Abstract of "Dynamic Simulations of Strongly Coupled Spin Ensembles", by Charles Snider, Brown University, May 2022

We develop an efficient simulation package of nuclear magnetic resonance spin echo experiments in order to study the effects of strong electronic spin correlations on the dynamics of the nuclear spin ensemble. A mean-field model can be used to study long range correlated electronic phases through their hyperfine interaction with nuclear spins. We first review specific aspects of the simulation package developed for our exploration. We discuss its structure, accuracy, and the technical merits of the various approximations used to model the spin ensemble. We then explore the dynamics of the interacting nuclear ensemble and discuss the key behaviors of the system. In particular, we classify the types of temporal asymmetry the interaction induces in the system as well as a pulse-dependent shift in the spectral domain. Using these results, we discuss how careful measurement of the pulse-dependent shift can be used to extract information about the anisotropy of the interaction and how these results represent a novel tool for the examination of exotic NMR signatures in strongly correlated systems.

Dynamic Simulations of Strongly Coupled Spin Ensembles

by

Charles Snider

A dissertation submitted in partial fulfillment of the
requirements for the Degree of Doctor of Philosophy
in the Department of Physics at Brown University

Brown University

May 2022

This dissertation by Charles Snider is accepted in its present form
by the Department of Physics as satisfying the
dissertation requirement for the degree of Doctor of Philosophy.

Date _____ _____
Professor Vesna Mitrović, Advisor

Recommended to the Graduate Council

Date _____ _____
Professor Dmitri Feldman, Reader

Date _____ _____
Professor Brad Marston, Reader

Approved by the Graduate Council

Date _____ _____
Andrew Campbell,
Dean of the Graduate School

Curriculum Vitae

Charles Snider

Education

1. **Brown University** (2017 - Present, *In Progress*): Doctor of Philosophy in Physics
2. **Brown University** (2020): Master of Science in Physics (*Transitional*)
3. **Brown University** (2012-2016): Bachelor of Science in Physics

Publications

1. S. Carr, **C Snider**, D. E. Feldman, C. Ramanathan, J. B. Marston, V. F. Mitrović, "Signatures of electronic correlations and spin-susceptibility anisotropy in nuclear magnetic resonance", under review
2. **C. Snider**, S. Carr, D. E. Feldman, C. Ramanathan, J. B. Marston, V. F. Mitrović, "Dynamic Simulations of Strongly Coupled Spin Ensembles", to be submitted to Computer Physics Communications

Presentations

1. "Spin Echo Simulations and Supervised Machine Learning", APS March Meeting, 2021 (virtual)

Acknowledgments

I would like to thank my advisor, Professor Vesna Mitrović, for her support through the last five years. She gave me the freedom and flexibility to explore the research I was passionate about, and she supported my remote status which allowed me to move with my then-girlfriend, now soon-to-be wife, to Pennsylvania. Without her support, I would not be where I am in my career or my life, and for that I will always be grateful.

I would also like to thank Dr. Stephen Carr for all his help over the years. His suggestions helped craft this project into what is is today, and his advice has been invaluable in getting me to this point. I would like to thank Professors Dmitri Feldman and Brad Marston of Brown University and Professor Chandrasekhar Ramanathan of Dartmouth College for their ideas and input from the earliest days of this project.

Finally, I would like to dedicate this dissertation to my grandparents, John and Jane Bell and Charles and Joan Snider, whose support made all of this possible and who believed in me before I did.

Contents

List of Tables **viii**

List of Figures **ix**

1 Introduction **1**

1.1 NMR and Strongly Correlated Electrons 3

1.2 Overview 10

2 Theoretical Background **13**

2.1 NMR Basics 13

2.1.1 The RF Pulse 16

2.1.2 The π Pulse and the Spin Echo 18

2.1.3 The Knight Shift 21

2.1.4 Relaxation 22

2.2 Model 23

2.2.1 The Nuclear Ensemble 26

2.2.2 Pulses and Time Evolution 27

2.2.3 Interactions in the Rotating Frame 28

2.2.4 Magnus Expansion 31

2.2.5 Effective Time Scale Ω 36

2.2.6 Liouville Space 38

3 Simulation **42**

3.1 Simulation Structure 46

3.1.1 Local Magnetization: $M_{d,L}^{(ij)}$ 48

3.1.2 Matrix Exponent 52

3.2 Performance 56

3.3 Validation 60

3.4 Baseline Numerical Error 62

3.5 Convergence 64

3.5.1 Convergence Relative to n 66

3.5.2 Convergence Relative to dt 71

3.5.3 The Magnus Expansion 76

4 Results **79**

4.1 Global Coupling 79

4.2 Local Interactions 89

4.2.1 A Quantum Treatment 97

4.2.2 A Note on Dissipation 100

4.3 Γ as a Universal Scale 102

4.4 Application to Experiment 108

5 Conclusions and Future Research **110**

List of Tables

3.1 Profiling of Julia simulation 56

3.2 Profiling of GPU-Julia simulation 57

3.3 Profiling of CUDA simulation 58

4.1 Fit coefficients for scaling of ω for different linewidths Γ 106

List of Figures

2.1 The spin echo 20

3.1 Error ε of CUDA simulation relative to GPU-Julia simulation versus ω 61
3.2 Baseline numerical error 64
3.3 Convergence relative to ensemble size n 68
3.4 Drift of error ε of echoes of different n versus time 70
3.5 Convergence versus time step dt 74
3.6 Drift in error ε for echoes with different dt versus time 75
3.7 Magnus expansion versus high sampling rate 77

4.1 Spin projections and phase locking 80
4.2 Interaction dynamics and their occurrence during the spin echo . . . 83
4.3 Echo behaviors for a global interaction 84
4.4 Spectral behaviors for a global interaction 85
4.5 Echo behaviors for a global interaction versus flip angle 87
4.6 Spin echo spectral behaviors for a global interaction versus flip angle and interaction anisotropy 88
4.7 FID spectral behaviors for a global interaction versus flip angle and interaction anisotropy 89
4.8 Echo behaviors for a local interaction for varying ξ 90
4.9 First moment of echoes for a local interaction versus ω, ξ 92

4.10 Echo behavior for varying ω_z and ξ with $\omega = 3.15\Gamma$. Echo amplitudes are not normalized to accentuate the attenuation from the introduction of ω_z. 93
4.11 Spectral behavior for a local interaction for varying ξ 94
4.12 Spectral behavior for a local interaction for varying ξ, ω_z 96
4.13 Spectral behavior for a local interaction with $\xi = 2.0$ for varying flip angle and interaction anisotropy 97
4.14 Spectral behaviors for a global interaction with dissipation 101
4.15 Error ε for equivalent echoes at different linewidths Γ 104
4.16 Comparison of equivalent echoes for different linewidths Γ 105
4.17 Echo comparison for equivalent echoes at different linewidths Γ . . . 107

Chapter 1

Introduction

The study of strongly correlated electron systems is the main focus of condensed matter physics. These systems can play host to a great variety of unique and complex physical phenomena that can be challenging both to theoretically model and to experimentally identify. Much experimental work has gone into charting out the location of and conditions for novel electronic and magnetic phases arising from these strong correlations, and theoretical explorations have brought to the forefront the potential for a wide variety of new phenomena.

One topic of correlated electronic systems to which much attention is given is that of unconventional superconductivity. Such unconventional superconductors display behaviors that diverge from that of normal BCS-type superconductors, and the driving cause of superconductivity in these systems– that is, the mechanism of Cooper pairing– is still a subject of much speculation. In particular, two types of unconventional superconductivity that have garnered a lot of attention are triplet superconductivity, such as that once believed to be present in Sr_2RuO_4, and the so-called "Fulde-Ferell-Larkin-Ovchinnikov" (FFLO) state [1].

In triplet superconductors, the spin state of the Cooper pairs is given by a superposition of the symmetric spin triplet states, rather than the antisymmetric singlet

state characteristic of BCS superconductivity. Because of symmetry constraints on the overall state of the Cooper pairs, the formation of triplet spin states has profound implications on the orbital component of the wavefunction [2]. The triplet spin states can result in novel behavior of the electronic susceptibility in these superconductors, and the presence of higher order orbital states in the Cooper pairs can lead to momentum-space nodes in the gap [2].

Like with triplet superconductors, the FFLO state exhibits nodes in the superconducting gap. However, unlike triplet superconductors, whose nodes occur in momentum space, the FFLO state has nodes in real space on account of Cooper pairs forming with non-zero momentum [3]. The regime in which the formation of an FFLO state is possible is extremely restrictive, and it can be highly sensitive to impurities [3]. The severe limitations of the formation and stability of the FFLO state limits the materials in which its formation is viable and makes it difficult to experimentally observe.

The presence of strongly correlated electrons, and the exotic phases that result, can result in complex dynamics that are difficult to measure and experimentally classify. Nuclear Magnetic Resonance (NMR) has has proven to be one of the more successful tools in studying these systems on account of its unique access to the electronic states through their hyperfine interactions with the nuclear spins. However, NMR measurements of strongly correlated electrons systems can exhibit apparently anomalous signatures which defy easy explanation [4]. Through the use of numerical simulations of strongly correlated electrons in the context of NMR spin echo experiments, we can develop an understanding of how these unique signatures can be indicative of the presence of strongly correlated electrons, and how we can use the presence of these features to extract key information about the structure of the electronic state.

1.1 NMR and Strongly Correlated Electrons

Nuclear Magnetic Resonance (NMR) has been, and continues to be, a powerful tool in the study of strongly correlated electrons. NMR has a unique ability to probe the properties of the electrons in these materials on account of its access to their states via the hyperfine interactions between the electrons and nuclei. Because the characteristic frequencies of the nuclei are much smaller than those of the electrons, NMR offers a probe of the electronic states which minimally disturbs the state itself. The central measurements of NMR are sensitive to to the spin and orbital states of the electrons, and as such NMR measurements are capable of extracting information about the electronic spin susceptibility and giving valuable insight into the electronic structure. Given its strengths in studying the properties of electrons in these systems, it is unsurprising that NMR has been widely used in identifying and qualifying the features of unconventional superconductors.

One of the most powerful measurements for the analysis of the electronic spin states in NMR is the Knight shift. The Knight shift refers to an overall shift in the nuclear frequency spectrum caused by the magnetization of conduction electrons and driven by the hyperfine interactions between the electrons and nuclei [5]. Measurement of the Knight shift, dependent as it is on the electronic spin susceptibility, is a common tool to study strongly correlated materials. The manner in which the Knight shift changes relative to temperature through a superconducting transition can yield valuable information about the electronic spin susceptibility. The emergence of standard BCS type superconductivity results in a sharp drop and eventual complete disappearance of the Knight shift as $T \to 0$ from the disappearance of the electronic spin susceptibility in the superconducting state. However, in exotic superconductors, such as triplet superconductors, the triplet spin state of the Cooper pairs results in non-zero spin susceptibility below the superconducting transition and a persistence of finite Knight shift at $T \to 0$. The nature of the lingering Knight shift can also yield

information about the type of triplet state present in the material [2]. The Knight shift can additionally yield information about the presence of the FFLO state, as the formation of the FFLO state relies on the polarizability of the superconducting electrons and consequently requires different behavior in the spin susceptibility and Knight shift than a standard BCS superconducting state [3]. In this manner, by examining the change in the Knight shift across a superconducting transition and the emergence of apparently anomalous changes in the Knight shift, one can both probe the nature of the electronic spin state and detect the emergence of unusual strongly correlated modes.

Although much of our upcoming discussion focuses on changes in the Knight shift, it is not the only NMR tool relevant to the detection of strongly correlated electrons. Another powerful NMR measurement for the detection of exotic electronic correlations is the spin lattice relaxation time T_1. Dependent as it is on fluctuations in the local magnetic field driven by electronic behavior, T_1 is sensitive to the emergence of strongly correlated phenomena, and its measurement can be a further indicator of exotic dynamics. Spin lattice relaxation, characterized by T_1, is a dissipative process driven by the magnetic field generated by the electron's polarizability in the presence of the Zeeman field. It will therefore be related to the dynamic susceptibility and can serve as another probe of the electronic spin state. Additionally, changes in the electronic state such as in the emergence of superconductivity can drive unusual phenomena in measurements of T_1. In typical BCS superconductors, T_1 exhibits a "coherence peak" below the critical temperature T_c [6]. A simple calculation for a metal can relate $1/T_1$ to the density of states at the Fermi energy [7]. In a BCS superconductor, a gap opens in the density of states at the Fermi energy and the density of states increases dramatically at the edge of the gap. At the transition temperature, this sudden increase in the density of states results in the value of $1/T_1$ suddenly increasing, creating the coherence peak. This feature is explicitly tied to

the predictions of BCS theory, and divergence from this behavior can indicate of the presence of unconventional superconductivity.

Ever since the discovery of the superconductor Sr_2RuO_4 and the first suggestion that it may host triplet superconductivity, it has garnered a large amount of attention [8, 9]. Theoretical justifications have been developed for the presence of triplet superconductivity in Sr_2RuO_4 [10, 11], and a large amount of experimental research has attempted to identify the presence of this exotic type of superconductivity in Sr_2RuO_4, to varying degrees of success [12–16]. NMR measurements in particular of the Knight shift and spin lattice relaxation T_1 have already been widely used in the search for triplet superconductivity in Sr_2RuO_4. NMR studies of both Ru and ^{17}O in Sr_2RuO_4 have shown a constant Knight shift across the superconducting phase transition, the lack of a T_1 coherence peak below T_c, and temperature dependence characteristic of unconventional triplet superconductivity [17–19]. However, recent ^{17}O NMR examinations contest the claim of triplet superconductivity, demonstrating that the constancy of the Knight shift through the phase transition arises from heating from the RF pulses causing an unintentional suppression of the superconducting state and the dominance of the normal state down to low temperature [20, 21]. These contradictory results demonstrate the pulse-dependence of the Knight shift, showing that the higher power pulses, which result in Knight shifts akin to the normal state, induce heating in the sample and destroy the superconducting state. The revelation of the effect of heating on these samples additionally demonstrates the difficulty in measuring these exotic phases, and the ease at which small experimental factors can be destructive to them in ways that are difficult to identify.

Much like Sr_2RuO_4, the heavy fermion superconductor $CeCoIn_5$ has received a lot of attention for its potential hosting of the FFLO state [22–25]. NMR measurements, in particular measurements of the Knight shift, have also had a fair degree of use in the search for experimental evidence of the elusive FFLO state. The behavior of the

Knight shift has allowed the identification of higher order symmetry of the superconducting order parameter indicating the presence of unconventional superconductivity, agreeing with evidence from other experimental techniques [26–28]. Observations of the behavior of the shift as a function of temperature and field have been used to chart the phase diagram of $CeCoIn_5$; these measurements show strong evidence of two phase transitions, the second of which may be indicative of the emergence of the FFLO state [29]. Measurements of the spin susceptibility via the Knight shift indicate a notable departure from normal superconducting behavior [30]. However, the evidence for the presence of the FFLO state in $CeCoIn_5$ is not yet conclusive, indicating the difficulty in probing strongly correlated materials [31]. Additionally, measurements of delicate and exotic superconducting states such as the FFLO state are susceptible to heating effects in the same manner we saw with Sr_2RuO_4, inducing an emergence of the normal phase that can be misinterpreted as the presence of a new phase [32–36].

We have seen several cases in which strong electronic correlations give way to unique and complex physics. The emergent phenomena driven by these strong electronic correlations can be diverse and unpredictable, and they can give way to exotic phases which defy easy examination. In the specific case of unconventional superconductivity, understanding the complicated electronic dynamics and the bonding mechanism that drives the formation of Cooper pairs in these systems is crucial to understanding the emergence of superconductivity outside the BCS paradigm. The two cases we have already discussed, triplet superconductivity and the FFLO state, both host exotic anisotropic superconducting gaps (one in momentum space, one in real space), and understanding the roll the complex electronic dynamics play in driving the emergence of these exotic phases would not only help elucidate the various mechanisms by which superconductivity can arise, but also help chart out the wider landscape of the dynamics of strongly correlated electrons. However, merely measur-

ing the presence of these exotic phases, much less extracting information about the dynamics of the electrons within, has proven extremely difficult. Not only are the phases rare, they are sensitive to a wide range of potential experimental failures, and they can be highly unstable relative to temperature and magnetic field and prone to breakdown.

Our discussion has already laid the groundwork for why NMR can be a potent tool for the study of systems with strongly correlated electrons. Its unique approach to probing the electronic structure and access to information about the spin state of the electrons through the dependence of the Knight shift and T_1 on the static and dynamic electronic susceptibility respectively makes it an attractive method for the examination of these systems. However, the wide variety of potential behaviors emerging from the presence of strongly correlated electrons make experimental design difficult, and the myriad ways in which the experimental approach itself can interact with these novel phases can lead to "false positives", as we saw with Sr_2RuO_4. The extreme delicacy with which one must treat these systems makes experimental verification difficult, but it also opens the door to the implementation of novel techniques currently under-utilized or under-developed.

A particularly under-developed field in the study of NMR is the use of simulation to understand the dynamics of large, interacting spin ensembles. A number of conventional and widely used packages exist, but their focus is generally unsuited for problems in solid state physics, especially dynamic simulations. Many NMR simulation packages focus on chemical or biological applications [37–40]. Several prioritize spectral simulations or focus on the effect of quadrupole interactions rather than the types of exotic, long range correlations we aim to study [41–43]. Some focus on either powder or liquid state NMR, and can only accomodate systems of spins far too small for our purposes [44–47]. Those simulations which are more broadly applicable or geared towards dynamic simulations are programmed using software such as

Mathematica or languages like Python and are too inefficient for the ensemble sizes we require [48–51]. Other simulation packages in NMR are typically out-dated or no longer being developed [52, 53].

NMR simulations of strongly correlated electrons and long range dynamics represent a difficult field of numerics given the necessity for large ensembles in order to capture a sufficient sampling of the frequency environments which drive the core measurements of NMR. To this end, very few simulations exist which are suited for the study of strongly coupled nuclear spins. However, there exists immense potential for the development of dynamic simulations in NMR by taking advantage of several key tools. Basic models, such those using the mean field approximation, can provide great insight into the dynamics of systems despite their simplicity. Additionally, by reducing the size of the Hamiltonian with a mean field model or another dimensionality-reducing technique, we can take advantage of the high degree of optimization offered by massively parallel computing using GPUs.

GPU programming is known to have a large amount of potential for the development of modern physical simulation [54]. It has primarily seen application to Monte Carlo methods, which are unsuited for dynamic simulations [55–59]. GPU programming has seen some application topics in spin dynamics, but these applications are not applicable for NMR modelling or focus on molecular dynamics or semi-classical approaches [60–63]. Application of GPU programming to the simulation of dynamic single spin systems in the quantum picture with time dependent Hamiltonians– the type of system we are concerned with– has only very recently begun to be developed [64]. However, these simulations come with several limitations, in particular a limited scope to specifically unitary dynamics as well as an inability to model systems where the time dependence of the Hamiltonian is not known in advance, such as systems whose interactions induce a self-dependent back-action. Additionally, these simulations rely on a high degree of analytic calculation to facilitate a high degree

of efficiency and accuracy. While such calculations can be very useful within their scope, they inherently limit the application of the simulation a very specific class of problem. A more numerically driven approach, which can be more widely applied within its realm of focus which also takes advantage of GPU tools, remains a relatively unexplored regime.

In the development of our program, we aimed to fill this gap in the current landscape of dynamic spin simulations. We set out to develop a efficient program which could be applied to a wide range of problems within in the dynamics of single-spin systems or effectively isolated spins (i.e. those whose interactions can be treated on the mean field level). Our intent was to make a simulation which could model spins in the presence of non-unitary dynamics, such as from thermal effects, and could also be adjusted to be applied to different Hamiltonians without an excess of re-programming. To these ends, we primarily focused on numerical techniques over analytic tools in the key calculations of the system, such as the matrix exponential. We also prioritized efficiency, using a large degree of GPU programming and ultimately writing the simulation directly in CUDA to minimize computation time. In this manner, we aimed to create a simulation that could be applied to a large variety of problems in NMR and could be quickly adapted and re-used to study new interactions.

By developing a custom simulation package that takes full advantage of the accelerated computation possible with GPU integration in tandem with a theoretically simple yet robust model of the complicated electronic dynamics, we can explore the possibilities for how unforeseen or potentially misunderstood signatures in NMR studies of these exotic systems can be indicative of long range electronic correlation. With these new tools, combined with a knowledge of NMR experimental technique, we can develop new interpretations of and expectations for the ways strongly correlated electrons can exhibit their unique dynamics in the context of the spin echo. With these new interpretations, we can propose novel experimental routines specifically designed

to probe the behavior of the nuclear spins interacting in the presence of strongly correlated electrons and identify key features about the structure of the electronic state.

