and R. D. *STARucn*. 2016. URL: <https://sourceforge.net/projects/starucn/>.

- [Ber02] W. Bernreuther. *CP violation and baryogenesis*. 2002. arXiv: hep-ph/0205279 [hep-ph]. URL: <https://arxiv.org/abs/hep-ph/0205279>.
- [Bev+19] G. Bevilacqua et al. “Restoring narrow linewidth to a gradient-broadened magnetic resonance by inhomogeneous dressing.” In: *Physical Review Applied* 11 (2 Feb. 2019), p. 024049. DOI: 10.1103/PhysRevApplied.11.024049. URL: <https://link.aps.org/doi/10.1103/PhysRevApplied.11.024049>.
- [Bev+20] G. Bevilacqua et al. “Harmonic fine tuning and triaxial spatial anisotropy of dressed atomic spins.” In: *Physical Review Letters* 125 (9 Aug. 2020), p. 093203. DOI: 10.1103/PhysRevLett.125.093203. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.125.093203>.
- [Bev+22] G. Bevilacqua et al. “Harmonic dual dressing of spin-1/2 systems.” In: *Physical Review A* 105 (2 Feb. 2022), p. 022619. DOI: 10.1103/PhysRevA.105.022619. URL: <https://link.aps.org/doi/10.1103/PhysRevA.105.022619>.
- [BIS18] D. W. Booth, J. Isaacs, and M. Saffman. “Reducing the sensitivity of Rydberg atoms to dc electric fields using two-frequency ac field dressing.” In: *Physical Review A* 97 (1 Jan. 2018), p. 012515. DOI: 10.1103/PhysRevA.97.012515. URL: <https://link.aps.org/doi/10.1103/PhysRevA.97.012515>.
- [Bla+09] S. Blanes et al. “The magnus expansion and some of its applications.” In: *Physics Reports* 470.5 (2009), pp. 151–238. ISSN: 0370-1573. DOI: <https://doi.org/10.1016/j.physrep.2008.11.001>. URL: <https://www.sciencedirect.com/science/article/pii/S0370157308004092>.
- [BLS11] R. Barbieri, P. Lodone, and D. M. Straub. “CP Violation in Supersymmetry with Effective Minimal Flavour Violation.” In: *Journal of High Energy Physics* 05 (2011), p. 049. DOI: 10.1007/JHEP05(2011)049. arXiv: 1102.0726 [hep-ph].
- [Blü+91] R. Blümel et al. “Dynamical localization in the microwave interaction of rydberg atoms: the influence of noise.” In: *Physical Review A* 44 (7 Oct. 1991), pp. 4521–4540. DOI: 10.1103/PhysRevA.44.4521. URL: <https://link.aps.org/doi/10.1103/PhysRevA.44.4521>.
- [BNN77] S. R. Barone, M. A. Narcowich, and F. J. Narcowich. “Floquet theory and applications.” In: *Physical Review A* 15 (3 Mar. 1977), pp. 1109–1125. DOI: 10.1103/PhysRevA.15.1109. URL: <https://link.aps.org/doi/10.1103/PhysRevA.15.1109>.

- [Bra+12] G. Branco et al. “Theory and phenomenology of two-Higgs-doublet models.” In: *Physics Reports* 516.1 (2012). Theory and phenomenology of two-Higgs-doublet models, pp. 1–102. ISSN: 0370-1573. DOI: <https://doi.org/10.1016/j.physrep.2012.02.002>. URL: <https://www.sciencedirect.com/science/article/pii/S0370157312000695>.
- [Bur24] D. Burgess Jr. *Self-diffusion and binary-diffusion coefficients in gases*. Technical Note (NIST TN). National Institute of Standards and Technology, 2024. DOI: [10.6028/NIST.TN.2279](https://doi.org/10.6028/NIST.TN.2279). URL: https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=957163.
- [C C69] C. Cohen-Tannoudji, S. Haroche. “Absorption et diffusion de photons optiques par un atome en interaction avec des photons de radiofréquence.” In: *Journal de Physique* 30.2-3 (1969), pp. 153–168. DOI: [10.1051/jphys:01969003002-3015300](https://doi.org/10.1051/jphys:01969003002-3015300). URL: <https://doi.org/10.1051/jphys:01969003002-3015300>.