1.2 Overview

NMR has been used widely in the examination of strongly correlated materials. Measurements of these strongly correlated phases bring with them a variety of difficulties partially due to the strict limits on the stability and presence of these novel phases. In these situations, the emergence of unusual signatures in these materials is not uncommon. We have already seen how the emergence of pulse-dependent Knight shifts can occur in unconventional superconductors, possibly on account of unintentional heating effects in the case of Sr_2RuO_4, but also in other potentially unconventional superconductors such as FeSe [4]. However, it is not unusual to observe asymmetric and multi-peak behavior in the time domain, as well. Even more so than unusual features in the spectral domain, time domain asymmetry is typically attributed to experimental failure. We propose that both of these behaviors, both temporal and spectral, rather than being indicative of experimental failure or unintended phase transitions, may be indicative of the presence of a strongly correlated electron mode inducing unique dynamics in the nuclear spin ensemble.

To demonstrate the possibility of a new understanding of the role strongly correlated electrons play in inducing anomalous signatures in NMR measurements, we develop a mean field model of highly coupled nuclear spins and examine its dynamics using a custom simulation package. We propose a mechanism of coupling between nuclear spins enhanced by the presence of long range correlations between electrons. We lay out a simple model of such an interaction in the mean field limit with the goal of isolating the ways such an interaction exhibits itself in the spin echo signal,

both in the time and frequency domains. In order to study this model efficiently, we develop a custom simulation which takes advantage of GPU programming to rapidly simulate large ensembles of spins.

Using this simple model and our custom simulation, we then explore the dynamics of the nuclear ensemble and demonstrate how such interactions can result in the exact pulse-dependent Knight shift behavior observed in FeSe and other systems in which anomalous Knight shifts have been measured in controlled environments where the potential for heating induced effects is significantly reduced [4]. Furthermore, through a careful analysis of the Knight shift as well as several experimental techniques, we outline how information about the anisotropy of the electronic spin susceptibility and relative length scale of the electronic correlations can be extracted from the anomalous shifts. With these results, we demonstrate how the use of even simple models in tandem with efficient numerical simulations represents a new tool with which the dynamics of strongly correlated materials can be studied and understood.

We begin our discussion with a review of basic NMR theory. We introduce the fundamentals of the Knight shift, spin lattice relaxation T_1, and how they relate to the electronic structure. We introduce the spin echo protocol and how the dephasing of the spins from Larmor precession can be reversed via the use of RF pulses. We next turn our attention to the specifics of our model, introducing a Hamiltonian based on a theory of an effective interaction between the nuclear spins mediated by hyperfine interactions with the electrons and amplified by the strong correlations between the electrons. We review the specific methods used to calculate the dynamics of the spins, discussing key topics such as the rotating frame and approximations used in time evolving the system.

With the physical background of the model established, we introduce the simulation packages developed for the calculation of the spin echo. We review the structure of the simulation, the different programming languages used, and we examine several

key algorithms and their accuracy. We review the performance of the simulation and its different versions.

After discussing the structure of the simulation and its numerical accuracy, we turn to the accuracy of our model. We perform several validating tests to ensure the observed behavior is consistent across different simulation versions and not sensitive to differences in specific algorithms, precision types, or other numerical differences. We examine the convergence of the model with regards to the temporal resolution and ensemble size in order to chart out appropriate regimes of accuracy with regards to these key artificial parameters.

Following a discussion of the theoretical and computational foundations of our model, we examine the dynamics of the nuclear ensemble under the influence of our interaction. We discuss the explicitly the spin dynamics and how they give way to the unique temporal and spectral features observed. We examine key features of both the time domain signals and the spectra observed under the influence of our model across a wide range of experimental and theoretical parameters in the global interaction limit before discussing how the introduction of localization affects these observed behaviors.

Finally, we discuss how these types of simulations can be used to interpret experimental results. We demonstrate how we can extract information about the relative strength and anisotropy of the electronic susceptibility from anomalous shifts in the spectra of strongly correlated materials, and how the use of simulations can aid in the understanding and identification of unique phases observed using NMR.

Chapter 2

Theoretical Background

Before we proceed to a discussion of the simulation, we discuss the fundamentals of nuclear magnetic resonance within the context of our exploration. NMR is a rich field with a wide variety of phenomena to examine. In the following discussion, however, we focus primarily on those key elements which are relevant to our model. We introduce the fundamental mechanics of NMR and the structure and origin of the spin echo protocol. To give context to our previous discussion, we briefly examine the origins of the Knight shift and spin lattice relaxation and how they are tied to the electronic structure. We then turn to the specifics of our model, its Hamiltonian, and the specifics of how we approximate the various features of a spin echo experiment.

2.1 NMR Basics

NMR is primarily concerned with measuring the splitting of the nuclear energy levels in the presence of a magnetic field. This splitting is typically induced with the use of an applied external magnetic field B_0, although in some cases interactions in the system are sufficient to induce an observable splitting of the nuclear spin energy levels (i.e. nuclear quadrupole resonance) and no additional field is necessary. NQR does not apply to spin-$1/2$ particles (the focus of our simulations) so we do not discuss it

in depth and rather direct the interested reader to a relevant source [65]. For our discussion, we assume the presence of an externally applied field B_0.

In the presence of B_0, the energy levels of the nuclei split according to spin from the effect of Zeeman splitting

$$H = -\gamma \mathbf{I} \cdot \mathbf{B}_0 \tag{2.1}$$

typically, B_0 is treated as being aligned with the z axis. The energy levels split according to the spin state in the z basis, with the energy gap between levels given by $\Delta = \gamma \hbar B_0$. In NMR, this is typically written in terms of the *Larmor frequency* $\Delta = \hbar\nu$ for $\nu = \gamma B_0$ the Larmor frequency. This frequency, aside from denoting the spacing of the levels, additionally sets the frequency at which a planar spin state, for spin-1/2 either $|\pm x\rangle$ or $|\pm y\rangle$, precesses about the z axis. Consider the following spin-1/2 state

$$|\psi\rangle = \frac{1}{\sqrt{2}}|+\rangle + \frac{e^{i\varphi}}{\sqrt{2}}|-\rangle \tag{2.2}$$

for $|+\rangle, |-\rangle$ the z basis states which denotes a state fully "in the plane", i.e. with $\langle I_z \rangle = 0$. Under the action of the Zeeman Hamiltonian, this state evolves as

$$|\psi(t)\rangle = e^{i\nu t/2}\left[\frac{1}{\sqrt{2}}|+\rangle + \frac{e^{-i(\nu t+\varphi)}}{\sqrt{2}}|-\rangle\right] \tag{2.3}$$

We can quickly calculate the expectation values to demonstrate the oscillatory effect

$$\langle I_x \rangle = \frac{1}{2}\cos(\nu t + \varphi) \qquad \langle I_y \rangle = -\frac{1}{2}\sin(\nu t + \varphi) \tag{2.4}$$

In a real material, each spin will experience a slightly different local field on account of dipole-dipole interactions with other nuclear spins. The effect of these dipole-dipole

interactions between the nuclear spins is to create a distribution of energy gaps about the unperturbed average $\Delta = \gamma\hbar B_0$. This distribution is often roughly Lorentzian in shape.

Although in a strict sense NMR measures the distribution of energy gaps in a material, it is done so by taking advantage of the way these energy differences set the precession frequencies of each nuclear spin. Consider an ensemble of spin-$1/2$ nuclei all of which, at $t = 0$, are aligned with the x axis (i.e. $\langle I_x \rangle = 1/2$). The coherence of the spins creates a net magnetization in the sample which can be measured (we omit many of the experimental details of this measurement in order to focus our discussion on the theoretical elements). As the spins precess about the z axis, all with slightly different precession frequencies, the overall coherence dissipates and the spins dephase.

In a simplified sense, we can treat the signal $S(t)$ as simply being equal to the transverse magnetization, which is specified by the expectation value of the raising operator $I_+ = I_x + iI_y$ averaged over the spin ensemble. Using Equation 2.4, we can see the signal is given by

$$S(t) = \frac{1}{2} \int d\nu \, p(\nu) e^{-i\nu t} \tag{2.5}$$

where $p(\nu)$ represents the amplitude of the spectrum at that frequency ν. In this form, we can see the signal is clearly the Fourier transform of the frequency spectrum and results in exponential decay of the signal with characteristic time set by the inverse of the spectral width. This exponentially decaying signal is known as the *Free Induction Decay* or just the *FID*. The inverse is also true: the frequency spectrum is the Fourier transform of the observed signal. This process is the fundamental technique of modern NMR. In the most simple terms, it involves measuring the net magnetization of the sample allowed to evolve under an externally applied field B_0 and applying a Fourier

transformation to the time domain signal to extract the distribution of frequencies– and thus the distribution of energy level gaps– of the system.

2.1.1 The RF Pulse

In the previous section, we treated the spins as being aligned with the plane and explored the resulting dynamics. However, a fundamental element of NMR is the tool by which the spins are aligned in the plane: the Radio Frequency (RF) pulse, named for the range of frequencies typical of systems studied with NMR. In their ground state, the spins favor the $+z$ state on account of its lower energy from the Zeeman splitting. In a real system, the population difference is small and can be calculated using basic statistical mechanics, but for simplicity we shall assume all spins are in the ground state.

We now consider the evolution of a spin-$1/2$ particle in the $+z$ state under the action of an oscillating magnetic field applied orthogonal to the external field B_0. Let the oscillating field be circularly polarized (in a real system, the oscillating field is only applied along a single axis, but that system is not analytically solvable). The Hamiltonian for this spin is

$$H = -\nu I_z - \omega \left(I_x \cos \nu' t - I_y \sin \nu' t \right) \tag{2.6}$$

with $\omega = \gamma B$ representing the strength of the oscillating field. In matrix form, the Schrödinger equation can be written

$$\begin{pmatrix} \dot{a} \\ \dot{b} \end{pmatrix} = \frac{i}{2} \begin{pmatrix} \nu & \omega e^{i\nu' t} \\ \omega e^{-i\nu' t} & -\nu \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \tag{2.7}$$

for a spin $|\psi\rangle = a\,|+\rangle + b\,|-\rangle$. This relation gives us two coupled differential equations

$$\begin{aligned} \dot{a} &= \frac{i}{2}\left(\nu a + \omega e^{i\nu' t} b\right) \\ \dot{b} &= \frac{i}{2}\left(\omega e^{-i\nu' t} a - \nu b\right) \end{aligned} \tag{2.8}$$

Here, we can apply a change of variables $a \to \alpha = e^{-i\nu' t/2} a$ and $b \to \beta = e^{i\nu' t/2} b$. With these substitutions, our differential equations become

$$\begin{aligned} \dot{\alpha} &= \frac{i}{2}(\nu - \nu')\alpha + \frac{i}{2}\omega\beta \\ \dot{\beta} &= -\frac{i}{2}(\nu - \nu')\beta + \frac{i}{2}\omega\alpha \end{aligned} \tag{2.9}$$

Solving this differential equation is somewhat tedious, but not difficult. We can substitute in an ansatz

$$\begin{aligned} \alpha &= Ae^{i\mu t} + Be^{-i\mu t} \\ \beta &= Ce^{i\mu t} + De^{-i\mu t} \end{aligned} \tag{2.10}$$

for some frequency μ. We can find a condition on the value of μ by matching the coefficients of terms with the same exponential dependence, as the equations must be equal for any time t. With this process, we find

$$\mu = \frac{\sqrt{\omega^2 + (\nu - \nu')^2}}{2} \tag{2.11}$$

where we take the positive root since the two roots will be equivalent when applied to the solution. We can find the values of the coefficients A, B, C, and D by examining initial conditions. In this manner, we find

$$\begin{aligned} a(t) &= e^{i\nu' t/2}\left[a_0 \cos\mu t + \frac{i}{2\mu}\left[\omega b_0 + (\nu - \nu')\, a_0\right] \sin\mu t\right] \\ b(t) &= e^{-i\nu' t/2}\left[b_0 \cos\mu t + \frac{i}{2\mu}\left[\omega a_0 - (\nu - \nu')\, b_0\right] \sin\mu t\right] \end{aligned} \tag{2.12}$$

In the specific case of a spin initially in the $+z$ state with on-resonance pulsing ($\nu' = \nu$)

where $\mu = \omega/2$ we have

$$\begin{aligned} a(t) &= e^{i\nu t/2} \cos \omega t/2 \\ b(t) &= e^{-i\nu t/2} \sin \omega t/2 \end{aligned} \tag{2.13}$$

Applying the oscillating field for a fixed duration is called a *pulse*. From Equation 2.13, we can see that applying a pulse of length $t = \pi/2\omega$ will put the spin in a fully planar state (i.e. one where $|a| = |b|$ and $\langle I_z \rangle = 0$). This pulse is called the $\pi/2$ pulse and is typically the first step of an NMR measurement. Once the spins have been flipped into the plane, they precess and the spectrum can be measured as we discussed in Section 2.1.

Clearly, spins that are off-resonance ($\nu \neq \nu'$) will not end up fully flipped into the plane. However, in NMR the Larmor frequency ν is typically on the order of tens of MHz and the variation ν' is typically on the order of a few kHz. The variation from a perfect 90° flip is therefore typically very small.

2.1.2 The π Pulse and the Spin Echo

In some situations, a $\pi/2$ pulse followed by measurement of the FID is all that is required to extract the desired information about the system. However, for other measurements, such as relaxation times, or for situations where experimental factors such as ringing (a large amount of noise in the measurement near the pulse from acoustic feedback in the circuit) can make the FID either insufficient or too noisy to record. These issues can be solved with a different, two pulse protocol known as the *spin echo*.

Consider an ensemble of $s = 1/2$ spins initially in the $|+x\rangle$ state which have been allowed to freely precess for a time τ. For a spin with precession frequency equal to

the Larmor frequency, its state, up to an overall phase, will be

$$|\psi(\tau)\rangle = \frac{1}{\sqrt{2}}\left(|+\rangle + e^{-i\nu\tau}\,|\text{-}\rangle\right) \tag{2.14}$$

Using the result of Equation 2.12, if we apply an oscillating field to this spin, it will evolve as

$$|\psi(t+\tau)\rangle = \frac{1}{\sqrt{2}}\left[\left(\cos{}^{\omega t}\!/_{2} + ie^{-i\nu\tau}\sin{}^{\omega t}\!/_{2}\right)|+\rangle + \left(e^{-i\nu\tau}\cos{}^{\omega t}\!/_{2} + i\sin{}^{\omega t}\!/_{2}\right)|+\rangle\right] \tag{2.15}$$

If we apply the pulse for a duration $t = {}^{\pi}\!/_{\omega}$, or precisely twice the length of the ${}^{\pi}\!/_{2}$ pulse, phase factor switches to the $+z$ coefficient from the $-z$ coefficient. Extracting a total overall phase, we can write the new state as

$$|\psi(\tau + {}^{\pi}\!/_{\omega})\rangle = \frac{1}{\sqrt{2}}\left(|+\rangle + e^{i\nu\tau}\,|\text{-}\rangle\right) \tag{2.16}$$

If the spin is once again allowed to freely precess, the precession will now reduce the phase shift from the $|+x\rangle$ state, and at a time $t = 2\tau$ it will return to its initial state. Although spins with precession frequencies slightly off-resonance will not perfectly re-focus, the amount by which they remain dephased will scale with the value of $\nu - \nu'$ relative ν, and we have already established that ν is typically two to three orders of magnitude larger than $\nu - \nu'$. For the purposes of our discussion, we can ignore this variation.

This "re-focusing" effect of a pulse of length $t = {}^{\pi}\!/_{\omega}$, known as the π pulse, is the source of the spin echo. After the π pulse and after a time τ equal to the amount of time the spins originally were allowed to dephase, all spins return to the initial planar state they were in at the time of the ${}^{\pi}\!/_{2}$ pulse. This process results in a "second FID", mirrored about $t = 2\tau$. This peak is known as the *spin echo*.

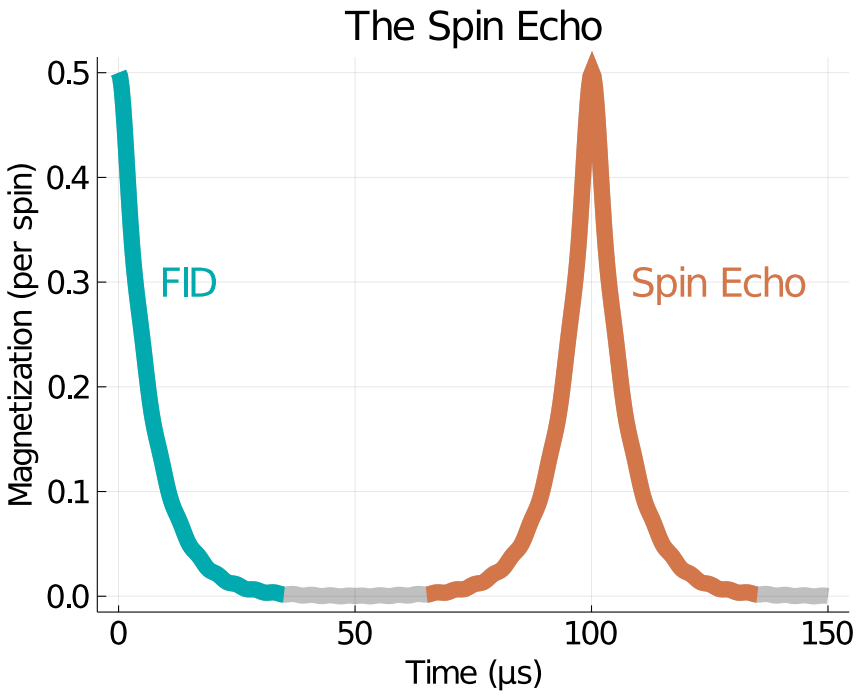

Figure 2.1: A basic simulated spin echo with $\tau = 50\mu s$, $\nu = 10$MHz, and a spectral width of $\Gamma = 50$kHz.

All together, the spin echo protocol consists of several steps. First, a $\pi/2$ pulse to flip the spins into the plane. This initial pulse is followed by a delay of length τ during which the spins de-phase and the observable signal decays exponentially. At $t = \tau$, the π pulse is applied, reversing the spins' phases and causing them to precess towards their initial state. Shortly before $t = 2\tau$, the signal exponentially grows to its initial value before exponentially decaying once again, creating the spin echo. The spin echo is one of the fundamental measurements in NMR, and it has served as the foundation for a huge number of complicated and nuanced pulse sequences which are capable of probing and refocusing a wide number of phenomena. For our purposes, we focus the spin echo in its most basic form: a $\pi/2$ pulse followed by free precession for a time τ, with a single re-focusing π pulse and an echo at $t = 2\tau$.

The pulses are often intuitively thought of as applying a rotation to the nuclear spins. This geometric interpretation is exact in the case of spin-$1/2$, but for higher spins it can still be intuitively useful. A common approximation of the pulse is therefore

to simply apply a rotation to the spins, as the dynamics during the pulse itself are often unimportant except from a technical perspective. To this end when we discuss pulsing, we typically discuss it in terms of flip angles and axis of rotation rather than the pulse duration or polarization.

Given that the echo comprises the same frequency behavior as the FID, the frequency spectrum of the nuclear system can be extracted from the spin echo in the same manner. Additionally, the echo's amplitude can be used to extract information about relaxation rates, which we discuss shortly.

2.1.3 The Knight Shift

In a real system, the main interest of NMR is not the distribution of energy levels from the Zeeman splitting, but the ways in which it diverges from this, either through overall shifts in the mean frequency, splitting of the peak, or broadening of the spectral width. The most relevant effect for our discussion, and the one we already introduced in Chapter 1, is the Knight shift K.

The Knight shift, originally identified by Walter Knight in 1949, describes an overall shift in the frequency from the hyperfine interaction of the nuclei with electrons [5]. The interaction between the nuclear and electronic spins introduces an additional term to the Hamiltonian

$$H_{hf} = \frac{\mu_0 \gamma \gamma_e \hbar^2}{4\pi} \sum_i \left[\frac{\mathbf{I} \cdot \mathbf{S}_i}{r_i^3} - \frac{3(\mathbf{I} \cdot \mathbf{r}_i)(\mathbf{S}_i \cdot \mathbf{r}_i)}{r_i^5} - \frac{8\pi}{3} \mathbf{I} \cdot \mathbf{S}_i \, \delta(\mathbf{r}_i) \right] \tag{2.17}$$

where we ignore the electron orbital contribution, since its effects are well known and do not encode physics relevant to our discussion [66]. In practice, these hyperfine interactions introduce an effective field generated by the electron moments which acts on the nuclei in addition to the externally applied field, resulting in an overall shift of the spectrum.

If we treat the effect of the hyperfine interaction as simply applying an effective field whose strength is proportional to the average magnetization of the electrons, i.e. $\langle \bar{\mathbf{S}} \rangle$, then we can see immediately how the Knight shift becomes a measure of the spin susceptibility. Because of its dependence on the polarization of the electronic spins, its magnitude will be proportional to the response of the electrons in a static field, that is the static susceptibility $\chi(\mathbf{q} = 0, \omega = 0)$, and its measurement can be used as a probe of the electronic state. In this way, information about the real part of the susceptibility, χ' can be obtained from the Knight shift.