- [Cav95] T. Cavallo. “On the methods of manifesting the presence, and ascertaining the quality, of small quantities, of natural or artificial electricity.” In: *A complete treatise on electricity, in theory and practice; with original experiments*. 4th ed. Vol. 3. C. Dilly, London, U.K., 1795, pp. 98–107.
- [Chr+64] J. H. Christenson et al. “Evidence for the 2π Decay of the K_2^0 Meson.” In: *Physical Review Letters* 13 (1964), pp. 138–140. DOI: [10.1103/PhysRevLett.13.138](https://doi.org/10.1103/PhysRevLett.13.138).
- [Cia+25] V. Cianciolo et al. *The Systematics and Operational Studies (SOS) Apparatus as a testbed for nEDM@SNS experiment*. 2025. arXiv: 2411.03337 [physics.ins-det]. URL: <https://arxiv.org/abs/2411.03337>.
- [Cla+18] S. Clayton et al. “Cavallo’s multiplier for in situ generation of high voltage.” In: *Journal of Instrumentation* 13.05 (May 2018), P05017. DOI: [10.1088/1748-0221/13/05/P05017](https://doi.org/10.1088/1748-0221/13/05/P05017). URL: <https://dx.doi.org/10.1088/1748-0221/13/05/P05017>.
- [Cla11] S. M. Clayton. “Spin relaxation and linear-in-electric-field frequency shift in an arbitrary, time-independent magnetic field.” In: *Journal of Magnetic Resonance* 211.1 (2011), pp. 89–95.
- [CM 12] C.M. Swank, A.K. Petukhov and R. Golub. “Correlation functions for restricted brownian motion from the ballistic through to the diffusive regimes.” In: *Physics Letters A* 376.34 (2012), pp. 2319–2324.
- [Com91] E. D. Commins. “Berry’s geometric phase and motional fields.” In: *American Journal of Physics* 59.12 (1991), pp. 1077–1080.

- [DB98] R. J. Donnelly and C. F. Barenghi. “The observed properties of liquid helium at the saturated vapor pressure.” In: *Journal of Physical and Chemical Reference Data* 27.6 (Nov. 1998), pp. 1217–1274. ISSN: 0047-2689. DOI: 10.1063/1.556028. eprint: https://pubs.aip.org/aip/jpr/article-pdf/27/6/1217/11342901/1217_1_online.pdf. URL: <https://doi.org/10.1063/1.556028>.
- [Dip19] R. Dipert. “Development of cryogenic detection systems for a search of the neutron electric dipole moment.” PhD thesis. Arizona State University, 2019.
- [DNN08] S. Davidson, E. Nardi, and Y. Nir. “Leptogenesis.” In: *Physics Reports* 466.4–5 (Sept. 2008), pp. 105–177. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2008.06.002. URL: <http://dx.doi.org/10.1016/j.physrep.2008.06.002>.
- [Eck+12] S. Eckel et al. “Time-dependent spin dressing using a ^3He atomic beam.” In: *Physical Review A* 85 (3 Mar. 2012), p. 032124. DOI: 10.1103/PhysRevA.85.032124. URL: <https://link.aps.org/doi/10.1103/PhysRevA.85.032124>.
- [Eng97] H. Engler. “The behavior of the QR-factorization algorithm with column pivoting.” In: *Applied Mathematics Letters* 10.6 (1997), pp. 7–11. ISSN: 0893-9659. DOI: [https://doi.org/10.1016/S0893-9659\(97\)00098-0](https://doi.org/10.1016/S0893-9659(97)00098-0). URL: <https://www.sciencedirect.com/science/article/pii/S0893965997000980>.
- [Far+94] K. Farakos et al. “The electroweak phase transition at $m_H \simeq m_W$.” In: *Physics Letters B* 336.3–4 (Sept. 1994), pp. 494–501. ISSN: 0370-2693. DOI: 10.1016/0370-2693(94)90563-0. URL: [http://dx.doi.org/10.1016/0370-2693\(94\)90563-0](http://dx.doi.org/10.1016/0370-2693(94)90563-0).
- [Fil22] B. Filippone. “Overview of neutron EDM (nEDM) searches.” Presented at Fundamental Symmetries, Neutrons, and Neutrinos Town Meeting, UNC Chapel Hill. 2022. URL: <https://indico.phy.ornl.gov/event/209/contributions/1104/>.
- [Flo83] G. Floquet. “Sur les équations différentielles linéaires à coefficients périodiques.” fr. In: *Annales scientifiques de l’École Normale Supérieure* 2e série, 12 (1883), pp. 47–88. DOI: 10.24033/asens.220. URL: <http://www.numdam.org/articles/10.24033/asens.220/>.
- [Fur+17] D. Furse et al. “Kassiopeia: A modern, extensible C++ particle tracking package.” In: *New Journal of Physics* 19.5, 053012 (May 2017), p. 053012. DOI: 10.1088/1367-2630/aa6950. arXiv: 1612.00262 [physics.comp-ph].
- [Gav+94] M. Gavela et al. “Standard model CP-violation and baryon asymmetry (ii). finite temperature.” In: *Nuclear Physics B* 430.2 (1994), pp. 382–426. ISSN: 0550-3213. DOI: <https://doi.org/10.1016/0550->

- 3213(94) 00410–2. URL: <https://www.sciencedirect.com/science/article/pii/S0550321394004102>.
- [GH98] M. Grifoni and P. Hänggi. “Driven quantum tunneling.” In: *Physics Reports* 304.5 (1998), pp. 229–354. ISSN: 0370-1573. DOI: [https://doi.org/10.1016/S0370-1573\(98\)00022-2](https://doi.org/10.1016/S0370-1573(98)00022-2). URL: <https://www.sciencedirect.com/science/article/pii/S0370157398000222>.
- [GL94] R. Golub and S. K. Lamoreaux. “Neutron electric-dipole moment, ultracold neutrons and polarized ^3He .” In: *Physics Reports* 237.1 (1994), pp. 1–62. ISSN: 0370-1573. DOI: [https://doi.org/10.1016/0370-1573\(94\)90084-1](https://doi.org/10.1016/0370-1573(94)90084-1). URL: <https://www.sciencedirect.com/science/article/pii/0370157394900841>.
- [GLW57] R. L. Garwin, L. M. Lederman, and M. Weinrich. “Observations of the failure of conservation of parity and charge conjugation in meson decays: the magnetic moment of the free muon.” In: *Physical Review* 105 (4 Feb. 1957), pp. 1415–1417. DOI: [10.1103/PhysRev.105.1415](https://doi.org/10.1103/PhysRev.105.1415). URL: <https://link.aps.org/doi/10.1103/PhysRev.105.1415>.
- [Goe58] G. Goertzel. “An algorithm for the evaluation of finite trigonometric series.” In: *The American Mathematical Monthly* 65.1 (1958), pp. 34–35. ISSN: 00029890, 19300972. URL: <http://www.jstor.org/stable/2310304> (visited on 05/29/2024).
- [Gon+21] F. M. Gonzalez et al. “Improved neutron lifetime measurement with UCN τ .” In: *Physical Review Letters* 127 (16 Oct. 2021), p. 162501. DOI: [10.1103/PhysRevLett.127.162501](https://doi.org/10.1103/PhysRevLett.127.162501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.127.162501>.
- [Gre+79] G. L. Greene et al. “Measurement of the neutron magnetic moment.” In: *Physical Review D* 20 (9 Nov. 1979), pp. 2139–2153. DOI: [10.1103/PhysRevD.20.2139](https://doi.org/10.1103/PhysRevD.20.2139). URL: <https://link.aps.org/doi/10.1103/PhysRevD.20.2139>.