2.1.4 Relaxation

As we have already mentioned, the spin echo is also used to extract information about relaxation rates of the nuclear spin ensemble. Although we do not examine relaxation effects in our exploration, it is informative to explore how they can be used to measure information about the electronic state in order to contextualize our discussion in Chapter 1. Although there are two primary relaxation times measured in NMR, T_1 and T_2, we are primarily interested in T_1.

T_2 is often called *spin-spin relaxation* driven by interactions between nuclear spins which introduces an irreversible decoherence of the spins. This decoherence cannot be refocused using a spin echo and typically emerges as an attenuation of the echo's amplitude. In many solid state systems, T_2 is on the order of tens or hundreds of microseconds. In liquid state NMR, very long coherence times are possible, on the order of seconds.

More relevant to our motivating discussion is T_1, known as *spin-lattice relaxation*. Spin lattice relaxation arises from interactions between the nuclei and electrons which drives transitions between the nuclear energy levels and ultimately causes the ensemble to relax toward thermal equilibrium. Spin lattice relaxation is a dissipative process requiring the absorption of energy to drive transitions between energy levels.

We therefore would expect T_1 to be related to the imaginary part of the dynamic susceptibility in the static field limit i.e. $\omega \to 0$. An explicit relationship of the spin-lattice relaxation time T_1 to the electron susceptibility is far more difficult to demonstrate than with the Knight shift. That said, T_1 can be explicitly related to the dynamic susceptibility via

$$\frac{1}{T_1 T} \propto \lim_{\omega_0 \to 0} \sum_{\mathbf{q}} F(\mathbf{q}) \frac{\chi''(\mathbf{q}, \omega_0)}{\hbar \omega_0} \tag{2.18}$$

for χ'' the imaginary part of the susceptibility [66, 67]. This explicit relation makes clear the dependence of T_1 on the susceptibility and its usefulness in studies of electronic structure via NMR.

Perhaps more instructive, but less applicable, is a relationship between the Knight shift and T_1 known as the *Korringa relation*. The Korringa relation, although it only applies in relatively ideal circumstances, makes clear the related nature of T_1 and the Knight shift via their shared dependence on the electronic structure. The Korringa relation is given by [68]

$$T_1 T K^2 = \frac{\hbar}{4\pi k_B} \left(\frac{\gamma_e}{\gamma} \right)^2 \tag{2.19}$$

2.2 Model

Our model seeks to emulate the interactions between the nuclear spins mediated by the local electron magnetization via hyperfine interactions. The interaction between the nuclear spins from hyperfine interactions with local electrons will scale as the square of the hyperfine interaction strength and will typically be very weak. In the presence of a collective excitation, such as in a system with strongly correlated electrons, the strong coherence of the electrons can magnify the strength of this indirect interaction and result in noticeable changes in the dynamics of the system.

With this model, we primarily aim to study relatively long range interactions between nuclear spins driven by strong electronic correlations. In such situations, with long range interactions and ensembles with a large number of spins, the influence of highly local forces, such as J coupling or dipolar couplings, which typically extend only to nearest and next-nearest neighbors, become less relevant. The effects of such interactions are well studied and well understood in NMR, whereas our exploration focuses on the effect of the proposed long range correlations and novel methods of their detection. To that end, to isolate the effect of long rage interactions between the spins, we do not include these terms in the Hamiltonian.

In order to reduce computational complexity and efficiently model a large ensemble of spins, we use a mean field approach to model the interaction between nuclear spins. Given the nature of the phenomena we aim to model, we must simulate a large ensemble of spins, making the use of more exact approaches nearly impossible. However, our model focuses typically on longer range correlations, resulting in a large degree of averaging and strengthening the mean field approximation. Additionally, the introduction of the localization factor allows us to still consider the influence of spatial variations to a degree by applying the mean field approximation locally rather than globally.

In the mean field limit, the electronic degrees of freedom are integrated out of the model to generate an effective interaction between nuclear spins, mediated by hyperfine interactions between nuclei and electrons. This effective interaction on the nuclear spin i can be written as

$$H_E^{(i)} = -\gamma \mathbf{I}^{(i)} \cdot \mathbf{M}^{(i)} \tag{2.20}$$

where $\mathbf{M}^{(i)}$ is the effective magnetic field felt at site i due to nearby nuclear spins. The magnitude of M depends on both the strength of the coupling between the electrons

and the nuclei and the distance between the nuclei. We write the effective field $\mathbf{M}^{(i)}$ as a sum over couplings to nearby nuclei, j,

$$M_d^{(i)} = \alpha_d \sum_{j \neq i} f(|\mathbf{r}_i - \mathbf{r}_j|) \langle I_d^{(j)} \rangle \quad \text{for } d = x, y, z \tag{2.21}$$

This form of $\mathbf{M}^{(i)}$ displays how the electronic action facilitates a back-action of the nuclei on themselves; the magnitude of the action on a nuclear spin i depends on the states of nearby spins. The full Hamiltonian for a given spin i, after including the Zeeman term, becomes

$$H^{(i)} = -\nu_i I_z^{(i)} - \gamma \sum_{d=x,y,z} \alpha_i I_d^{(i)} M_d^{(i)} \tag{2.22}$$

with $M_d^{(i)}$ as defined in Equation 2.21.

Our model has several tuneable parameters that dictate the behavior of the system: the axis-dependent spin coupling strengths α_d and the details of the localization function f, which can vary in both form and degree of localization. The parameter α_d represents primarily the strength of the initial coupling of the electron mode to the nuclei. The localization function f captures the spatial extent of the nuclear interaction, and is generally related to the structure of the electronic susceptibility. The combination of the two defines the overall "strength" of the interaction.

The localization function f can be set to have any structure necessary to mimic the systems being studied. In our discussion, we focus on a Gaussian localization scheme where

$$f(|\mathbf{r}_i - \mathbf{r}_j|) \equiv f(r, \xi) = e^{-(r/\xi)^2} \tag{2.23}$$

The Gaussian localization aims to mimic an electronic excitation with a gap; other potential localization schemes, such as a power law $f(r, p) = 1/r^p$ or an RKKY type

$f(r, \xi) = x^4(x \cos x - \sin x)$ for $x = 2\xi/r$ could be used to model systems such as one without a gap, or in a simple metal respectively [69–71].

2.2.1 The Nuclear Ensemble

In order to study the general features that can occur due to our proposed spin-spin interaction, we restrict ourselves to a simple system: a two dimensional square lattice of spin-1/2 nuclei with a periodic boundary condition. Using a mean field model, as opposed to a product space representation for the quantum states, allows us to study large ensembles of spins, as the number of degrees of freedom scale as $\mathcal{O}(N)$ rather than $\mathcal{O}(x^N)$. Restricting ourselves to two dimensions allows us to increase the length-scale of the lattice without excessive computation time, allowing us to study both short-range and long-range interactions. We generally ignore temperature dependent effects and initialize this system in a ground state with all spins pointing along the z axis, although temperature dependent effects can be reintroduced using the Lindblad terms, as we discuss in a later section.

Due to small variations in the local nuclear environment, in real materials the spins have slightly different precision times under an applied magnetic field. This observation, that the spectrum of the spins is not constrained to a single frequency, is the key insight in the development and implementation of the spin echo protocol. We use a Lorentzian distribution for the spectrum centered at 10MHz, a gyromagnetic ratio $\gamma = 1\text{MHz/T}$ and a field of $B_0 = 10\text{T}$. Although we vary the linewidth at times, the dynamics generally scale in a predictable way with the linewidth. For most of our experiments, we therefore set the linewidth of the spectrum at 50kHz, which is typical for high-quality crystalline samples at low temperatures. In assigning frequencies to each spin, we randomly sample a range of frequencies from 9.75MHz to 10.25MHz, with probability weighted by the Lorentzian spectrum. Although we use instantaneous pulses, as described below, in real experiments the duration of the

pulse sets a limitation on the bandwidth, motivating our cutoff of the spectrum over a window of 0.5MHz, corresponding to $\pi/2$ pulses roughly of length $0.5\mu s$.

2.2.2 Pulses and Time Evolution

We implement our RF pulses by assuming instantaneous pulses. The time scale of the experiment is on the order of $100\mu s$, whereas the pulses occur on the order of $1\mu s$. The influence of the interaction on the ensemble during the pulse is therefore negligible. To implement the pulses, we simply apply rotation matrices to the density matrix of each spin

$$\rho_f = e^{iR\varphi/\hbar}\rho_0\, e^{-iR\varphi/\hbar} \tag{2.24}$$

where R is the appropriate spin matrix for the axis of rotation, typically I_x or I_y, and φ the flip angle, ρ_0 is the spinor's density matrix before the pulse and ρ_f its density matrix after. When we introduce pulsing with $\varphi \neq \pi/2$, we also introduce basic phase cycling to eliminate the lingering FID after the π pulse. We use a basic $(0,1)-(0,3)$ phase cycling, with the $\pi/2$ pulse along x and the π pulse along $\pm y$.

We model a simple spin echo experiment that consists of a $\pi/2$ and a π pulse separated by echo time τ. We apply both pulses along the x axis, resulting in a signal predominantly aligned with the y axis. Although we will refer to these two pulses as $\pi/2$ and π pulses, we do not restrict ourselves to spin rotation angles of strictly $\pi/2$ and π, nor necessarily in the ratio $1:2$.

The time evolution of the density matrix is dictated by

$$\frac{d\rho}{dt} = \frac{i}{\hbar}[\rho, H] \tag{2.25}$$

Because of the Zeeman term in the Hamiltonian, the spins will rotate rapidly about the z axis. This rapid oscillation can be eliminated by working in a rotating reference

frame in which the x and y axes rotate about the z axis at the central Larmor frequency. We can define the density matrix in this reference frame as

$$\rho_R = U^\dagger \rho U \longrightarrow \rho = U \rho_R U^\dagger \tag{2.26}$$

with $U = \exp(-i\nu I_z t/\hbar)$. Making this replacement, we find

$$\frac{d\rho_R}{dt} = \frac{i}{\hbar}[\rho_R, U^\dagger H U + \nu I_z] \tag{2.27}$$

We recover an "effective" Hamiltonian in the rotating frame, H_R such that

$$\frac{d\rho_R}{dt} = \frac{i}{\hbar}[\rho_R, H_R] \quad \text{with} \quad H_R = U^\dagger H U + \nu I_z \tag{2.28}$$

Using this transformation, our Hamiltonian from Equation 2.22 becomes

$$H_R^{(i)} = -(\nu_i - \nu) I_z^{(i)} - \gamma \sum_{d=x,y,z} \alpha_d U^\dagger I_d^{(i)} U M_d^{(i)} \tag{2.29}$$

The density matrix ρ_R still undergoes Larmor precession, but at a frequency $\Delta\nu \equiv \nu_i - \nu$, which is typically on the order of kHz rather than MHz, resulting in a period on the same order as the experiment time, around $100\mu s$. By eliminating the high-frequency behavior induced by the Larmor precession, we drastically reduce the temporal resolution required in our calculation.

2.2.3 Interactions in the Rotating Frame

The transition to the rotating frame may induce high-frequency behavior in the interaction term. To perform our time evolution, we propagate by small time steps dt, using average Hamiltonian theory to create an effective Hamiltonian on each time step. As we shall see, in the rotating frame the interaction Hamiltonian oscillates

rapidly (on the order of the Larmor frequency) about a slowly varying average value. If we choose our time steps dt to be sufficiently small such that the slowly varying magnitude of the interaction can be treated as constant on any given time step, we can use average Hamiltonian theory to eliminate the high frequency terms entirely. By showing the high frequency terms disappear in the effective Hamiltonian, we will be free to set the time step much larger than would otherwise be required by the terms which oscillate at frequencies near the Larmor frequency. In order to characterize the "slowly varying" and "rapidly varying" parts of the interaction and set limits on our choice of dt, we must carefully examine how the interaction term behaves in the rotating frame.

The case $d = z$ in Equation 2.29 is uninteresting, as I_z commutes with U. In this case, the validity of our approximation depends simply on the value of α_z. For $d = x, y$, U has the effect of rotating the spin operator about the z axis by an amount νt:

$$\begin{aligned} U^\dagger I_x U &= I_x \cos \nu t - I_y \sin \nu t \\ U^\dagger I_y U &= I_y \cos \nu t + I_x \sin \nu t \end{aligned} \tag{2.30}$$

By taking advantage of the cyclic property of the trace, $M_d^{(i)}$ can also be rewritten relative to the rotating density matrix ρ_R:

$$M_d^{(i)} = \sum_{j \neq i} f(r_{ij}) \langle U^\dagger I_d^{(j)} U \rangle_R \tag{2.31}$$

where $r_{ij} \equiv |\mathbf{r}_i - \mathbf{r}_j|$ and $\langle I_d^{(j)} \rangle_R$ indicates the expectation value taken with respect to $\rho_R^{(j)}$. For simplicity, moving forward we define $\langle - \rangle \equiv \langle - \rangle_R$, unless otherwise noted. By combining the expressions in Equations 2.30 and 2.31 and using some basic algebraic manipulation, we can rearrange the expression $\alpha_x U^\dagger I_x^{(i)} U M_x^{(i)} + \alpha_y U^\dagger I_y^{(i)} U M_y^{(i)}$, separating the oscillatory (those with an explicit $\sin n\nu t$ or $\cos n\nu t$) and non-oscillatory

terms. With these adjustments, our Hamiltonian in the rotating frame becomes

$$\begin{aligned} H_R^{(i)} = &- (\Delta\nu + a_z^{(i)}) I_z^{(i)} \\ &- (a_{x+}^{(i)} + a_{x-}^{(i)} \cos 2\nu t - a_{y-}^{(i)} \sin 2\nu t) I_x^{(i)} \\ &- (a_{y+}^{(i)} - a_{y-}^{(i)} \cos 2\nu t - a_{x-}^{(i)} \sin 2\nu t) I_y^{(i)} \end{aligned} \tag{2.32}$$

with coefficients a_i defined by

$$\begin{aligned} a_z^{(i)} &= \gamma \alpha_z \sum_{j \neq i} f(r_{ij}) \langle I_z^{(j)} \rangle \\ a_{d\pm}^{(i)} &= \frac{\gamma}{2} (\alpha_x \pm \alpha_y) \sum_{j \neq i} f(r_{ij}) \langle I_d^{(j)} \rangle \end{aligned} \tag{2.33}$$

We can make several substitutions here which will more strongly relate the different terms to the physical quantities of interest. We introduce an anisotropy factor η defined by

$$\eta \equiv \frac{\alpha_x - \alpha_y}{\alpha} \tag{2.34}$$

for $\alpha \equiv \alpha_x + \alpha_y$ as well as terms $\omega_d^{(i)}$ defined by

$$\omega_d^{(i)} \equiv \sum_{j \neq i} f(r_{ij}) \langle I_d^{(j)} \rangle \tag{2.35}$$

Finally, we define $\omega_\perp^{(i)} \equiv \omega_x^{(i)} + i\omega_y^{(i)}$. With these substitutions, the Hamiltonian becomes

$$\begin{aligned} H_R^{(i)} = &- \left(\Delta\nu + a_z^{(i)}\right) I_z^{(i)} \\ &- \frac{\gamma\alpha}{4} \left(\omega_\perp^{(i)} \left(1 + \eta e^{2i\nu t}\right) + \bar{\omega}_\perp^{(i)} \left(1 + \eta e^{-2i\nu t}\right) \right) I_x^{(i)} \\ &- \frac{i\gamma\alpha}{4} \left(\bar{\omega}_\perp^{(i)} \left(1 - \eta e^{-2i\nu t}\right) - \omega_\perp^{(i)} \left(1 - \eta e^{2i\nu t}\right) \right) I_y^{(i)} \end{aligned} \tag{2.36}$$

2.2.4 Magnus Expansion

The Hamiltonian in the rotating frame, $H_R^{(i)}$, only has explicit time dependence in the case that $\eta = 0$, the case of isotropic planar coupling. In the case of anisotropic planar coupling, where $\eta \neq 0$ and the Hamiltonian has explicit time dependence, we can approximate the dynamics using the Magnus expansion [72].

The Magnus expansion has a wide array of applications in solid state physics and NMR and is a powerful tool to examine the dynamics of time dependent systems [73]. However, in applying the Magnus expansion, we have to take several precautions. Average Hamiltonian Theory is known to fail under certain circumstances, most notably systems acted upon by oscillating fields [74–79]. It can be shown that even for a simple time-independent system, use of the interaction representation, combined with an average Hamiltonian approach wherein the integration period is exactly equal to the period of induced time dependence introduced with the interaction picture, can cause the average Hamiltonian theory to fail [74]. Accordingly, in systems with periodic time dependence, such as ours, great care must be taken in using an average Hamiltonian approach over the period of the time dependence.

The benefits of using such an approach can still be significant, as the time dependence will integrate out exactly in the first order "average" term, and the complexity of higher order terms is vastly reduced. In order to mitigate the possibility for such a transformation to produce incorrect results, we take several precautions. Firstly, we perform much of our analysis using an interaction which is isotropic in the plane, i.e. $\alpha_x = \alpha_y$. In this case, the time dependence is identically zero and we have no need to use the Magnus expansion at all. Secondly, in application of the Magnus expansion, we use numerical techniques to examine its accuracy by comparing it to a direct implementation of the full, time dependent Hamiltonian with a sampling sufficiently small to capture the high frequency dynamics of the system.

Because of the self-referential nature of the Hamiltonian introduced through the ω_d

terms, the Hamiltonian cannot be explicitly integrated with respect to time. However, we can approach this issue in two ways: we could use a perturbative approach, using a non-interacting echo or an echo with isotropic coupling to approximate the values of ω_d and integrate the resulting expression, or we could approximate the ω_d as constant on the interval of integration. We use the latter approach, as our time steps will be sufficiently small such that this approximation holds.

Since we are time evolving the density matrices and not state vectors, to properly use the Magnus expansion we should elevate the system to Liouville space. Ultimately, we do proceed in Liouville space; however, since the commutation relations of the spin matrices are equivalent in Hilbert and Liouville space, we can calculate the coefficients of the Magnus expansion in the lower dimensional Hilbert space and elevate the terms to Liouville space after the fact.

In the our technical discussions, we typically separate the examination of the impacts of α_z and α_x, α_y, as they interact with the system in fundamentally different ways. For simplicity, given our current focus on the impact of planar anisotropy, we proceed with $\alpha_z = 0$. The α_z terms can be reintroduced by simple replacement of $\Delta\nu \to \Delta\nu + a_z^{(i)}$. In order to examine how the anisotropy of the coupling affects the system, we derive the second and third order terms in the Magnus expansion for this case. When introducing dissipation, the terms become more complicated, given the case of dissipation requires full use of the Liouville space representation. However, a Hilbert space approach still gives accurate results for those terms which do not depend on the dissipation, so we approach the impact of the dissipation terms separately.

Given the oscillatory nature of the time dependence, proper selection of the size of the time step dt (which is also the period T over which the Magnus expansion is evaluated) can drastically simplify the resulting expressions. With this in mind, we use time steps $dt = T = n\pi/\nu$ which equates to integrating over one full period of the time dependent terms.

The Magnus expansion gives an approximate solution to the equation

$$\frac{dv}{dt} = Av \tag{2.37}$$

for some matrix A and vector v. The solution is given in series

$$v(T) = e^{\bar{A}(T)}v(0) \tag{2.38}$$

We have $v \to \rho$ and $A \to {}^{i\hat{L}}/_{\hbar}$ for the Liouvillian $\hat{L}$. Given the lack of dissipation, we can perform the calculation with H instead of $\hat{L}$, being careful to restore one factor of ${}^{i}/_{\hbar}$ for every copy of H that appears. The additional factors ${}^{1}/_{\hbar}$ will precisely be cancelled by factors $\hbar$ gained from the commutation of spin operators. For simplicity, moving forward we suppress factors (i) in the factors ω, with the understanding that each is dependent on a reference particle.

The first order term in the Magnus expansion is given by

$$H^{(i)}_{R,1} = \frac{i}{\hbar}\int_0^T H^{(i)}_R(t)dt \tag{2.39}$$

When integrating over one period, this clearly simply integrates out the time dependence. Integrating this term from 0 to $T = n\pi/\nu$, we find

$$\begin{gathered} H^{(i)}_{R,1} = a^{(i)}_z I_z + a^{(i)}_x I_x + a^{(i)}_y I_y \\ a^{(i)}_z = -\frac{iT}{\hbar}\Delta\nu\,, \quad a^{(i)}_d = -\frac{iT}{\hbar}\left(\frac{\gamma\alpha\omega_d}{2}\right) \end{gathered} \tag{2.40}$$

where we immediately note the lack of terms scaling with η. To first order in the Magnus expansion, assuming the bounds are chosen to coincide with the periods of the Larmor precession, the dependence on the anisotropy disappears. This features is representative of the fact that the anisotropy's contribution to the dynamics of the system is, on average, zero in the approximation that ω_d is constant. Contributions

from the anisotropy could be included in the first order term by choosing bounds that do not coincide with the period of the Larmor precession, but doing so does little except to drastically increase the complexity of the higher order terms. A different choice of bounds could help safeguard against the issues with the average Hamiltonian approach, but we choose to use a numerical approach to check for divergent behavior rather than use the more cumbersome bounds.