- [HK92] M. Hellmund and J. Kripfganz. “The decay of the sphaleron.” In: *Nuclear Physics B* 373.3 (1992), pp. 749–758. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/0550-3213\(92\)90274-F](https://doi.org/10.1016/0550-3213(92)90274-F). URL: <https://www.sciencedirect.com/science/article/pii/S055032139290274F>.
- [Hoo19] A. Hook. “TASI Lectures on the Strong CP Problem and Axions.” In: *Proceedings of Science TASI2018* (2019), p. 004. DOI: [10.22323/1.333.0004](https://doi.org/10.22323/1.333.0004). arXiv: 1812.02669 [hep-ph].
- [HS95] P. Huet and E. Sather. “Electroweak baryogenesis and the standard model CP violation.” In: *Physical Review D* 51.2 (1995), pp. 379–394.

- [Kat+03] H. Katori et al. “Ultrastable optical clock with neutral atoms in an engineered light shift trap.” In: *Physical Review Letters* 91 (17 Oct. 2003), p. 173005. DOI: 10.1103/PhysRevLett.91.173005. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.91.173005>.
- [KS15] G. A. Kazakov and T. Schumm. “Magic radio-frequency dressing for trapped atomic microwave clocks.” In: *Physical Review A* 91 (2 Feb. 2015), p. 023404. DOI: 10.1103/PhysRevA.91.023404. URL: <https://link.aps.org/doi/10.1103/PhysRevA.91.023404>.
- [Lam+24] N. Lambert et al. *QuTiP 5: The Quantum Toolbox in Python*. 2024. arXiv: 2412.04705 [quant-ph]. URL: <https://arxiv.org/abs/2412.04705>.
- [LE07] C.-P. Liu and J. Engel. “Schiff screening of relativistic nucleon electric-dipole moments by electrons.” In: *Physical Review C* 76.2 (Aug. 2007). ISSN: 1089-490X. DOI: 10.1103/physrevc.76.028501. URL: <http://dx.doi.org/10.1103/PhysRevC.76.028501>.
- [Leu25] K. Leung. “Neutron EDM with superfluid 4He (and 3He) experiment.” Electric Dipole Moments: Experimental and Theoretical Horizons. 2025. URL: <https://indico-tdli.sjtu.edu.cn/event/4106/>.
- [LG05] S. K. Lamoreaux and R. Golub. “Detailed discussion of a linear electric field frequency shift induced in confined gases by a magnetic field gradient: implications for neutron electric-dipole-moment experiments.” In: *Physical Review A* 71 (3 Mar. 2005), p. 032104. DOI: 10.1103/PhysRevA.71.032104. URL: <http://link.aps.org/doi/10.1103/PhysRevA.71.032104>.
- [LG09] S. K. Lamoreaux and R. Golub. “Experimental searches for the neutron electric dipole moment.” In: *Journal of Physics G: Nuclear and Particle Physics* 36.10 (Sept. 2009), p. 104002. DOI: 10.1088/0954-3899/36/10/104002. URL: <https://dx.doi.org/10.1088/0954-3899/36/10/104002>.
- [LL80] L. Landau and E. Lifshitz. “Chapter I - the fundamental principles of statistical physics.” In: *Statistical Physics*. Ed. by L. Landau and E. Lifshitz. Third Edition. Oxford: Butterworth-Heinemann, 1980, pp. 1–33. ISBN: 978-0-08-057046-4. DOI: <https://doi.org/10.1016/B978-0-08-057046-4.50008-7>. URL: <https://www.sciencedirect.com/science/article/pii/B9780080570464500087>.
- [Lud54] G. Luders. “On the equivalence of invariance under time reversal and under particle-antiparticle conjugation for relativistic field theories.” In: *Det Kongelige Danske Videnskabernes Selskab Matematisk-fysiske Meddelelser* 28N5.5 (1954), pp. 1–17.

- [Mag+10] S. Magdaleno-Adame et al. “Coil systems to generate uniform magnetic field volumes.” In: Oct. 2010.
- [Man+18] K. Manohar et al. “Data-driven sparse sensor placement for reconstruction: demonstrating the benefits of exploiting known patterns.” In: *IEEE Control Systems Magazine* 38.3 (2018), pp. 63–86. doi: 10.1109/MCS.2018.2810460.
- [McG90] D. D. McGregor. “Transverse relaxation of spin-polarized ^3He gas due to a magnetic field gradient.” In: *Physical Review A* 45.5 (1990), pp. 2631–2635.
- [McK+04] D. McKinsey et al. “Detecting ionizing radiation in liquid helium using wavelength shifting light collection.” In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 516.2 (2004), pp. 475–485. ISSN: 0168-9002. doi: <https://doi.org/10.1016/j.nima.2003.08.152>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900203024987>.
- [Mor23] M. J. Morano. “Development of hardware and simulations for the nEDM@SNS experiment.” PhD thesis. NC State, 2023.
- [MS16] N. R. S. Miftahur Rahman and L. K. Sadhu. “Glass transition of ABS in 3D printing.” In: *2016 COMSOL Conference in Boston*. 2016.
- [NIS] NIST. *Fundamental Physical Constants*. (<https://www.nist.gov/pml/fundamental-physical-constants>).
- [NP14a] N. Nouri and B. Plaster. “Sensitivity requirements for accessing interior magnetic field vector components in neutron electric dipole moment experiments via exterior boundary-value measurements.” In: *Journal of Instrumentation* 9.11 (2014), P11009. URL: <http://stacks.iop.org/1748-0221/9/i=11/a=P11009>.
- [NP14b] N. Nouri and B. Plaster. “Systematic optimization of exterior measurement locations for the determination of interior magnetic field vector components in inaccessible regions.” In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 767 (2014), pp. 92–98. ISSN: 0168-9002. doi: <https://doi.org/10.1016/j.nima.2014.08.026>. URL: <https://www.sciencedirect.com/science/article/pii/S0168900214009541>.
- [NVI25] NVIDIA Corporation. *cuFFT API reference*. 2025. URL: <https://docs.nvidia.com/cuda/cufft/index.html#accuracy-and-performance>.

- [OOR18] O. M. Ogreid, P. Osland, and M. N. Rebelo. “CP violation in extended Higgs sectors.” In: *Proceedings of Corfu Summer Institute 2017 “Schools and Workshops on Elementary Particle Physics and Gravity” — PoS(CORFU2017)*. Vol. 318. 2018, p. 044. DOI: 10.22323/1.318.0044.
- [PAC17] PACE. *Partnership for an Advanced Computing Environment (PACE)*. 2017. URL: <http://www.pace.gatech.edu>.
- [Pap65] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw Hill Book Company, 1965.
- [PC65] N. Polonsky and C. Cohen-Tannoudji. “Interprétation quantique de la modulation de fréquence.” In: *Journal de Physique* 26 (7 July 1965), pp. 409–414. DOI: 10.1051/jphys:01965002607040900.
- [Pen+04a] J. M. Pendlebury et al. “Geometric-phase-induced false electric dipole moment signals for particles in traps.” In: *Physical Review A* 70 (3 Sept. 2004), p. 032102. DOI: 10.1103/PhysRevA.70.032102. URL: <http://link.aps.org/doi/10.1103/PhysRevA.70.032102>.
- [Pen+04b] J. Pendlebury et al. “Geometric-phase-induced false electric dipole moment signals for particles in traps.” In: *Physical Review A* 70.3 (2004), p. 032102.
- [Pha23] N. S. Phan. “High voltage studies and electrode development for the nedm@sns experiment.” nEDM2023 Workshop. 2023. URL: <https://indico.miserver.it.umich.edu/event/5/contributions/48/>.
- [Pig+15] G. Pignol et al. “Frequency shifts and relaxation rates for spin-1/2 particles moving in electromagnetic fields.” In: *Physical Review A* 92 (5 Nov. 2015), p. 053407. DOI: 10.1103/PhysRevA.92.053407. URL: <https://link.aps.org/doi/10.1103/PhysRevA.92.053407>.