The second order term can be found by integrating the commutator

$$H_{R,2}^{(i)} = -\frac{1}{2\hbar^2}\int_0^T dt \int_0^t dt' \left[H_R^{(i)}(t), H_R^{(i)}(t')\right] \tag{2.41}$$

which, when integrated, gives

$$\begin{aligned} H_{R,2}^{(i)} &= b_z^{(i)} I_z^{(i)} + b_x^{(i)} I_x^{(i)} - b_y^{(i)} I_y^{(i)} \\ b_z^{(i)} &= \frac{iT}{\hbar}\left[\frac{(\gamma\alpha)^2\eta\left(\omega_\perp^2 - \eta\omega_\perp\bar{\omega}_\perp + \bar{\omega}_\perp^2\right)}{16\nu}\right] \\ b_d^{(i)} &= \frac{iT}{\hbar}\left[\frac{\eta\gamma\alpha\omega_d\Delta\nu}{4\nu}\right] \end{aligned} \tag{2.42}$$

We note a few significant features of the second order term. Aside from the scaling with η, the second order terms introduce a mixing of the interaction terms. The planar corrections scale not only with α and ω_d, but additionally scale with the Larmor shift $\Delta\nu$. Conversely, the out-of-plane coupling scales with the planar variables α, η, and $\omega_\perp$. More importantly, the overall scale of the second order terms relative to the first is set by the ratio of the interaction strength with the Larmor frequency. Given the interaction strength is limited by the linewidth Γ, the scaling of the higher order terms in the Magnus expansion can be said to be, at best, Γ/ν, which is typically on the order of $1/100$ or less. The impact of these terms will therefore be minimal, if at all visible. Given the timescale of the interaction is on the order of kHz and the Larmor frequency is on the order of MHz, the minimal impact of the planar anisotropy is expected. The spins precess far too quickly for any difference in the x

and y couplings to accumulate and cause any real difference in the dynamics.

To confirm the scaling of the Magnus expansion, we can additionally calculate the third order term. The third order term is found by integrating the double commutators

$$\begin{aligned} H_{R,3}^{(i)} &= \frac{-i}{6\hbar^3}\int_0^T dt\int_0^t dt'\int_0^{t'} dt''\Big[C(t,t',t'')+C(t'',t',t)\Big] \\ C(t_1,t_2,t_3) &\equiv [H_R^{(i)}(t_1),[H_R^{(i)}(t_2),H_R^{(i)}(t_3)]] \end{aligned} \tag{2.43}$$

This integration yields the third and final term of our analysis

$$\begin{aligned} H_{R,3}^{(i)} &= c_z^{(i)}I_z^{(i)}+c_x^{(i)}I_x^{(i)}-c_y^{(i)}I_y^{(i)} \\ c_z^{(i)} &= \frac{iT}{\hbar}\left[-\frac{(\gamma\alpha)^2\Delta\nu(\omega_\perp^2+\bar{\omega}_\perp^2-2\eta\omega_\perp\bar{\omega}_\perp)}{32\nu^2}\right] \\ c_x^{(i)} &= \frac{iT}{\hbar}\left[\frac{\gamma\alpha\eta(16\Delta\nu^2-2(\gamma\alpha)^2P_\omega+(\gamma\alpha)^2\eta Q_\omega)}{128\nu^2}\right] \\ c_y^{(i)} &= \frac{iT}{\hbar}\left[\frac{\gamma\alpha\eta(16\Delta\nu^2+2(\gamma\alpha)^2P_\omega+(\gamma\alpha)^2\eta Q_\omega)}{128\nu^2}\right] \\ P_\omega &\equiv \bar{\omega}_\perp^2-\omega_\perp^2-\eta\omega_\perp\bar{\omega}_\perp\,, \quad Q_\omega \equiv \omega_\perp^2-2\eta\omega_\perp\bar{\omega}_\perp \end{aligned} \tag{2.44}$$

The third order terms are further reduced by a factor proportional to the interaction relative to the Larmor frequency, as we saw with the second order term.

The corrections from the second and third order terms in the Magnus expansion only apply for planar anisotropic interactions, when $\alpha_x \neq \alpha_y$. Given we expect the impact of the planar anisotropy to be negligible, we focus much of our analysis on isotropic planar interactions, where $\alpha_x = \alpha_y$, in which case the corrections from the second and third order terms in the Magnus expansion do not apply. This approach is also equivalent to a first order Magnus approximation. We will return to the impact of these corrections later in our discussion.

2.2.5 Effective Time Scale Ω

The Magnus expansion is only valid if the norm of the Hamiltonian is much smaller than the time interval over which the expansion is calculated. One such norm is given by [80]

$$\Omega = \left[\mathrm{Tr} \left[H_R^{(i)} \right]^2 \right]^{1/2} \ll 1/T \tag{2.45}$$

This condition can also be used as a limit on the validity of our assumption that the coefficients $\omega_d^{(d)}$ are constant on the period T of the Magnus expansion, which is also the time step dt. For an isotropic planar interaction, we find

$$\Omega = \frac{1}{\sqrt{2}} \sqrt{(\Delta\nu + \gamma\alpha_z\omega_z^{(i)})^2 + (\gamma\alpha|\omega_\perp|/2)^2} \tag{2.46}$$

In order to put an upper bound on the value of Ω, we assume the local spin averages $\omega_d^{(i)}$ to be constant and equal to their maximal value, where all expectation values are equal to the spin s. Additionally, each individual particle will have a different value of the parameter $\Delta\nu$, so to create a parameter that captures the dynamics of the system as a whole, we substitute $\Delta\nu$ with an "average" value, i.e. the linewidth Γ. We additionally define quantities ω and ω_z which we call the "interaction weight" defined by

$$\begin{aligned} \omega &\equiv (\alpha_x + \alpha_y) \sum_{j \neq 1} f(r_{1j}) \\ \omega_z &\equiv \alpha_z \sum_{j \neq 1} f(r_{1j}) \end{aligned} \tag{2.47}$$

which represents a "maximal", spin independent value of the $\omega_d^{(i)}$ parameters introduced previously, scaled by α. With these substitutions, we arrive at a parameter

Ω

$$\Omega = \sqrt{(\Gamma+\gamma s\omega_z)^2/2 + (\gamma s\omega)^2/8} \tag{2.48}$$

which defines the overall "frequency" of the interaction, when interpreted as a rotation (for $\alpha \equiv \alpha_x + \alpha_y$ and $\alpha_x = \alpha_y$). With this definition, the condition for our approximation of the time evolution to be valid is given by

$$\Omega \ll 1/T \tag{2.49}$$

where $T \equiv dt$ is our time step and the period of the Magnus expansion. In the case of an isotropic planar interaction, our Hamiltonian reduces to one nearly identical to the original Hamiltonian in the laboratory frame, but with a shifted Larmor frequency:

$$H_R^{(i)} = -\left(\nu_i - \nu\right) I_z^{(i)} - \gamma \sum_{d=x,y,z} \alpha_d M_{d,L}^{(i)} I_d^{(i)} \tag{2.50}$$

where we require $\alpha_x = \alpha_y$ and we have defined the "local magnetization"

$$M_{d,L}^{(i)} = \sum_{j\neq i} f(r_{ij}) \langle I_d^{(j)} \rangle_R \quad \text{for } d = x, y, z, + \tag{2.51}$$

where $d = +$ corresponds to the raising operator I_+, which corresponds to the complex valued planar magnetization. This parameter is equivalent to our parameters $\omega_d^{(i)}$ introduced earlier. We see that in the rotating frame, for an isotropic interaction, the behavior of the x and y axes is not truly distinct. The action of these two terms in the model combine to form one effective term with strength additive of the individual strengths α_x, α_y. For the specific case of a spin 1/2 particle, for which we perform our

analysis, this Hamiltonian in matrix form is

$$H_R^{(i)} = \frac{\hbar}{2}\begin{pmatrix} -\Delta\nu - \gamma\alpha_z M_{z,L}^{(i)} & -\frac{\gamma}{2}\alpha M_{+,L}^{(i)} \\ -\frac{\gamma}{2}\alpha\overline{M}_{+,L}^{(i)} & \Delta\nu + \gamma\alpha_z M_{z,L}^{(i)} \end{pmatrix} \tag{2.52}$$

for $\overline{M}$ indicating the complex conjugate and $\alpha \equiv \alpha_x + \alpha_y$.

2.2.6 Liouville Space

We model our system in Liouville space, which allows us to introduce dissipative terms to the system. In Liouville space, the $d \times d$ density matrix is flattened into a $d^2 \times 1$ vector; in our simulation, we use row major order to make this conversion. On account of this convention, it can quickly be shown that left and right multiplication by some operator $\mathcal{O}$ are written

$$\begin{aligned} \mathcal{O}\rho &\longrightarrow (\mathcal{O} \otimes \mathbb{I})\rho_L \\ \rho\mathcal{O} &\longrightarrow (\mathbb{I} \otimes \mathcal{O}^T)\rho_L \end{aligned} \tag{2.53}$$

With this convention in mind, the time evolution equation for our density matrix becomes

$$\begin{gathered} \frac{d\rho_R^{(i)}}{dt} = \frac{i}{\hbar}[\rho^{(i)}, H_R^{(i)}] \longrightarrow \frac{d\rho_{R,L}^{(i)}}{dt} = \frac{i}{\hbar}\hat{L}_R^{(i)}\rho_{R,L}^{(i)} \\ \text{with} \\ \hat{L}_R^{(i)} = \mathbb{I} \otimes \overline{H}_R^{(i)} - H_R^{(i)} \otimes \mathbb{I} \end{gathered} \tag{2.54}$$

For a time independent Hamiltonian, the above equations have the solution

$$\rho_{R,L}^{(i)}(t) = e^{i\hat{L}_R^{(i)}t/\hbar}\rho_{R,L}^{(i)}(0) \tag{2.55}$$

Our Hamiltonian, which depends on the expectation values of the spins in the ensemble, is not time independent. However, as previously established, we can approximate

the Hamiltonian as time independent on time steps dt so long as the condition in Equation 2.49 is met. Explicitly, we time evolve our density matrices in steps dt

$$\rho_{R,L}^{(i)}(t+dt) \approx e^{i\hat{L}_R^{(i)}(t)\,dt/\hbar}\rho_{R,L}^{(i)}(t) \tag{2.56}$$

In an experimental setting typical of the systems we are modelling, $T_1 \gg \tau$ and can be generally be ignored for a single spin echo measurement. However, it is not uncommon that $T_2 \sim \tau$. To that end, we include dissipative effects in our simulation. For general dissipation, the time evolution of the density matrix is dictated by

$$\frac{d\rho}{dt} = \frac{i}{\hbar}[\rho, H] - \frac{1}{2}\sum_n \left(L_n^\dagger L_n \rho + \rho L_n^\dagger L_n\right) + \sum_n L_n \rho L_n^\dagger \tag{2.57}$$

where L_n are the Lindblad jump operators [81]. Although we are primarily concerned with the dephasing of spins via T_2 processes (i.e. $L_n = I_z$), we also include terms that account for emission ($L_n = I_-$), absorption ($L_n = I_+$), which allow us to include T_1 and thermal effects in the simulation if we desire. Taking into account our row-major ordering and the self-adjoint nature of the spin operators, the inclusion of dissipation results in an additional term

$$\hat{J} = \sum_{d=z,+,-} \frac{\Gamma_d}{\hbar^2}\left(-\frac{1}{2}\left[I_d^2 \otimes \mathbb{I} + \mathbb{I} \otimes \overline{I}_d^2\right] + I_d \otimes \overline{I}_d\right) \tag{2.58}$$

With this term included, we arrive at the final equation which dictates time evolution in our simulation:

$$\rho_{R,L}^{(i)}(t+dt) = e^{\left(i\hat{L}_R^{(i)}(t)/\hbar + \hat{J}\right)dt}\rho_{R,L}^{(i)}(t) \tag{2.59}$$

With the dissipation included, we have to re-examine the Magnus expansion. The terms calculated in the Hilbert space representation will still be valid, but there will

be additional terms which represent the mixing of the dissipation and the interaction. As before, the first order term is simply the time-averaged term and is equivalent to the isotropic case. The second order term is identical to the second order term from the Hilbert space representation, plus an additional term that accounts for the mixing of the dissipation. This additional term has factors that look like

$$\frac{\gamma\alpha\omega\eta}{\nu}\sum_{i=1}^{3}\beta_i\Gamma_i \tag{2.60}$$

where the β_i are either ± 1 or ± 3 depending on the term and i. We see the exact same scaling as in the previous examination, as well as a mixing of the dissipation and interaction terms. Given the dramatic increase in complexity of the calculation in Liouville space, as well as the extreme scaling of the terms, we only include the first and second order contributions from dissipation. With the form of the Hamiltonian given, the third order term can be calculated relatively quickly in software which is capable of symbolic integration, such as MATLAB [82]. However, it is cumbersome to the point of being impractical to include, and its contribution will be multiple orders of magnitude smaller than even the second order term.

Although our discussion primarily focuses on non-dissipative cases of the interaction, we still perform our calculations in Liouville space. We do this for several reasons. Primarily, the introduction of dissipative effects such as from phase decoherence can help frame the dynamics of the model in a setting more indicative of an actual experiment. As we shall see Chapter 4, certain emergent phenomena driven by the interaction rely on a high degree of coherence of the spin ensemble and can be sensitive to dissipative effects such as would be present in a real system from relaxation and thermal effects. Future applications of the simulation and similar models would likely focus more on the effects of dissipation in the system. Additionally, because the simulation is primarily coded in CUDA, much of the structure has to be hard coded.

It behooves us to structure the majority of the simulation in Liouville space from the beginning, rather than attempt to add it in later, as the Liouville representation can be used to model dynamics that can be modelled in Hilbert space, but the inverse is not true.

Chapter 3

Simulation

Numerical simulations are a relatively common tool for the study of nuclear magnetic resonance. There are several widely-used simulation packages available (e.g. SPINEVOLUTION or SIMPSON) which are designed to be flexible tools for the simulation of NMR systems [44–47]. Many of the simulations packages emphasize generality and accessibility, intending to present to the user the ability to simulate a wide range of NMR systems with minimal need for direct programming.

The aforementioned simulation packages primarily focus on powdered or liquid state NMR; additionally, the packages limit the user in terms of the available interactions that can be included, typically only allowing the tuning of the parameters of several standard NMR terms rather than the introduction of novel interactions between the spins. When it comes to the simulation of single-crystal solid state systems, especially those which explore the effect of novel interactions, the landscape of available simulations becomes much more barren.

Direct simulations of single-crystal solid state systems can rapidly become intractable, as the introduction of interactions between the spins results in 2^N scaling which rapidly exceeds current computational capability for anything more than a handful of spins. Even in the case of a mean field model, which eliminates the 2^N

scaling, computation times can rapidly become excessive for large ensembles of spins.

Recently, strides have been made in the efficient simulation of dynamic, time-dependent spin dynamics [64]. While the simulation is highly efficient and very accurate, and represents a major development in the use of GPU programming for spin dynamics, it is ultimately poorly suited for many NMR applications and in particular is fundamentally unable to model non-unitary dynamics or Hamiltonians whose time dependence is not known in advance.

One of the primary goals of our work is the development of a simulation that can accurately and efficiently study large ensembles of spins with novel interactions. The program focuses on the simulation of large ensembles of spins whose interactions can be modelled on the mean field level. The particular form of the interaction is not essential to the broad structure of the simulation; however, the interaction Hamiltonian is hard coded to optimize performance, and changing the interaction on a structural level requires directly changing the simulation code, a task that is more tedious than difficult. We additionally take advantage of GPU computing– a tool not widely seen in the world of time-domain NMR simulations– to accelerate the program, allowing for highly efficient simulations of large ensembles of spins.

In the process of developing the simulation package and studying the model, we wrote three versions of a spin echo simulation [83]. The first iteration, written entirely in Julia and utilizing no GPU programming (referred to moving forward as simply *the Julia simulation*), was primarily used in our initial exploratory investigations of the model; it is essentially a direct translation of the model described above, and it primarily serves as a tool to quantify the performance of more optimized programs [84]. We primarily use the other two simulations for our experiments, one written in Julia with some GPU functionality (which we call the *GPU-Julia simulation*) and one written in CUDA (the *CUDA simulation*) [85–87]. In the following sections, we discuss these two simulations in some depth, focusing on a few key calculations, their

accuracy, and performance.

Although both the GPU-Julia and CUDA simulations model the same spin echo, each programming language offers unique benefits and complications. We primarily use the Julia code to prototype new features, although it also serves as a metric against which to compare the performance and accuracy of the CUDA simulation. Once a new feature or model is tested using the GPU-Julia simulation, it can be added to the more efficient (but more unwieldy) CUDA simulation for the generation of large data sets.

Julia's primary benefit is its accessibility. It offers a robust and optimized library of scientific functions; given its similarity in style to Python, translating the mathematics of the model into a working simulation is direct and relatively simple. These factors make it easy to initially create the simulation as well as to add new features to the simulation or to adjust existing elements of the code.

The main cost of Julia is its performance. Julia primarily uses CPU computing, with some GPU compatibility, and options for parallelization are limited. Although many of the available libraries in Julia are well optimized for CPU computing, the structure of our simulation is not. Our model is a mean field model, which practically translates to a large number of small, independent calculations. Parallelizing over a small number of powerful CPU cores offers negligible gain in terms of computation time, as the time cost of moving memory between cores is not recouped by the $4\times$, $6\times$, or even $8\times$ parallelization offered by a CPU. Julia's GPU libraries are not complete enough to remedy this problem entirely, but they do offer a significant increase in performance (around $5\times$) for the most costly step of the simulation. Although Julia's CUDA library does offer tools for the programming of custom CUDA functions in Julia, it offers little to no benefit for our simulation over using CUDA natively in C++.

To study the spin-interaction model in great depth requires larger data sets and

a more optimized simulation. In terms of performance, very few languages perform GPU parallelization better than properly implemented CUDA. With CUDA, we can explicitly choose how to split up a function across the thousands of available threads on the GPU. For a given particle, each step of our simulation is simple, requiring little more than 4×4 matrix operations. With complete control over the parallelization structure, we can execute the simulation almost entirely in parallel, with each thread performing the calculation for one of the spins. Although the individual cores are much slower than those on a CPU, the massive parallelization results in speed-ups of 10-20$\times$ compared to our GPU-implementation in Julia.

This high performance comes at a cost. CUDA and C++ offer only extremely limited libraries for the common operations necessary for scientific computation. In particular, there are almost no linear algebra libraries that are well suited to our simulation. CUDA does have a BLAS-based linear algebra library, but it is designed for operations on large matrices and is not well suited for the small matrices of our model. Almost every function in our CUDA simulation is therefore adjusted to be suitable for our model, which makes creating the simulation and adding new features a much slower process than in Julia. Of particular note is the lack of a matrix exponential, as even CUDA's cuBLAS library does not have a matrix exponential function. A robust discussion of the matrix exponential algorithm is beyond the scope of this paper, and the implementation of such an algorithm represents a major undertaking in the development of a CUDA-based simulation [88].

In the following sections, we review the basic structure of the Julia and CUDA simulations, some of the more unique elements of their implementation, and their performance. We review the two non-trivial calculations involved in the simulations, namely the calculation of $M_{d,L}^{(i)}$ and the calculation of the matrix exponential in the CUDA simulation. We then examine the performance of the three simulations (Julia, GPU-Julia, and CUDA) and examine the accuracy of the handwritten algorithms in

the CUDA simulation by comparing its results to those from the GPU-Julia simulation.

3.1 Simulation Structure

Both simulations implement the same model and calculate the output of a basic spin-echo experiment. The calculated signal and output of the simulations is the net magnetization along each axis, M_d, given by

$$M_d = \sum_{i=0}^{N} \langle I_d^{(i)} \rangle_R \quad \text{for } i = z, + \tag{3.1}$$

where we save the value of the "rotating" expectation value, as defined in Section 2.2.3. This value is equivalent to the laboratory frame value for $d = z$, and for $d = +$ it differs by a simple phase factor $e^{-i\nu t}$. Laboratory frame values for M_x and M_y can be recovered by $M_x = \mathrm{Re}\,(e^{-i\nu t} M_+)$ and $M_y = \mathrm{Im}\,(e^{-i\nu t} M_+)$.

To perform the simulation, we first generate and initalize an ensemble of density matrices as well as a number of variables that are held constant throughout the experiment before walking through the basic steps of the spin echo: we apply an RF pulse to flip the spins an angle θ (i.e. a rotation of each density matrix by an angle θ), time evolve under our model for a time τ, apply a second RF pulse to flip the spins by an angle φ, and then time evolve for a time 2τ, saving the values of M_z, M_+ at each time step and after each pulse.

Although in principle every variable in the system can be adjusted freely, the simulations are designed to accommodate easy adjustment of the most commonly changing variables: the interaction strengths α, the length scaling function f, the dissipation coefficients Γ_d, the lattice dimensions n_x, n_y, the echo time τ and the time step dt, and the pulse angles θ and φ. Adjusting other variables in the system typically requires hard-coding the new values or re-programming some elements of

the simulation to more seamlessly loop over the new variables, a task that is more tedious than difficult.

The "free" parameters for a given set of trials (those mentioned previously over which the program is designed to iterate) are loaded from a parameter file (a simple text file) at the start of an experiment. In the GPU-Julia simulation, the parameters are loaded from a single parameter file, and the simulation performs a simple loop over each set of parameters. In the CUDA simulation, the parameters are loaded from a pair of parameter files: one denoted "simulation" parameters and one denoted "echo" parameters. Simulation parameters are those upon which memory allocation depends, such as n_x or τ, and echo parameters are those upon which it does not. In the CUDA simulation, the loop over trials is performed with a nested loop, where the outer loop iterates over the simulation parameters and the inner loop over echo parameters. This structure allows us to perform the memory allocation once per set of echo parameters by reusing and overwriting memory from completed trials, minimizing the time spent on memory management.

The key optimization of the GPU-Julia and CUDA simulations is the massive parallelization possible when using a mean-field model. Each of the thousands of thread on a GPU is typically quite weak computationally compared to one of the cores of a CPU, and consequently they favor small, simple calculations. This preference results in GPU computing being highly efficient at processes which can be divided into a large number of very simple, independent calculations.

In the GPU-Julia simulation, we only have limited control over how the calculation is parallelized on the GPU. However, Julia's CUDA package has highly-optimized parallelization schemes for many basic matrix operations built-in. Without direct control of the thread assignments, we can optimize the parallelization of those calculations performed on the GPU by taking advantage of these built-in optimized operations. By structuring a given calculation in such a way that it can be written as a single

operation on a matrix, or between several matrices, we can ensure that the calculation is parallelized optimally. This re-structuring comes at a cost, however, as we will see later.

In CUDA, we have direct and complete control over how a calculation is divided up across the threads on the GPU. Given our model consists of thousands of independent manipulations of small matrices (typically two or four dimensional), we parallelize the system by particle: we assign each particle to a single thread which performs the full calculation for said particle for a single time step.

In the next sections, we review some of the important calculations performed in the simulations which are typically both non-trivial and relevant to the performance of the simulation. For the Julia and GPU-Julia simulations, the only calculation of interest is the local magnetization $M_{d,L}^{(i)}$. For the CUDA simulation, we additionally must discuss the calculation of the matrix exponential, as our simulation uses a hand-written version of the scaling-and-squaring algorithm for its calculation. In the following sections, we discuss these calculations, how they vary across simulation versions, and the accuracy of the hand-written exponentiation algorithm compared to the built-in versions available in Julia.

3.1.1 Local Magnetization: $M_{d,L}^{(ij)}$

The calculation of the local magnetization uses an $n_x \times n_y$ matrix S known as a "stencil", where n_x and n_y are the lengths of the lattice along the x and y axis, respectively. Moving forward, we shift from "particle number" indexing to a system by which we index particles based on their x and y coordinates. $M_{d,L}^{(i)}$ becomes $M_{d,L}^{(ij)}$. Our previous notation $f(r_{ij})$ becomes

$$f(r_{ij}) \to f(r_{ij,k\ell}) \tag{3.2}$$

where ij, $k\ell$ refer to each particle's x and y coordinates. The stencil S is defined by

$$S_{ij} \equiv f(r_{00,ij}) \tag{3.3}$$

with indexing now starting at zero. Additionally, $S_{00} \equiv 0$, a condition which prevents self-coupling. The stencil S is, by default, "centered" at $(0, 0)$. By taking advantage of our periodic boundary conditions, we can shift the stencil to be centered at any given point (k, ℓ). Letting $S^{(k\ell)}$ designate a stencil centered at (k, ℓ), we find the following transformation between S and $S^{(k\ell)}$:

$$\begin{aligned} S_{ij}^{(k\ell)} &= S_{i'j'} \\ i' &= (i-k) \bmod n_x \\ j' &= (j-\ell) \bmod n_y \end{aligned} \tag{3.4}$$

where we define the modulo of a negative value periodically with

$$(-a) \bmod n \equiv n - a \tag{3.5}$$

Additionally, we use an $n_x \times n_y$ matrix M_d^{eval} which contains the expectation value of I_d of each particle on the lattice. We define M_d^{eval} by

$$\left(M_d^{eval}\right)_{ij} \equiv \langle I_d^{(ij)} \rangle_R \tag{3.6}$$

where the indices (i, j) refer to a given particle by its x, y coordinates as before. With these definitions in place, we can see that the definition of $M_{d,L}^{(ij)}$ as given in Equation 2.51 is equivalent to the sum of the element-wise multiplication of M_d^{eval} and $S^{(ij)}$.

The fundamental process used to compute $M_{d,L}^{(ij)}$ is the same across all three versions of the simulation (Julia, GPU-Julia, and CUDA), although the specifics of the implementation vary. In all three cases, we calculate the value of S a single time at

the start of a trial, as the calculation represents one of the slower processes in the simulation. In the Julia and GPU-Julia simulations, we also calculate the full set of shifted stencils $S^{(ij)}$ for all particles prior to performing the time evolution. The set of shifted stencils is saved to a four dimensional array $\hat{S}$, defined by

$$\hat{S}_{ijk\ell} = S^{(ij)}_{k\ell} \tag{3.7}$$

In the Julia simulation, the process is implemented directly with a simple loop over the particles, with each iteration performing the element-wise multiplication and sum between $S^{(ij)}$ and M_d^{eval}. This process is done solely on the CPU and is by far the slowest step of the simulation, accounting for over 90% of the computation time of the Julia simulation.

In the GPU-Julia simulation, we perform the calculation of $M_{d,L}^{(ij)}$ on the GPU. As discussed previously, there is little to no benefit to attempting to write custom CUDA kernels in Julia, and our focus therefore lies in structuring the calculation in such a way as to take advantage of the highly optimized large-matrix operations built into Julia's CUDA libraries. In particular, by restructuring M_d^{eval}, we can perform the calculation of $M_{d,L}^{(ij)}$ for all particles with two matrix operations without any need for loops. Each of these operations will be automatically optimally parallelized by the CUDA library's default parallelization structures. However, accessing these optimized functions comes at the cost of dedicating a non-trivial amount of computation time to re-structuring the calculation to make it compatible with these functions. Additionally, there is the further loss of efficiency due to the need to move memory back and forth between the GPU and CPU before and after the calculation of $M_{d,L}^{(ij)}$.

In order to re-structure the calculation, we first create a four dimensional array

$\hat{M}_d^{eval}$ where each slice is a copy of M_d^{eval}:

$$\left(\hat{M}_d^{eval}\right)_{ijk\ell} \equiv \left(M_d^{eval}\right)_{k\ell} \tag{3.8}$$

The element-wise multiplication step for all particles can then be performed with a single element-wise multiplication between $\hat{M}_d^{eval}$ and $\hat{S}$. Performing the sum is equally simple, using *mapreduce* on the two dimensional slices of the resulting four dimensional array. Julia's CUDA library automatically uses routines from CUDA's cuBLAS libraries for these operations, routines which are designed to take advantage of the GPU's potential for large-scale parallelization and can be assumed to be highly optimized. Despite the additional need for memory management and the creation of $\hat{M}_d^{eval}$, for a 100×100 lattice this restructuring results in a nearly ten-fold reduction in computation time, including the time spent on the additional steps.

The implementation of the $M_{d,L}^{(ij)}$ calculation in the CUDA simulation is the most direct of the three approaches, as the abilities to customize the parallelization structure and directly manage memory give us more freedom with how to structure the calculation. When calculating $M_{d,L}^{(ij)}$, each particle is given an index (which is also its thread index). CUDA does not facilitate multidimensional arrays well, so all arrays are flattened in row-major order. With this in mind, we can convert the spin/thread index, which we will call q, into a pair of indices which point to the particles position on the lattice:

$$\begin{aligned} j &= q \bmod n \\ i &= \frac{q-j}{n} \end{aligned} \tag{3.9}$$

where we are operating on a square lattice with $n_x = n_y \equiv n$. We then define two new sets of indices: k, ℓ and k', ℓ'. These indices are related by

$$\begin{aligned} k' &= (k+i) \bmod n \\ \ell' &= (\ell+j) \bmod n \end{aligned} \tag{3.10}$$

In this relation, i, j refer to particle position, k, ℓ refer to the stencil, and k', ℓ' refer to M_d^{eval}. Adding i, j to k, ℓ, modulo the lattice size, applies the shift previously calculated in $\hat{S}$. We calculate $M_{d,L}^{(ij)}$ by summing over k, ℓ:

$$M_{d,L}^{(ij)} = \sum_{k=0}^{n-1} \sum_{\ell=0}^{n-1} S_{k\ell} \big(M_d^{eval}\big)_{k'\ell'} \tag{3.11}$$

The ability to customize the parallelization structure allows this version of the calculation to bypass the need to calculate $\hat{M}_d^{eval}$ or $\hat{S}$ as in the GPU-Julia version while still allowing us to parallelize over the indices i and j. The CUDA simulation's calculation improves upon the performance of the GPU-Julia calculation by an additional twenty times for a 100×100 lattice.

3.1.2 Matrix Exponent

Although Julia has a built in matrix exponent, CUDA lacks such a function. To calculate the matrix exponent in the CUDA simulation, we use a handwritten version of the scaling-and-squaring algorithm, the same algorithm used by both Julia and MATLAB [84, 88]. The details this algorithm are beyond the scope of our discussion, but we will briefly outline its steps to give context to the details of our implementation.

We choose to use the scaling-or-squaring algorithm over other methods such as diagonalization or analytic methods because of combined desires for both accuracy and broad applicability. For unitary dynamics, the Lie-Trotter formula along with a knowledge of the analytic properties of the exponentiation of basis elements of $\mathfrak{su}(N)$ (i.e. anti-Hermitian matrices) can be used to great effect to create highly accurate analytic expressions which can be used to approximate the matrix exponent, bypassing a large number of computational steps [64, 89]. However, this method fails when non-unitary dynamics are included and the Hamiltonian cannot be written using a basis of $\mathfrak{su}(N)$. Methods such as diagonalization can be highly attractive, especially

for low dimensional matrices where analytic forms can be calculated. However, numerical diagonalization methods can become highly inaccurate when the eigenvalues of the matrix in question are too similar, as may be the case for a slowly oscillating spin during a weak phase of the interaction in our model [89]. Additionally, exact diagonalization rapidly becomes intractable for larger matrices, and even in the case of 4×4 matrices for a spin-$1/2$ particle in our model the analytic forms of the eignevalues become extremely cumbersome. Of the many methods for matrix exponentiation, the scaling-and-squaring method, although not perfect, remains one of the most accurate and widely applicable methods [89].

The scaling-and-squaring algorithm uses a Padé approximant to approximate the matrix exponent. In order to improve accuracy, the algorithm first reduces the norm of the matrix A to be exponentiated by scaling the matrix by a factor 2^s, with the choice of s dependent on the initial norm of the matrix A. The numerator P_m and denominator Q_m of the Padé approximant are then calculated (each a polynomial function of the matrix A of order m), and a linear solver is used to solve for the approximant itself. The resulting approximant is squared a number of times s to reverse the initial scaling, resulting in an approximation of e^A.

Although CUDA has accurate and efficient linear solvers via the cuSOLVER library, the solvers in this library are optimized around parallelizing over a single large matrix operation rather than an ensemble of small matrix problems, much like the cuBLAS libraries [87]. In order to parallelize over the ensemble, we use a handwritten Gaussian elimination and backwards substitution algorithm for our linear solver. Although such an algorithm scales extremely poorly, our simulation is designed for the modelling of independent, single-particle density matrices, which are typically two or four dimensional. The poor scaling is therefore not an issue in this context.

In order to examine the accuracy of our matrix exponential algorithm, we compared it to Julia's built in algorithm for matrices with elements from multiple orders

of magnitude. We focused on 4×4 matrices for the analysis, as the majority of our examination of this model occurs for spin $1/2$ systems (and the simulation is performed in Liouville space). To generate a matrix M, we first generate a set of four complex eigenvalues and save them in a diagonal matrix D. We then conjugate this matrix by a unitary matrix U, generated via exponentiation of a random Hermitian matrix H:

$$\begin{aligned} D &= \mathrm{diag}\,(\lambda_1, \ldots, \lambda_4) \\ U &= e^{iH} \\ M &= UDU^\dagger \end{aligned} \tag{3.12}$$

This process gives us a random complex matrix M with a known set of eigenvalues. We exponentiate this matrix using our CUDA algorithm and compare the result to the same matrix exponentiated using Julia's built in function. We analyze matrices whose eigenvalues fall in the ranges $10^{-p}, 10^{-p+1}$, for p ranging from 9 to 2. We limit ourselves to eigenvalues whose magnitude is strictly less than 1 due to the reliance of both the Julia algorithm and the CUDA algorithm on high order products of the matrix being exponentiated, resulting in poor behavior of both algorithms for matrices whose elements are larger than 1. We additionally analyze a set of matrices with mixed value eigenvalues, drawn at random from the full range 10^{-9} to 10^{-1}.

The error of the exponentiation of a given matrix using the CUDA algorithm, $U^{(c)} = e^M$, relative to the same calculated using the Julia algorithm, $U^{(j)}$ was calculated using the relative error of the real and imaginary parts separately. In particular, the error of a given matrix element $(U^{(c)})_{ij}$ has two values, one for the real part and one for the imaginary part, given by

$$\delta_{r,ij} = \frac{\mathrm{Re}\,(U^{(c)})_{ij} - \mathrm{Re}\,(U^{(j)})_{ij}}{\mathrm{Re}\,(U^{(j)})_{ij}} \tag{3.13}$$

and the equivalent $\delta_{i,ij}$ for the imaginary part. The "overall error" for a given matrix

is given by considering both the average across all matrix elements of the real and imaginary parts as well as by considering the average *magnitude* of the error of the real and imaginary parts across all matrix elements.

There were no identifiable trends in the error relative to the magnitude of the largest eigenvalue of the matrix. Across all regimes studied, the error was consistently small. On average, both the real and imaginary parts showed errors as defined above on the order of 10^{-8}, or $10^{-6}\%$. When considering the magnitude of the error, the result is similarly strong, with the real part having an average on the order of 10^{-6} and the imaginary part being on the order 10^{-7}. The distribution of error about these averages was similarly small, with each having a standard deviation on the order of $10^{-7}, 10^{-8}$. Overall, our handwritten matrix exponential algorithm in CUDA is extremely accurate relative to the Julia algorithm and can be treated as equivalent.

Although most calculations across both versions of the simulation are performed in single precision (as GPU performance is greatly improved by using single precision), the calculation of the matrix exponent in the Julia version of the simulation is performed in double precision. Although this somewhat reduces the performance of the Julia simulation, it improves accuracy– when applied to matrices stored in single precision with small matrix elements (e.g. eigenvalues on the order of 10^{-9} to 10^{-8}), Julia's matrix exponential shows significant divergence the CUDA algorithm, MATLAB's matrix exponential algorithm, and even from the same Julia function calculated in double precision. In order to avoid errors from cases like these, we perform the matrix exponentiation in double precision in Julia, converting back to single precision for the other calculations in the simulation.

3.2 Performance

To profile the simulation, we run fifty trials of each simulation with a range of values of Ω, as defined in Equation 2.49, in order to account for potential fluctuations in performance for different simulation parameters. Each trial was run on a 100×100 lattice, the standard lattice size we use for our simulations, with a sampling rate $1/dt \approx 3$MHz. We first examine the performance of the Julia simulation to establish a baseline against which to compare the improvements gained in the GPU-Julia and CUDA simulations.

Calculation/Process	Time per Execution (ms)	% of Total Time
Calc. Constants	381.161	0.177
M_d^{eval}	0.158	0.018
M_d	0.070	0.008
$M_{d,L}^{(ij)}$	722.676	81.421
H	30.802	3.456
U	131.007	14.698
Propagate	1.987	0.223
Total Time	419825	–

Table 3.1: Profiling of the major steps of the Julia simulation. % of total time is calculated by multiplying the time per execution by the number of executions (typically either 1 or $3n_\tau$), divided by the total time for the trial.

The performance of the Julia simulation is summarized in Table 3.1. The "Calculation of Constants" step summarizes the total time spent on variables that remain constant throughout the trial and are typically only calculated once or twice, primarily the stencils S and $\hat{S}$. These steps can be some of the slowest on a per-execution basis, but they are performed a fraction of the times as the other steps.

On average, trials took almost $420s$ each. By far the most dominant contribution to this poor performance is the calculation of $M_{d,L}^{(ij)}$ for $d = z, +$. The poor performance of this step is unsurprising, given that it essentially must iterate over 100,000,000 lesser calculations. The performance of this step is only better than di-

rectly serially performing 100,000,000 multiplications and 100,000,000 additions (on complex floats) by a factor of around 10 (tested locally, a single multiplication or addition of two complex floats takes around $25ns$, whereas performing the $M_{d,L}^{(ij)}$ calculation for a single particle takes around $50\mu s$, i.e. around $5ns$ per spin).

Calculation/Process	Time per Execution (ms)	% of Total Time
Memory Management	16.340	7.812
Calc. Constants	340.524	0.671
M_d^{eval}	0.245	0.118
M_d	0.124	0.059
$M_{d,L}^{(ij)}$	9.481	4.551
H	35.578	17.008
U	141.359	67.575
Propagate	4.615	2.206
Total Time	98528	–

Table 3.2: Profiling of the major steps of the GPU-Julia simulation. % of total time is calculated by multiplying the time per execution by the number of executions (typically either 1 or $3n_\tau$), divided by the total time for the trial.

The GPU-Julia simulation is essentially the same as the Julia simulation, with the exception of the calculation of $M_{d,L}^{(ij)}$, a fact we see reflected in the similar per-execution performance of many of the GPU-Julia simulation's steps, contained in Table 3.2. By moving the calculation of $M_{d,L}^{(ij)}$ to the GPU, however, we see a significant boost in performance. The calculation of $M_{d,L}^{(ij)}$, including the additional time spent calculating $\hat{M}_d^{eva}$, shows an improvement of almost 80 times. We suffer the additional loss of needing to manage memory between the GPU and the CPU, but this loss amounts to an additional $16ms$ per calculation of $M_{d,L}^{(ij)}$, but this loss is more than made up for by the improvement in the calculation itself. This improvement, as well as some minor gains in some of the stencil calculations accomplished by performing them on the GPU, reduce the per-trial time from $420s$ to $98.5s$, an improvement of almost 4.5 times.

However, given the majority of the remaining computation time in the GPU-Julia

simulation is in the calculation of the propagator, there is little room for further improvement in performance. CUDA does not have a built in matrix exponential function, as mentioned previously, so any additional improvements to the GPU-Julia simulation necessitate some amount of direct CUDA programming, at which point it is more optimal to translate the entire simulation into CUDA to additionally eliminate some of the overhead accrued by using hybrid GPU-CPU programming.

Calculation/Process	Time per Execution (ms)	% of Total Time
Memory Management	193.216	4.674
Calc. Constants	2.788	0.067
M_d^{eval}	0.008	0.088
M_d	1.126	12.882
$M_{d,L}^{(ij)}$	3.543	40.372
H	0.048	0.548
U	3.594	40.943
Propagate	0.037	0.426
Total Time	4134	–

Table 3.3: Profiling of the major steps of the CUDA simulation. % of total time is calculated by multiplying the time per execution by the number of executions (typically either 1 or $3n_\tau$), divided by the total time for the trial.

The CUDA simulation shows marked improvements across the board, with an overall reduction in per-trial computation time of almost 25 times relative to the GPU-Julia simulation, from over $98s$ to just over $4s$ per trial. Memory management is by far the most costly per-execution step, but given memory management occurs on average less than once per trial (since it is performed in the outer loop, a large fraction of the trails re-use the same set of allocated memory), it has little impact on the overall computation time. Parity has been achieved between the two most costly calculations, U and $M_{d,L}^{(ij)}$, with each consuming around 40% of the computation time. Performance decreased in the calculation of the signal, M_d, as this calculation is inherently serial and by nature performs poorly on the GPU. However, the calculation of M_d consumes only 12% of the total computation time, and methods such as *atomicAdd* exist that

could be used to improve its performance, with some additional development.