- [Pig19] G. Pignol. “A magic magnetic field to measure the neutron electric dipole moment.” In: *Physics Letters B* 793 (June 2019), pp. 440–444. ISSN: 0370-2693. DOI: 10.1016/j.physletb.2019.05.014. URL: <http://dx.doi.org/10.1016/j.physletb.2019.05.014>.
- [PR00] M. Pospelov and A. Ritz. “Theta vacua, QCD sum rules, and the neutron electric dipole moment.” In: *Nuclear Physics B* 573.1 (2000), pp. 177–200. ISSN: 0550-3213. DOI: [https://doi.org/10.1016/S0550-3213\(99\)00817-2](https://doi.org/10.1016/S0550-3213(99)00817-2). URL: <https://www.sciencedirect.com/science/article/pii/S0550321399008172>.
- [R G08] R. Golub, C.M. Swank, S.K. Lamoreaux. “Electric dipole moment searches: Effect of linear electric field frequency shifts induced in confined gases, II.” In: *arXiv:0810.5378v3* (2008).

- [R G91] R. Golub, D. Richardson, S. K. Lamoreaux. *Ultra-cold neutrons*. IOP Publishing Ltd, 1991. ISBN: 0-7503-0115-5.
- [Ram50] N. F. Ramsey. “A molecular beam resonance method with separated oscillating fields.” In: *Physical Review* 78 (6 June 1950), pp. 695–699. DOI: 10.1103/PhysRev.78.695. URL: <https://link.aps.org/doi/10.1103/PhysRev.78.695>.
- [Red57] A. G. Redfield. “On the theory of relaxation processes.” In: *IBM Journal of Research and Development* 1.1 (1957), pp. 19–31. DOI: 10.1147/rd.11.0019.
- [RN17] C. Rackauckas and Q. Nie. “DifferentialEquations.jl—a performant and feature-rich ecosystem for solving differential equations in Julia.” In: *Journal of Open Research Software* 5.1 (2017).
- [Sak] A. D. Sakharov. “Violation of CP-invariance, C-asymmetry, and baryon asymmetry of the universe.” In: *In the Intermissions...* Pp. 84–87. DOI: 10.1142/9789812815941_0013. eprint: https://www.worldscientific.com/doi/pdf/10.1142/9789812815941_0013. URL: https://www.worldscientific.com/doi/abs/10.1142/9789812815941_0013.
- [Sch+17] W. Schreyer et al. “PENTTrack—a simulation tool for ultracold neutrons, protons, and electrons in complex electromagnetic fields and geometries.” In: *Nuclear Instruments and Methods in Physics Research Section A* 858 (2017), pp. 123–129. DOI: 10.1016/j.nima.2017.03.036. arXiv: 1610.06358 [physics.ins-det].
- [Sea92] V. F. Sears. “Neutron scattering lengths and cross sections.” In: *Neutron News* 3.3 (1992), pp. 26–37. DOI: 10.1080/10448639208218770. eprint: <https://doi.org/10.1080/10448639208218770>. URL: <https://doi.org/10.1080/10448639208218770>.
- [Sha49] C. Shannon. “Communication in the presence of noise.” In: *Proceedings of the IRE* 37.1 (1949), pp. 10–21. DOI: 10.1109/JRPROC.1949.232969.
- [SHD91] S. D. Stoller, W. Happer, and F. J. Dyson. “Transverse spin relaxation in inhomogeneous magnetic fields.” In: *Physical Review A* 44 (11 Dec. 1991), pp. 7459–7477. DOI: 10.1103/PhysRevA.44.7459. URL: <https://link.aps.org/doi/10.1103/PhysRevA.44.7459>.
- [Sli90] C. P. Slichter. “Spin-lattice relaxation and motional narrowing of resonance lines.” In: *Principles of Magnetic Resonance*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1990, pp. 145–218. ISBN: 978-3-662-09441-9. DOI: 10.1007/978-3-662-09441-9_5. URL: https://doi.org/10.1007/978-3-662-09441-9_5.

- [SPG16] C. M. Swank, A. K. Petukhov, and R. Golub. “Random walks with thermalizing collisions in bounded regions: physical applications valid from the ballistic to diffusive regimes.” In: *Physical Review A* 93 (6 June 2016), p. 062703. DOI: 10.1103/PhysRevA.93.062703. URL: <http://link.aps.org/doi/10.1103/PhysRevA.93.062703>.