Overall, performance improved from $420s$ per trial to $4s$ per trial in the transition from Julia to CUDA, a gain of almost 100 times. The CUDA simulation additionally has more room to grow, as we use several inefficient algorithms that could be improved with additional development. Better parallelization schemes could be incorporated, as our scheme uses at most 10,000 threads per calculation. However, there are limits to the improvements available from increased parallelization. Local memory requirements, based on the number of variables declared on each thread during the calculations and the number of threads in a thread block, limit the number of threads that can be launched on a single calculation. Additionally, the total number of threads on the GPU set an absolute upper bound on parallelization. For example, a reasonably priced commercially available GPU, such as the GTX 1080 Ti, on which much of the development of this simulation was performed, has *at most* 57,344 threads available [90]. Further gains on a single GPU are limited, and higher level parallelization schemes likely merit cross-GPU parallelization, despite the additional overhead of managing memory between multiple GPUs.

Performance of our CUDA simulation is additionally on par with other CUDA driven spin dynamics simulations [64]. Although the fundamental structure and goals of the simulation in Reference [64] is different than ours, we can make a rough comparison of the execution times in each case. We assume relatively ideal conditions, with the simulation being maximally parallelized on a single GPU. We assume complete utilization of 50,000 threads based on the number of available cores in a GTX 1080 Ti, as we discuss above. In this situation, based on the computation time and simulation parameters provided by A. Tritt *et al.*, we roughly estimate a computation time of $2\mu s$ per spin per recorded time step (i.e. 100ms per simulation of a single spin, based on given results from a GTX 1070, with 50,000 time steps– each recorded time step seems to require its own GPU thread, and the remaining fraction of the

division of time evolution occurs in serial on a given GPU thread, with a 100ns integration time resulting in 20 additional divisions of the time step on each thread). The profiling data for $n = 10,000$ presented in Table 3.3 can be reasonably assumed to be representative of a maximally parallelized calculation. In this context, a computation time of 4s per echo with 500 time steps and 10,000 spins results in a per spin per time step computation time of $0.8\mu s$.

3.3 Validation

In order to validate the accuracy of the handwritten algorithms in the CUDA simulation as a whole, and to make sure the small variations between the two simulations do not affect results, we examine the error between the spin echoes generated by the CUDA simulation to those generated by the GPU-Julia simulation. Although the individual algorithms in the two simulations have only small variations between them, we do a comparison of the echo behavior to guarantee that these small errors do not propagate throughout a trial and accumulate to generate a noticeable difference. Furthermore, an examination of the echoes generated by each simulation can speak to the numerical stability of the model and serve as an indication of the degree to which the results of the model are independent from numerical fluctuations in the parameters and algorithms.

In order to analyze the similarity of the results of each simulation, we simulate a set of echoes with ω ranging from Γ up to 4.5Γ. We use the standard Γ scaling for our other parameters: $\tau = 5/2\Gamma$, $\Gamma dt = 1/10$, as well as an ensemble with $n = 10,000$ spins. We use a fixed frequency sampling for all trials, rather than resampling the frequency spectrum for each echo, to eliminate the variations between the echoes with the same parameters but different sampling.

We calculate the error ε between the signal from an echo calculated using the

CUDA simulation, $S^c(t)$, and the signal from an echo calculated using the GPU-Julia simulation, $S^j(t)$ by considering the magnitude of the error at each point, averaged over the whole echo.

$$\varepsilon = \frac{1}{L}\sum_{i=1}^{L}\left|S^c(t_i) - S^j(t_i)\right| \tag{3.14}$$

for L the length of the echoes in question. The value of ε, as well as its standard deviation, calculated using the individual errors of each point of each echo, is shown in Figure 3.1.

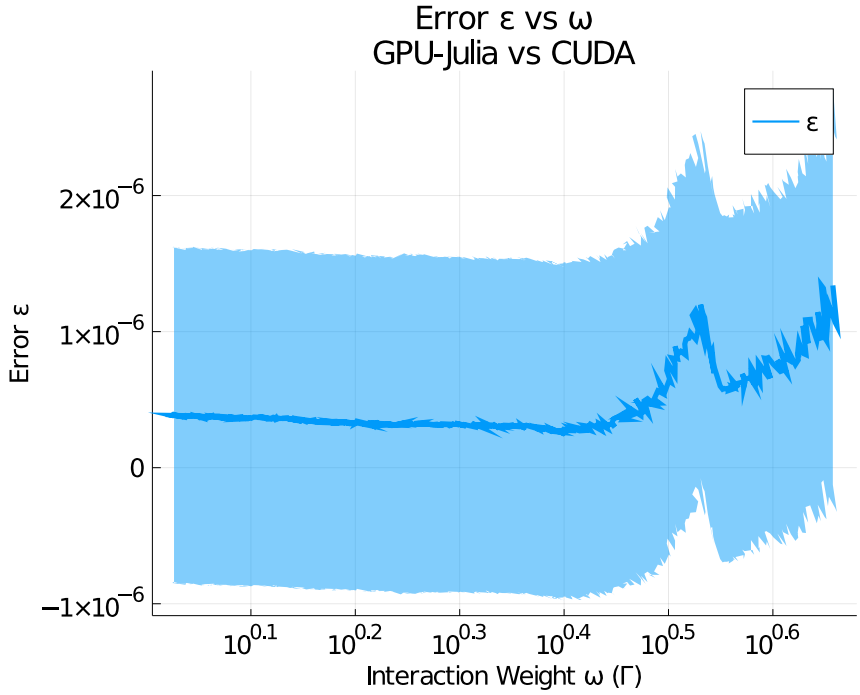

Figure 3.1: Error ε between echoes calculated by the CUDA simulation and the GPU-Julia simulation, plotted versus ω as a fraction of Γ. The shaded region represents one standard deviation about the error.

It is immediately apparent that the differences between the two simulations are negligible. The average error magnitude is on the order 10^{-7} and the spread in error is consistently on the order 10^{-6}. Even at multiple standard deviations from the mean, this represents an error of a fraction of a fraction of a percent, on the order of 0.001%. For all intents and purposes, the simulations can be considered identical. We

do see an increase in the error, although it remains on the order 10^{-6}, as ω increases towards values near Γ. In this region, the model enters a coupling regime where the interaction begins to compete with the dominant Larmor precession and the signature features of the model begin to emerge more strongly. This region is more sensitive to the exact values in the simulation, as the accuracy of the approximations used in the simulation become more important at higher values of ω. This uptick in the error is therefore expected given small differences in the numerics of the simulations, such as precision types and the solvers used in the matrix exponential algorithm. We recall the accuracy of our matrix exponential algorithm is on the order of 10^{-8} (Section 3.1.2) which is consistent with the error seen between the echoes of the two simulations.

3.4 Baseline Numerical Error

Before we turn to sources of error related to the model, we address one source of error that is unrelated to the behavior of the model specifically. The simulation uses floating point precision throughout to maximize performance. However, this use introduces the possibility for errors on account of the cutoff and rounding of single precision. In order to minimize this possibility, we use custom arithmetic operations built out of CUDA's intrinsic functions which gives us much more control over the rounding of the floating point values.

However, in one specific situation, there is still an observable error that arises from a combination of floating point point rounding and our data transfer protocols. In the case of extremely high sampling rate, that is very small dt, the vast number of matrix exponential products can result in an accumulated error which is small but observable. This primarily arises when comparing echoes with different but equivalently small time steps, for example an echo with sampling rate 1.5GHz and one with sampling

rate 2GHz. In this situation– which would almost never be relevant to an experimental application of the simulation, given the excessive computation time and diminishing returns in terms of improving accuracy– there is a "baseline error" in the accuracy of the echo shape.

In order to quantify this error, we use the matrix exponential function implemented in the simulation to compute a specific matrix in several ways. Specifically, we perform the following computation

$$e^A = \prod_{i=1}^{n} e^{A/n} \tag{3.15}$$

for $\log_2 n = 10$ up to 22 to represent the range of dt we study in Section 3.5.2. For our discussion, let E_n be the calculation of e^A using n products. We perform this calculation for 100 different matrices A all of which are complex valued and randomly generated with elements representative of the order of magnitude we see in our simulations. Additionally, the structure of A is the same as the Hamiltonian in our system. We then calculate the error ε via

$$\varepsilon(n) = \frac{1}{d^2}\left| \sum_{i,j}^{d} \left((E_n)_{ij} - (E_1)_{ij} \right) \right| \tag{3.16}$$

for d the dimension of the matrices, here 4. We consider the average value of ε over all 100 repetitions of this process and consider the result versus n. The results of this analysis are displayed in Figure 3.2.

The error is fit with the exponential function

$$\varepsilon(n) = c_1 \left(1 - e^{-n_2^c/c_3} \right) \tag{3.17}$$

Thus we see there is a baseline error on the order 10^-4 when an excessively small time step is used. This error *only* occurs for extremely high sample rates, specifically

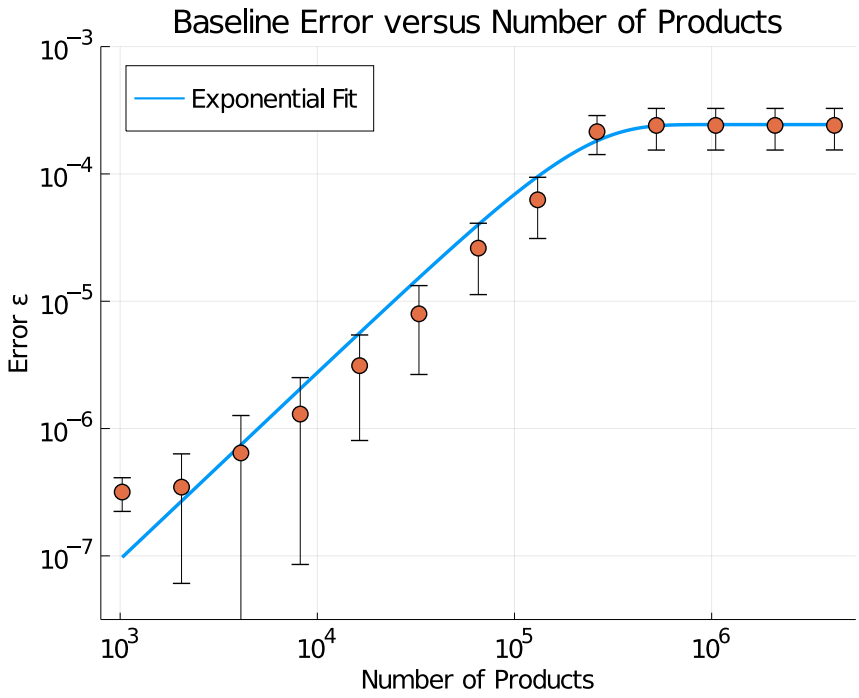

Figure 3.2: Baseline error ε versus n. The data is fit with an exponential fit $\varepsilon \propto 1 - e^{-n/c}$ for a constant c. The error bars give the 45% standard deviation of the error ε about its mean over the 100 samples. The error bars are scaled down to 45% their value on account of low n errors having a spread on the same order as their error and thus being unable to be plotted on a log scale.

those far beyond the necessary temporal resolution of the simulation. For smaller sample rates, this error exponentially disappears. For our simulations, we typically use a time step which would correspond to the low end of this scan, where the baseline error is on the order of 10^{-7} and can be safely ignored. However, we will see this error return in our discussion of the convergence relative to dt.

3.5 Convergence

In order to examine the numerical validity of the model, we study its convergence with respect to several key "artificial" parameters: in particular, n and dt. We study how well the simulation converges in the $n \to \infty$, $dt \to 0$ limits across a range of key physical parameters, in particular the linewidth Γ of the frequency spectrum and

the interaction strength (specified by α and the length scale ξ) in order to chart out the acceptable ranges for these numerical parameters (i.e. n and dt) in the various physical regimes of the model.

To study the convergence, we use a Gaussian length scaling function $f(r)$ defined by a length scale ξ:

$$f(r, \xi) = e^{-(r/\xi)^2} \tag{3.18}$$

Although the other two types of interaction scaling (RKKY and power law) result in slight variations in physical behavior, they do not have any numerical differences as long as their respective parameters are in the appropriate range. The results of the convergence with the Gaussian interaction can be safely generalized to the other forms of the interaction.

Given the phenomenological behavior of the model for $\xi > 5$ is generally determined not by particular values of the interaction strength α or length scaling ξ, but by the overall interaction weight ω (barring the emergence of phase locking effects at large ξ in a narrow band ω values, which we discuss in Chapter 4), and in order to reduce computation time to a realistic level, we do not vary α and ξ individually, but rather vary ω across a range of behavioral regimes and use a single pair of α and ξ to represent the regime. We generally use a value of ξ in the intermediate range where post-echo phase locking is not observed, but we allow some variation in order to scan a range of ξ as well. For a given value of ω, the value of ξ is kept constant for all values dt or n to ensure we are probing the same dynamics when performing the comparison.

In studying the convergence with respect to dt and n, we limit ourselves to isotropic interactions (i.e. $\alpha_x = \alpha_y$). In doing so, we eliminate any error accrued from the approximations of the Hamiltonian using the Magnus expansion and can focus solely

on the error that results from the ensemble size or time step. Coincidentally, as we shall see in a later section, higher order terms in the Magnus expansion, scaling as they do with $1/\nu$, will not alter the behavior of the model in any significant way, making the isotropic assumption essentially exact for any pair of values α_x, α_y to the limits of precision. Additionally, use of the Magnus expansion, at least for the terms we enumerate, puts limits on our choices of dt (i.e. that $dt = {}^{n\pi}/_{\nu}$), limiting our ability to study a range of values dt. We therefore on check isotropic pairs of $\alpha_x = \alpha_y$, or equivalently the first order approximation of the Hamiltonian in average Hamiltonian theory.

3.5.1 Convergence Relative to n

There are several parameters which interact with the size of the ensemble in non-trivial ways which can affect the convergence of the echo behavior with respect to the number of particles n. For a given n, increasing the linewidth Γ effectively reduces the sampling density of the dominant portions of frequency spectrum, which can adversely affect the accuracy of the echo. Additionally, given the interaction's dependence on a length scaling ξ, values of n which are too small can artificially increase the interaction weight ω by "cutting off" the tails of the interaction. In order to determine an appropriate value of n such that the overall behavior of the echo is sufficiently convergent to the $n \to \infty$ limit and such that errors from the interaction of n with Γ and ξ are minimized, we compute echoes for a range of Γ, ω, and n.

Additionally, as we shall see in Chapter 4, the dynamics of the system are inherently tied to the spin precession frequencies and thus the frequency sampling. In order to guarantee accurate behavior, the system must be sufficiently large that not only are all frequencies sampled in proportion to their spectral weight, but also that all frequency *environments* are sampled in proportion to their likelihood. The necessity to sample the frequency environments primarily becomes relevant in the case

of intermediate range interactions with strong ω. By examining the convergence of the system relative to ensemble size, we can identify an appropriate n whose error is within an acceptable tolerance.

As we will discuss in Section 4.3, the effective strength of the interaction is based on the value of ω relative to Γ, and the phenomenology of the echoes is consistent when ω and τ are both scaled with Γ. We therefore set the ranges which we scan based on the linewidth being studied in order to target equivalent behavioral regimes. We additionally scale the choice of dt with Γ to maintain a relatively consistent value of Ωdt. This helps reduce variation in the error on account of differences in effective temporal resolution, as the validity condition for the sampling rate is based on Ωdt and not just dt.

We study five linewidths from 10kHz up to 100kHz. Linewidths larger than 100kHz are not typically studied using single shot echoes, so we limit ourselves to this upper bound. We use $\tau = {}^{5}/_{2\Gamma}$, which corresponds to a value of $50\mu s$ for an ensemble with $\Gamma = 50$kHz. We vary ω from 1.5Γ up to 4.5Γ, a range which we choose based on observation of the values of ω which give the desired echo behaviors. Our variation in ω typically leads to values of ξ which are in the intermediate range. To attempt to limit the potential for n and ξ to interact adversely while still scanning small ensembles, we use a minimum lattice size of 25×25 ($n = 625$). We set the upper bound with a choice balanced between ensemble size and computational practicality, using a lattice of dimension 750 and $n = 525,000$. This upper bound of $n = 525,000$ will be used as the $n \to \infty$ approximation and the reference against which we calculate the error of the other echoes.

The error for a given echo is calculated by considering the average magnitude of error across the echo when compared to the reference echo. For an echo with a signal (absolute value) at time t_i of $S(t_i)$ and equivalent signal of the reference $S_{ref}(t_i)$, the

error ε is calculated by

$$\varepsilon = \frac{1}{L}\sum_{i=1}^{L}\left|S(t_i) - S_{ref}(t_i)\right| \tag{3.19}$$

for L the number of points in the echo of interest.

Unlike in the analysis of the convergence relative to dt, there is no convenient choice of universal convergence parameter which can be said to be constant relative to all Γ and ω. Because of the exponential relationship between ξ and Γ (via ω) and their non-trivial interplay with n, neither n nor n/Γ represent ideal parameters against which to study the convergence. Due to computational limitations, we choose to use a fixed range of n across all trials and examine the convergence of the model relative to it rather than Γ-scaled ensemble sizes.

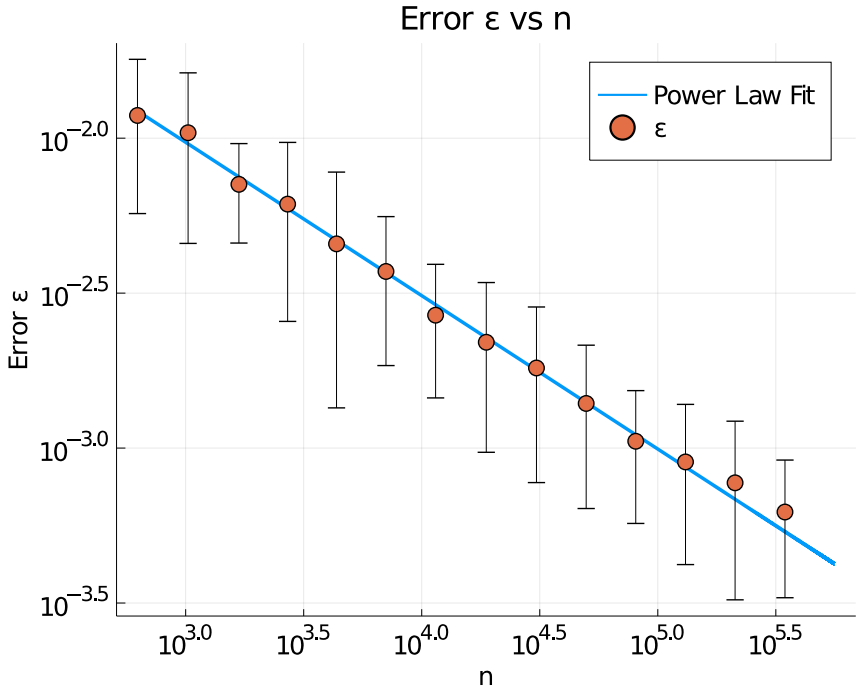

Figure 3.3: Average error magnitude ε for all ω, Γ vs n with a power law fit. Points represent the mean error across all ω, Γ at a given n and error bars represent the standard deviation of the error for that ensemble size.

A summary of the results of the convergence is presented in Figure 3.3. In Figure 3.3, each point represents the error ε averaged over all ω and Γ at a given n, and the

error bars are set by the standard deviation of the same. The data is fit with a power law

$$\varepsilon = an^{-b} \tag{3.20}$$

Power law convergence is consistent with the discrete random sampling approach we use to generate the frequencies of the ensemble. The fit has a power $b = 0.495 \pm 0.035$, indicating that the model on average converges as $1/\sqrt{n}$.

Importantly, we note that the error does not accumulate. Figure 3.3 shows the average magnitude of the error, but it does not encapsulate how the value of the error changes over the course of the echo. In short, it remains consistent, fluctuating near the value of the average error for the duration of the echo. In this sense, the echo does not exhibit any "drifting", or accumulation of error, over the course of the echo. The magnitude of the error near the end of the echo remains similar to the error near the beginning of the echo. The error at a given point in time, averaged over all values Γ and ω, is given in Figure 3.4.

The spread in the error is largely a feature of the variation in ω. The stronger interaction strengths are more susceptible to variation from the frequency sampling, as local fluctuations are accentuated by the stronger couplings. For higher linewidths with larger values of ω, this feature can become more dramatic, resulting in higher variation in the echo behavior compared to equivalent interactions at smaller linewidths. However, given the scaling of the dynamics with the linewidth, this higher fluctuation and error can be remedied by using smaller linewidths and sufficiently large ensembles where the fluctuations are less significant.