- [Ste+86] A. Steyerl et al. “A new source of cold and ultracold neutrons.” In: *Physics Letters A* 116.7 (1986), pp. 347–352. ISSN: 0375-9601. DOI: [https://doi.org/10.1016/0375-9601\(86\)90587-6](https://doi.org/10.1016/0375-9601(86)90587-6). URL: <https://www.sciencedirect.com/science/article/pii/0375960186905876>.
- [Sun22] X. Sun. “Measurements of beyond standard model interactions with the UCNA and nEDM@SNS experiments.” PhD thesis. California Institute of Technology, 2022.
- [Swa+18] C. M. Swank et al. “Spin-dressed relaxation and frequency shifts from field imperfections.” In: *Physical Review A* 98 (5 Nov. 2018), p. 053414. DOI: 10.1103/PhysRevA.98.053414. URL: <https://link.aps.org/doi/10.1103/PhysRevA.98.053414>.
- [Tak+20] A. Takenaka et al. “Search for proton decay via $p \rightarrow e^+\pi^0$ and $p \rightarrow \mu^+\pi^0$ with an enlarged fiducial volume in Super-Kamiokande I-IV.” In: *Physical Review D* 102 (11 Dec. 2020), p. 112011. DOI: 10.1103/PhysRevD.102.112011. URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.112011>.
- [Tat25] R. Tat. “Semiclassical Floquet-Markov master equation for Monte Carlo spin integration.” In: *Phys. Rev. A* 112 (1 July 2025), p. 013113. DOI: 10.1103/2bp2-hs9d. URL: <https://link.aps.org/doi/10.1103/2bp2-hs9d>.
- [TC04] M. Trodden and S. M. Carroll. *TASI lectures: introduction to cosmology*. 2004. arXiv: astro-ph/0401547 [astro-ph]. URL: <https://arxiv.org/abs/astro-ph/0401547>.
- [TS22] R. Tat and C. M. Swank. “Quantum control of two critically dressed spin-1/2 species in magnetic fluctuations.” In: *Phys. Rev. A* 105 (5 May 2022), p. 053120. DOI: 10.1103/PhysRevA.105.053120. URL: <https://link.aps.org/doi/10.1103/PhysRevA.105.053120>.
- [Tur60] G. Turin. “An introduction to matched filters.” In: *IRE Transactions on Information Theory* 6.3 (1960), pp. 311–329. DOI: 10.1109/TIT.1960.1057571.
- [WH97] T. G. Walker and W. Happer. “Spin-exchange optical pumping of noble-gas nuclei.” In: *Reviews of Modern Physics* 69 (2 Apr. 1997), pp. 629–642. DOI: 10.1103/RevModPhys.69.629. URL: <https://link.aps.org/doi/10.1103/RevModPhys.69.629>.

- [Wil67] J. Wilks. *The properties of liquid and solid helium*. Oxford, 1967.
- [WNH93] G. Wanner, S. Norsett, and E. Hairer. *Solving ordinary differential equations i*. Berlin Heidelberg New York: Springer, 1993.
- [Wu+57] C. S. Wu et al. “Experimental test of parity conservation in beta decay.” In: *Physical Review* 105 (4 Feb. 1957), pp. 1413–1415. doi: 10.1103/PhysRev.105.1413. URL: <https://link.aps.org/doi/10.1103/PhysRev.105.1413>.
- [ZCA12] T. Zanon-Willette, E. de Clercq, and E. Arimondo. “Magic radio-frequency dressing of nuclear spins in high-accuracy optical clocks.” In: *Physical Review Letters* 109 (22 Nov. 2012), p. 223003. doi: 10.1103/PhysRevLett.109.223003. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.109.223003>.
- [Zsi17] G. Zsigmond. “The MCUCN simulation code for ultracold neutron physics.” In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 881 (Sept. 2017). doi: 10.1016/j.nima.2017.10.065.