Based on these results, we use ensemble sizes of $n = 160,000$ when generating high quality echoes for direct analysis, as this guarantees the average error in the echo is less than 0.001 for a wide range of interaction strengths ω, or around 0.1% the amplitude of

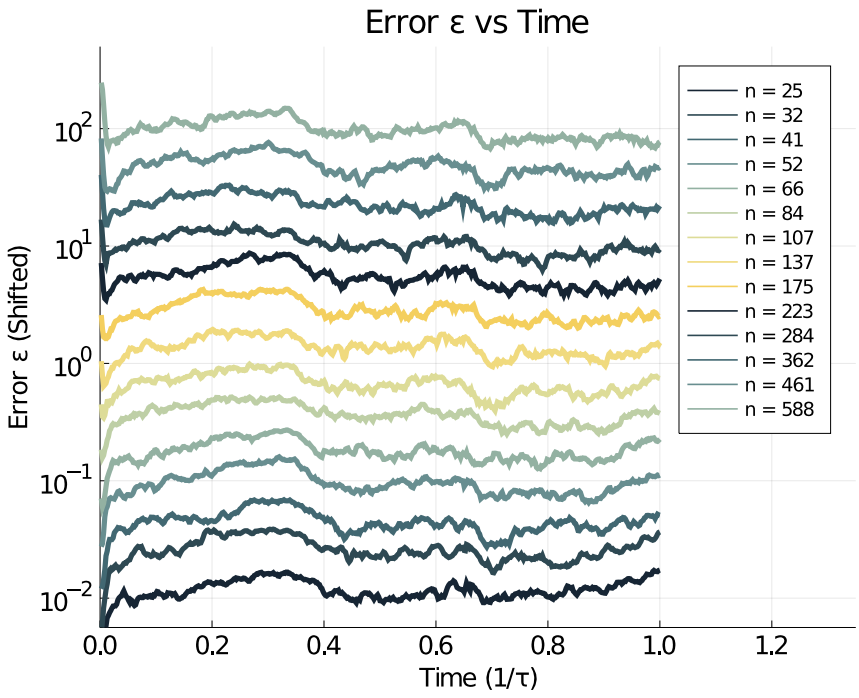

Figure 3.4: Error magnitude ε averaged over Γ, ω versus time. The error does not accumulate, i.e. the value at a time t_1 does not affect the error at a later time t_2. To improve readability, error values have been multiplicatively scaled to create a linear offset on the log scale.

the echo. Smaller ensembles such as $n = 10,000$ are occasionally used for generating large data sets where the broad phenomenology is the main goal (e.g. analysis of amplitudes) rather than direct analysis of the echoes. These echoes have much higher error, making the particulars of their shapes less reliable, but they are more efficient to calculate (around $3-5\times$ faster per spin), and the phenomenological behavior remains relatively consistent. As such, they can be used for high dimensionality sweeps where neither the spectrum nor echo shape is not the primary focus but the main focus is on "bulk" features such as amplitude or asymmetry. For a typical value of $\Omega dt \approx 1/10$, an ensemble of $n = 10,000$ results in computation times of around $4s$ per echo ($0.4ms$ per spin) and an ensemble of $n = 160,000$ of around 4 minutes per echo ($1.5ms$ per spin). The difference in per-spin efficiency arises from the higher ensembles reaching the limit of parallelization on a single GPU.

3.5.2 Convergence Relative to dt

The characteristic time scale of the dynamics of the model is set by the linewidth Γ and the interaction strength ω (Equation 2.47), both of which contribute to the parameter Ω, defined in Equation 2.48. For our approximation of the time evolution and application of the Magnus expansion to be accurate, dt must take on values relative to Ω such that the condition $\Omega dt \ll 1$ is met, as set in Equation 2.49. However, the degree by which Ωdt must be less than unity is not clear, especially since Ωdt is generally an overestimation of the strength of the dynamics. To appropriately map the accuracy of the model across the relevant parameters and identify an appropriate choice of dt we compute echoes for a range of Γ, Ω, and dt.

Given the allowed values of dt are set by the value of Ω, we define the range of dts over which we sweep relative to the value of Ω. We use a single set of dt values for each Γ across all Ωs, as the variation within the range of Ωs for a given Γ is only 10-15%. Additionally, since the behavior of the echoes is typically set not by the strict values of the interaction and experiment parameters (i.e. τ, ω) but by their values relative to Γ which we establish in Section 4.3, we scale these parameters as well, targeting equivalent behavioral regimes across the different linewidths. We will demonstrate later that re-scaling relative to Γ does not guarantee equivalent behavior across the different calculations, as the scaling of ω relative to Γ is not wholly linear. In particular, the scaling does not therefore numerically cancel when all relevant quantities are combined in the simulation.

For our choice of dt, we use a range bounded from below by $dt = {}^{8\pi\times10^{-5}}/\Gamma$ and above by $dt = {}^{\pi}/_{5\Gamma}$. The lower bound was set based on computational limitations. For a typical ensemble with $\Gamma = 50$kHz, these values result in a sampling rate $1/dt$ of 500kHz on the low end and 1.25GHz on the upper end. We use five values of Γ between 10kHz and 100kHz. Ensembles with higher values of Γ, (and even often those with Γ near 100kHz), are typically not experimentally studied using single echoes but

rather frequency sweeps, which our simulation does not model. We therefore limit ourselves with the upper bound of $\Gamma = 100$kHz. We range the values of ω from 1.5Γ up to 4.5Γ, values which were empirically observed to cover the range of behaviors we are targeting. Values of ω larger than 4.5 times the linewidth result in near destruction of the spin echo entirely and represent an interaction regime for which the spin echo is unsuited. We set the echo time τ as $5/2\Gamma$, corresponding to $50\mu s$ for a 50kHz linewidth. This value is large enough for both a full decay of the FID in weakly interacting systems and for phase accumulation and notable echo perturbation in strongly interacting cases. We use an ensemble of 10,000 spins for all trials, as the main quantity of interest is the error as a function of dt and not the accuracy relative to n. To eliminate fluctuations in the error on account of the sampling, we use a fixed frequency sampling for all echoes of a given linewidth.

Although our simulation is designed to perform calculations with dissipation and therefore both includes terms from the Lindblad equation and performs all computations in Liouville space, for the convergence data we set all dissipative terms to zero. This is a necessity due to the scale of the dissipation terms, which, when combined with extremely small values of dt result in a scale discrepancy between the dissipative terms and the non-dissipative terms in the Hamiltonian and result in large numerical fluctuations due to rounding errors from floating point arithmetic. This issue is not present at higher values of dt and only occurs at values of dt much smaller than any experimental applications of the simulation. However, those small values of dt are crucial for the $dt \to 0$ approximation, so we are forced to ignore dissipation when studing the convergence with respect to dt. Generally, the dissipation should not have an impact on the temporal convergence of the echoes, as the time scale of the dissipation (i.e. T_1, T_2) will be on the order of or longer than the entire experiment time and is therefore much longer than that of the interaction or FID.

The error for a given echo is calculated by considering the average magnitude

of error across the echo when compared to the reference echo. For each ω and Γ, we use the highest echo with the highest sampling rate (i.e. $dt = {}^{4\times10^{-5}}/_{\Gamma}$) as an approximation of the $dt \to 0$ limit, and calculate the error for all other dts with the same ω and Γ relative to it. Letting $S(t_i)$ denote the signal (absolute value) of the echo at time step t_i and $S_{ref}(t_i)$ denote the same for the reference echo, the error ε is given by

$$\varepsilon = \frac{1}{L}\sum_{i=1}^{L}\left|S(t_i) - S'_{ref}(t_i)\right| \tag{3.21}$$

for L the number of points in the echo of interest.

In Equation 3.21, S'_{ref} denotes the interpolated reference echo rather than the reference echo itself. Given the difference in their sampling rates, S_{ref} and S have a different number of points. In order to match the two echoes' sizes and calculate the error, we truncate S_{ref} to have the same number of points as S, using linear interpolation between points of S_{ref} to approximate its value at times that are between its recorded time steps. If S has a measured value at time t, and S_{ref} has values at $t_1 = t - \delta_1$ and $t_2 = t + \delta_2$, we calculate the value of S'_{ref} at t via

$$S'_{ref}(t) = S_{ref}(t_1) + \left(\frac{S_{ref}(t_2) - S_{ref}(t_1)}{\delta_2 + \delta_1}\right)\delta_1 \tag{3.22}$$

This interpolation scheme introduces a potential source of artificial error, especially at points in the echo with high curvature. However, the scale of this error is very small relative to the primary sources of error in the simulation, and it typically only becomes visible when comparing echoes both of which have very high sampling rates which are relatively similar in scale. This error can be eliminated by ensuring the reference echo has a much higher sampling rate than any of the echoes for which error is being calculated, if necessary. Typically, this error only becomes relevant in situations which are irrelevant to the practical applications of the simulation.

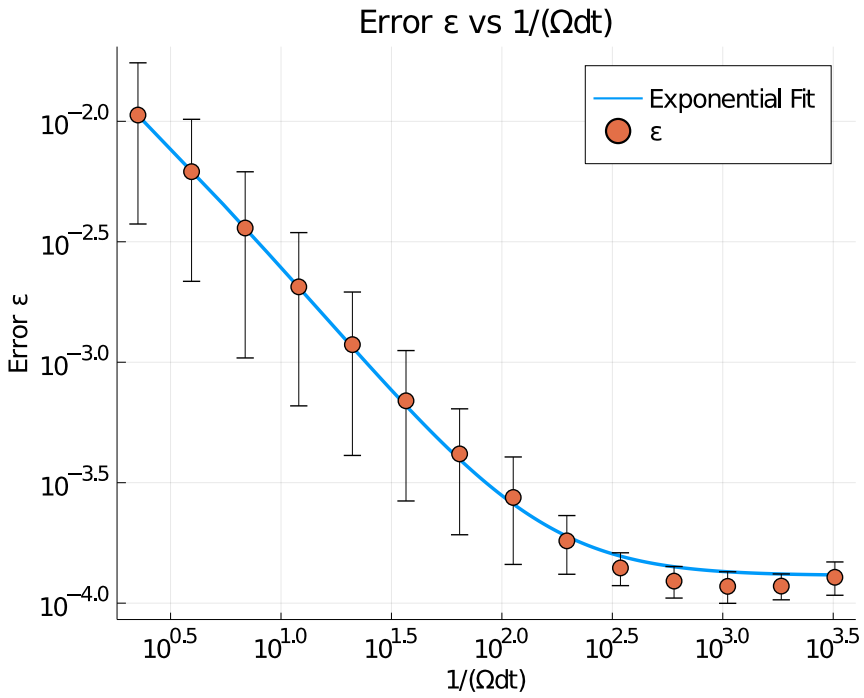

Figure 3.5: Average error magnitude ε for all ω, Γ vs $1/\Omega dt$. Points represent the mean error across all ω, Γ at a given $1/\Omega dt$ and error bars represent the standard deviation of the error for that sampling rate.

A summary of the error for all Γ, ω is presented in Figure 3.5. In Figure 3.5, the data points represent the mean value of ε across all values of Γ, ω that correspond to a given value of $1/\Omega dt$. The specific value of Ωdt varies for different values of ω, but the values of Ω only vary from the value $\Gamma/\sqrt{2}$ by a 10-15% across this range of ω. We therefore group terms based off the approximate value $\Omega \approx \Gamma/\sqrt{2}$ for clarity. The error bars denote the spread of values of ε for a given $1/\Omega dt$ and are given by the standard deviation of those εs.

The mean value of ε at each $1/\Omega dt$ is fit using a shifted exponential fit:

$$\varepsilon = ae^{-(x/b)^c} + d \tag{3.23}$$

for $x = 1/\Omega dt$. The value of the baseline error d is consistent with the numerical limits we observe in Section 3.4.

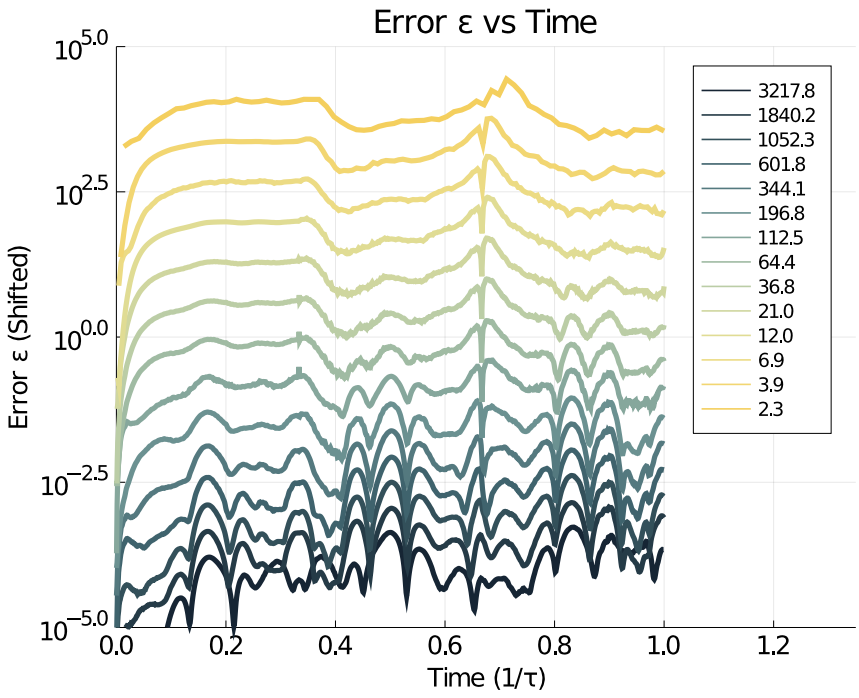

Figure 3.6: Error magnitude ε averaged over Γ, ω versus time. As in Section 3.5.1, the error does not accumulate. The legend indicates the value of $1/\Omega dt$ for the given error curve. The error curves have been multiplicatively shifted to improve visibility. The clumping near the highest values Ωdt is representative of the baseline error.

The spread in error at a given value of Ωdt (represented by the error bars in Figure 3.5) is broadly a feature of ω and not Γ. All linewidths study show trend of higher error at a given Γdt as ω increases, which is consistent with the model structure. In the absence of ω, the Hamiltonian is time independent and the approximations we use to time evolve the system become exact. For small ω, the model is only weakly time dependent, and the approximations remain strong and the relative error low. For larger values of ω, the Hamiltonian's time dependence becomes more pronounced, the approximations become less reliable, and the error at low $1/\Gamma dt$ correspondingly increases. The spread narrows as we increase $1/\Omega dt$, as the sampling rate becomes so much larger than the characteristic frequency of the interaction that any variation in the accuracy of the approximations for different values ω is more than accounted for by the temporal resolution. Additionally, as we discussed in Section 3.5.1, there

is no "drift" in the echo. The magnitude of the error does not have an increasing trend across the experiment, varying based on the features of the echo (e.g. the error is typically higher near the echo peak) and not the number of time steps that have passed.

The model shows strong convergence with respect to the choice of time step dt. The accuracy of the simulation relative to the choice of dt depends only on its value relative to the linewidth Γ, assuming we target regimes of equivalent behavior, with some fluctuation depending on the choice of ω (although the fluctuations across Γs are consistent). In order to guarantee an error of less than 0.1%, depending on the value of ω, a value of $1/\Omega dt$ between 5 and 20 is required. Typically, we target values of $1/\Omega dt \approx 10 - 15$, which results in a computation time of around 4-6s per echo for $n = 10,000$, or 4-6 minutes for $n = 160,000$.

3.5.3 The Magnus Expansion

As we discuss previously, the terms in the Magnus expansion scale with $1/\nu$ and are expected to be multiple orders of magnitude smaller than the first order term. Indeed, implementation of the higher order terms has next to no influence on the observed echo behavior, with total error between extremely high sample rate echoes (where the sample rate is sufficiently small to capture behavior on the scale of the Larmor frequency) and echoes using no higher order terms being essentially identical to the error observed when using the Magnus expansion. In all cases, the error was on the order of 0.1% of the echo amplitudes. The error in all cases is essentially identical to the expected error from the higher time step and established in Section 3.5.2.

This result follows naturally from the time scale of the different features of the interaction. The interaction axis, dependent as it is on the local magnetization of the spin ensemble, oscillates about the z axis with a frequency near or at the Larmor frequency. Even with dramatic differences in the two planar couplings, the time scale

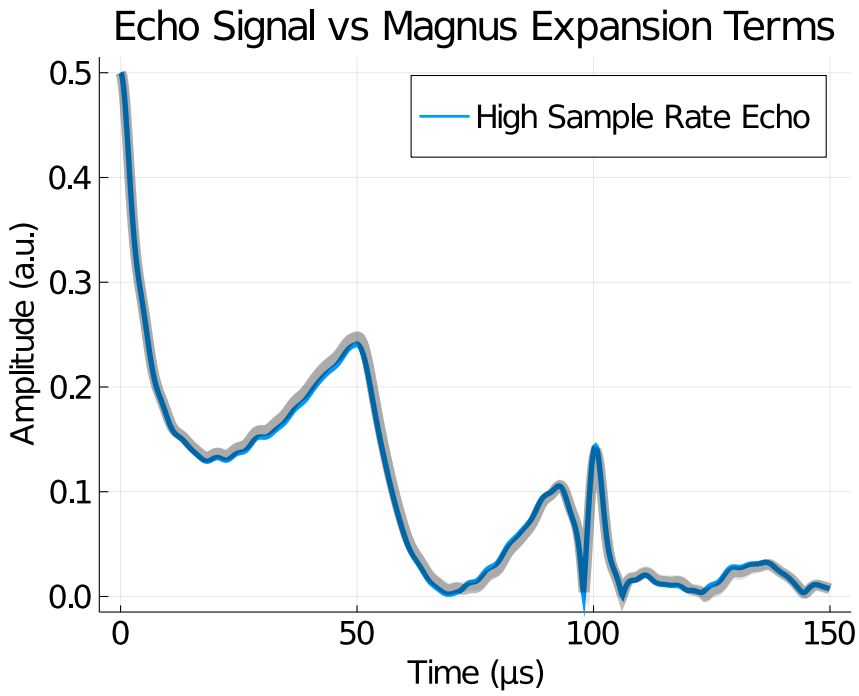

Figure 3.7: Comparison of echoes using 0 (exact Hamiltonian with time dependence) to 3 terms of the Magnus expansion (the first order term being the equivalent isotropic case) compared to an echo with extremely high sampling rate.

of the oscillation is so much larger than that of the interaction that any difference between the planar couplings is immediately washed out by the precession of the spins. However, as we observe in the spectral shifts induced by the interaction, the difference between the planar couplings is still accessible through re-orientation of the crystal relative to B_0.

The lack of any non-negligible contributions from the Magnus expansion allows us to wholly ignore the explicitly time dependent terms in the Hamiltonian and frees us to use an assumption of planar isotropy when examining the dynamics of the model, given the dynamics are solely determined by the relative strengths of the in-plane versus the out-of-plane coupling and not by the differences between the planar couplings. This additionally significantly reduces the computational complexity of parameters sweeps, as a single planar coupling strength is sufficient to capture all dynamics of the model. The lack of dependence on the explicitly time dependent

terms also allows us to use larger time steps without worry of incurring any additional error in the case of planar anisotropy.

Chapter 4

Results

Having reviewed the details of the model, we can now examine its effects and the manner in which it induces the unique temporal and spectral features we observe. In the spin-1/2 case, we can use a simple geometric interpretation of the spins' behavior to understand how this interaction induces the multipeak behavior observed in the echoes. Although the geometric argument does not apply to higher spin systems it offers a great deal of insight into the behavior of the system, and additionally the features observed in the echoes and spectra will generalize to higher spins on account of the dependence of the Hamiltonian only on the spin expectation values. We first discuss the most simple case– a global interaction where $\xi \to \infty$ and all spins experience the same interaction magnitude– before examining specifically how the introduction of the local coupling changes the behavior of the model.

4.1 Global Coupling

In the geometric picture, the presence of the I_x, I_y terms act to induce a rotation about the axis of the interaction in addition to the precession from the Zeeman term. We typically use pulse phasing $(0, 1) = (x, y)$ or $(0, 3) = (x, -y)$ which aligns the echo signal along the y axis, and the spins experience an additional rotation about

the y axis. The result of this additional rotation is to shift the precession axis away from the z axis and toward the y axis. Because the spins now precess about this tilted axis, their precession does not fully dephase them and there remains a lingering net magnetization along the interaction axis, as we see in Figure 4.1. This "phase locking" effect, and how it affects the behavior of the spins after the π pulse, is the primary driver of the multi-peak behavior observed in the echoes. For a strong enough interaction, the phase locking effect can result in spin coherence times much longer than the echo time.

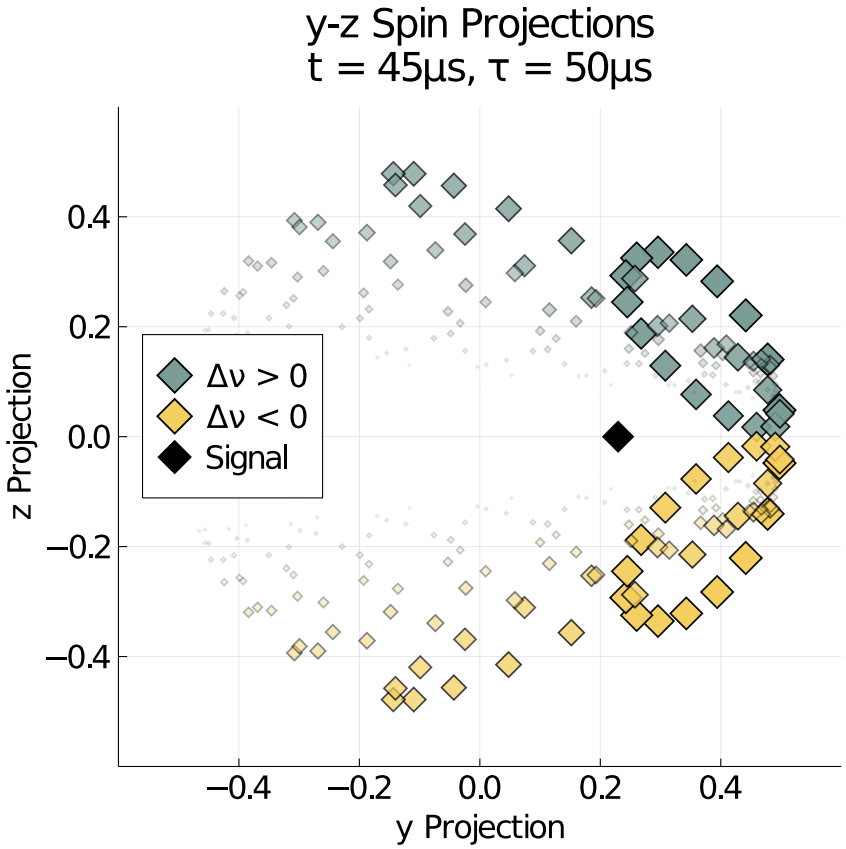

Figure 4.1: $\langle I_z \rangle$ vs $\langle I_y \rangle$ at $t = 40\mu s$ in an echo with $\tau = 50\mu s$ and $T_2^* = 1/\Gamma = 20\mu s$. Each marker indicates the location of a spin. The signal amplitude along the y axis is more than $4\times$ what is expected given exponential decay. Marker transparency and size is scaled with the spectral amplitude at that spin's $\Delta\nu$, with larger and darker spins having higher spectral weight.

The angle of the tilted precession axis varies for each spin and is set by the strength

of the interaction relative to the spin's Larmor frequency. Spins which oscillate more rapidly relative to the Larmor frequency have rotational axes more closely aligned with the z axis, whereas those spins which have precession frequencies near the Larmor frequency have an axis very close to the plane. The effect is symmetric about the y axis, with spins whose $\Delta\nu < 0$ having an axis below the plane, and spins with a positive shift aligning above the axis. Increasing the strength of the interaction more strongly pulls these axes toward the plane, accentuating the phase locking effect and increasing the magnitude of the persistent signal. Depending on the phase of a given spin and the sign of its $\Delta\nu$, the interaction also applies a shift to its planar precession frequency. The observed effect varies depending on the value of a spin's precession frequency and the orientation of the spin (i.e. the vector defined by its expectation values) relative to the interaction axis (in our case, the positive y axis).

In an actual spin echo, the phase locking effect is seen most prominently during the FID. After the initial pulse, the net magnetization is very strong and the spins begin to dephase about the z axis. The interaction's additional rotation about the y axis causes the spins to spiral out, gaining a non-zero z component as they rotate. Eventually, as the net magnetization decreases and the spins dephase, the precession axis for the spins partially stabilize, resulting in a nutation effect as the axes oscillate in some slice of the $y-z$ plane. This effect is observed in the time domain of the FID as persitent signal which oscillates about some non-zero average value. In the frequency domain, the interaction's rotation opposes the natural precession of the spins by an amount proportional to the value of $\langle I_z \rangle$ for the spin, resulting in a narrowing of the spectrum.

However, after the π pulse and the spins are flipped about the y axis, the z component of each spin changes sign, and the interaction no longer opposes the natural precession of the spins but reinforces it. This results in the spins rapidly dephasing and a sudden drop in the signal. As the precess and the signal decays, the interac-

tion dies off and the spins enter a period of free precession. However, the lingering coherence from the phase locking as well as the force of the interaction causes the positive ($\Delta\nu > 0$) and negative ($\Delta\nu < 0$) spins to precess in "clumps". Depending on the strength of the interaction, the degree of coherence from the phase locking, and the orientation of the spins at the π pulse, this "clumping" effect can result the spins re-cohering along the $-y$ axis, creating the peak we see prior to the echo.

Once the spins precess and re-cohere along the $-y$ axis, a second period of phase locking can emerge as the alignment of the spins and the interaction results in the interaction once again opposing the natural precession of the spins.

As $t \to 2\tau$, the spins begin to realign due to the reversal of the Larmor precession induced by the π pulse. This re-coherence favors alignment with the $+y$ axis, causing the coherence along the $-y$ axis to suddenly and rapidly decay and the net magnetization to suddenly change sign. This "snapping back" of the spins is destructive to the remaining interaction-induced coherence and typically prevents the emergence of more peaks after the echo. With a few exceptions, any peaks emerging after the echo are typically highly attenuated and negligible compared to the echo and pre-echo peak.

With this understanding, we can classify the different behaviors we observe in the time domain. The persistent signal in the FID occurs as a result of the phase locking effect, and the additional peaks occur as a result this lingering coherence at the π pulse. For a weak interaction, the axis about which the spins rotate during the FID is very nearly aligned with the z axis and the persistent coherence is minimal and insufficient result in additional peaks before the echo. However, during the echo, the interaction's strength grows significantly and the torque from the interaction causes a small degree of phase locking during the echo which results in a small secondary peak along the $-y$ axis after the echo, by the same mechanism as the pre-echo peaks for stronger interactions. However, this phase locking is too weak to induce additional

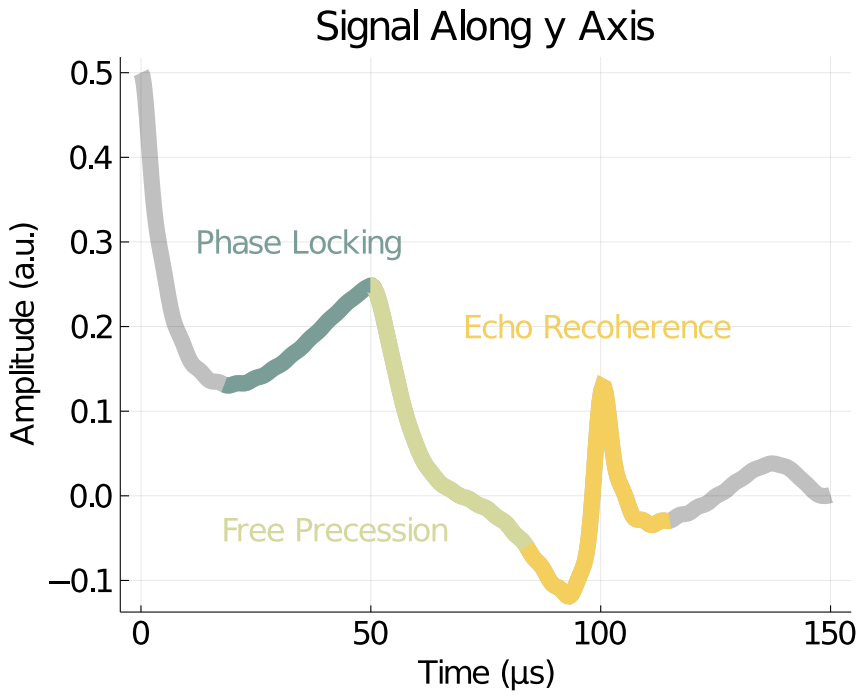

Figure 4.2: Echo phases under the influence of the interaction for $\omega = 3.05\Gamma$, $\Gamma = 50$kHz. The π pulse occurs at $t = 50\mu s$.

peaks beyond this single additional peak.

As the interaction strength increases, the coherence from the phase locking during the FID grows in strength, and the resulting initial peak prior to the echo begins to emerge. From here, stronger interactions cause a small variety of behaviors that depend on the complex relationship between the interaction strength and the spin orientations at the π pulse, but they are typically characterized by the increasing strength of the pre-echo peak relative to the echo. As the interaction grows in strength, the strength of the secondary phase locking coherence competes with and is disrupted by the strength of the recoherence from the Larmor precession until they begin to cancel entirely, and we see no echo signal at all but rather a sudden and complete destruction of the spin coherence at the echo time. Beyond this regime, the phase locking begins to dominate, and ultimately the system enters a regime wherein the Larmor precession plays almost no role and the phase locking effect is persistent for the entirety of the experiment. Such a regime only occurs with extremely strong in-

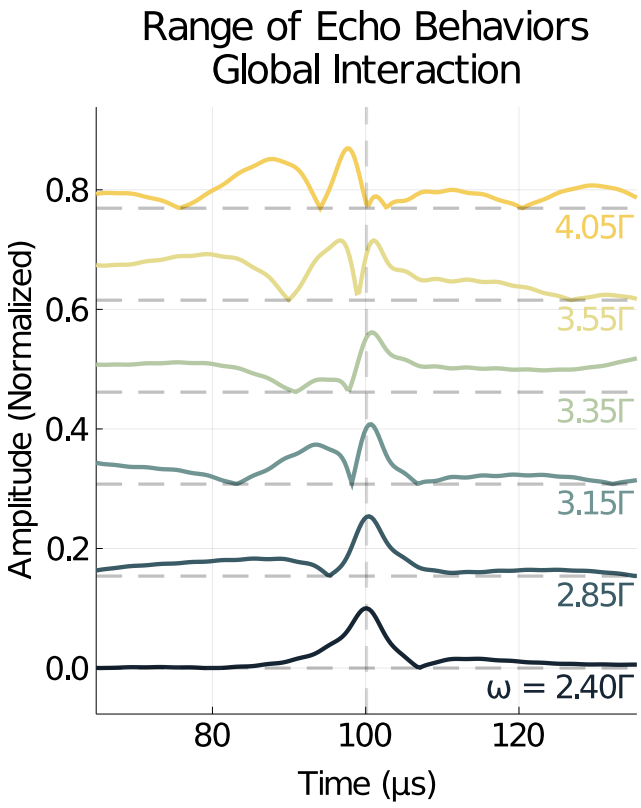

Figure 4.3: Variety of echo behaviors for a global interaction. Post echo phase locking occurs for global and nearly global echoes in a narrow region near $\omega = 3.35\Gamma$.

teractions that are not expected in a real system. A sample of several echo behaviors is given in Figure 4.3.

In the case of a global or nearly-global interaction, there is an additional region where the balance between the interaction strength and Larmor precession allows the re-emergence of phase locking along the positive y axis after the echo. Such behavior is only possible for a narrow region of interaction strengths in the long range limit and is unlikely to occur in a real system due to the effect of dissipation from thermal effects or other spin-spin interactions.

There are two many ways the interaction's effect is exhibited in the spectral domain. The first we have already introduced: a narrowing of the spectrum during the FID from the adversarial effect of the interaction relative to the natural Larmor precession. The other observed feature is an effect akin to hole burning during the

echo. The hole burning effect is primarily seen during the spin echo, when the sudden change in sign of the interaction from the echo re-coherence interacts with the post-pulse orientation of the spins to strongly shift the frequency of those spins which oscillate near $\Delta\nu = 0$.

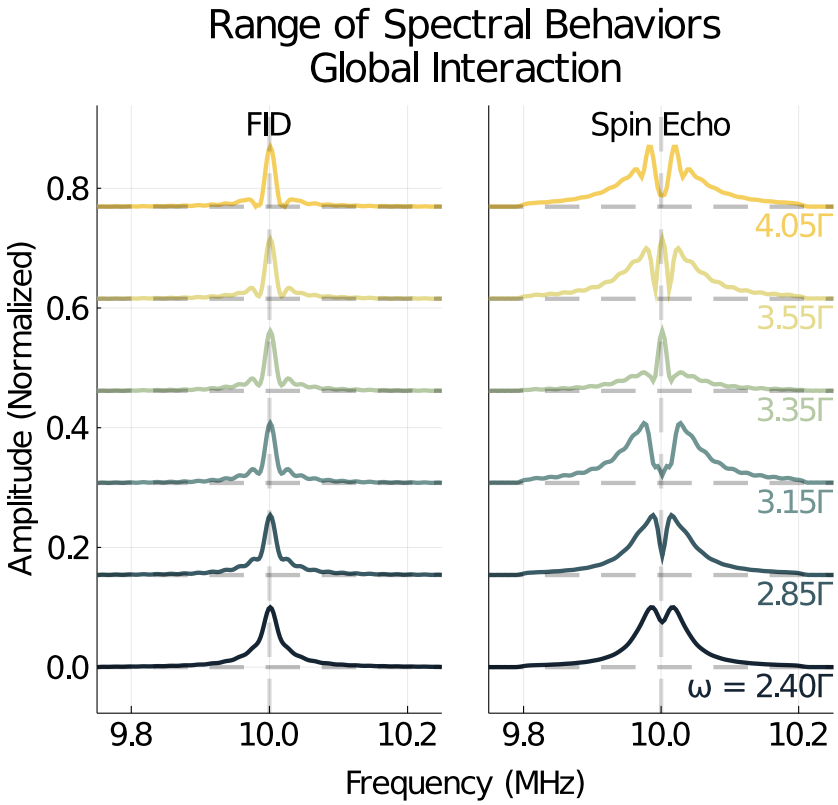

Figure 4.4: Variety of spectral behaviors for a global interaction for a spectrum obtained from the FID and the spin echo.

The spectra in Figure 4.4 capture much of the information we see in the echoes. In the FID, we see significant spectral narrowing from the effect of the phase locking. In the echo, we see the emergence of a hole burning effect in the spectra of the echo for strong enough ω. Around $\omega = 3.35\Gamma$, we see the emergence of a narrow peak near $\Delta\nu = 0$ and the echo's spectrum begins to resemble that of the FID, indicating the emergence of phase locking behavior during and after the spin echo. As ω grows, this effect dissipates and the hole burning effect returns.

The introduction of the z coupling in the global interaction has no effect when the spins are flipped perfectly into the plane. Given the interaction is global, the z coupling will only impact the dynamics of the spins when there exists a global shift away from $\langle I_z \rangle = 0$, i.e. when the spins are flipped less than ninety degrees.

An examination of the effect of the pulsing on the time domain is generally less informative than investigating the spectral effects. However, briefly discussing it in the global limit helps reinforce our understanding of what the more informative spectral features represent.

For pulsing in the range $\pi/2 \pm \pi/6$ the effect of the pulsing is generally to dampen the multi-peak behavior on account of the effective reduction of the interaction strength. For pulsing in the $\pi/6$ to $\pi/3$ range (or the equivalent over-pulsing), the multi-peak behavior vanishes almost entirely, and a post-echo phase locking effect is observed instead. This region of phase locking likely occurs from the global shift in z magnetization applying a more consistent torque to the spins resulting in a lower threshold for the phase locking to take effect. For pulsing less than $\pi/6$ the effect of the interaction on the time domain echo begins to disappear entirely as the driving force of the temporal behavior, the planar interaction, is strongly attenuated by the reduced in-plane projection of the ensemble magnetization. Introducing non-zero z coupling has little effect on this behavior except to dampen the phase locking effects at pulsing near $\pi/6$.

In the spectral domain, the introduction of z coupling will result in an overall pulse dependent frequency shift. A similar effect is seen from the influence of the planar interaction. When the spins are flipped exactly to 90°, they drift out of the plane symmetrically based on their value of $\Delta\nu$ and although each spin experiences a small torque resulting from the interaction, there is no net torque on the ensemble which shifts the spectrum as a whole. When the spins are not flipped fully into the plane, or are flipped past it, the consistent non-zero $\langle I_z \rangle$ of the spins breaks the symmetry

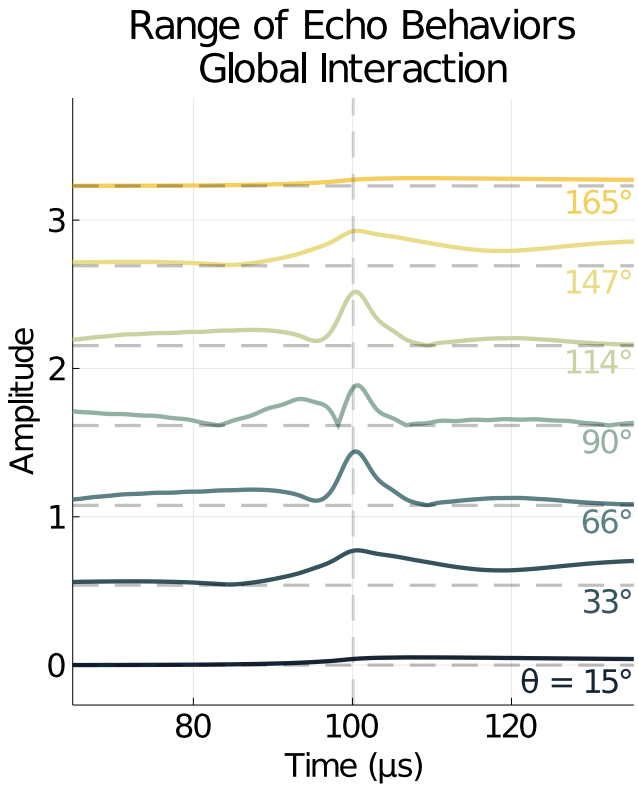

Figure 4.5: Variety of echo behaviors for a global interaction with $\omega = 3.15\Gamma$ for varying flip angles θ.

between the spins of different sign $\Delta\nu$ and the spins experience a net torque, resulting in an overall frequency shift. These effects oppose each other, although the size of the shift from the out-of-plane coupling is typically larger than the equivalent shift from the in-plane coupling at $\omega = \omega_z$. The exact form of the dependence of these shifts on the flip angle in the global case is derived in our companion work [69].

In the frequency domain of the spin echo, we can identify some of the same features and behaviors as we saw when looking at the perfect pulsing case. For the small flip angles which result in observable phase locking in the time domain, we see the same narrowing of the peak we see during the FID. As we get closer to 90°, we see the re-emergence of the hole burning effect. The general type of behavior observed (i.e. narrowing or hole burning) is generally unaffected by the additional shifts, but overall shift introduces a large amount of asymmetry in the observed effect, especially for

the hole burning. The FID differs from the echo in that it sees an overall reduction in the phase locking effect (i.e. the narrowing) as the flip angle gets further from $\pi/2$ rather than a period of increased locking at low flip angles. Both experience the pulse dependent shifts dependent on the relative values of ω_z and ω.

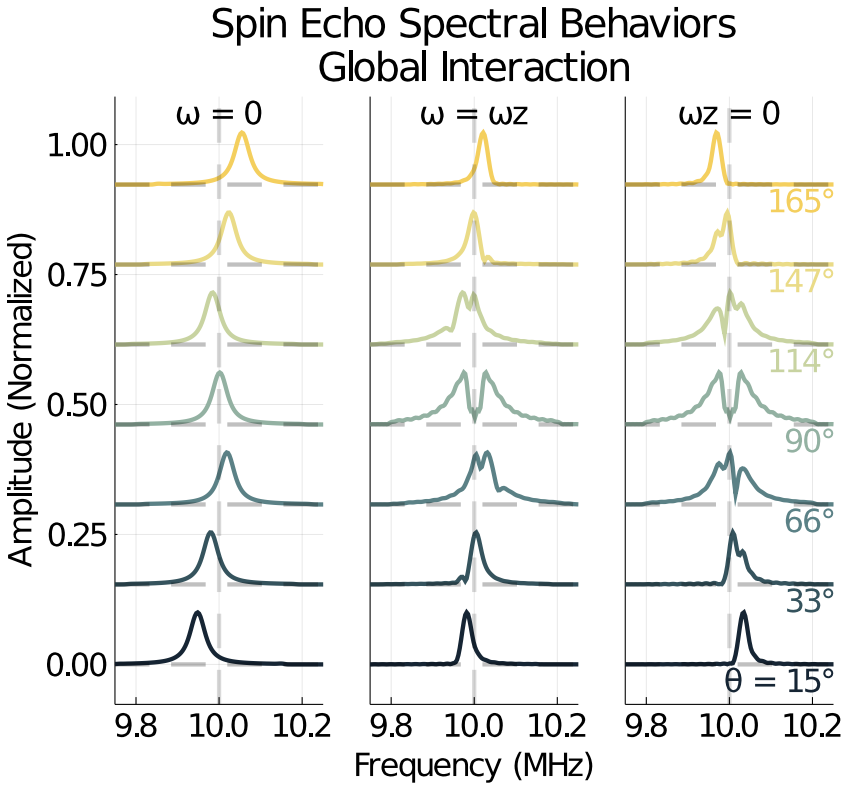

Figure 4.6: Variety of spectral behaviors observed in the spin echo for different flip angles for (ω, ω_z) equal to $(0, 3.15\Gamma)$, $(3.15\Gamma, 3.15\Gamma)$, and $(3.15\Gamma, 0)$.

We note that the FID is difficult to study in the pulse-dependent case, as the inherently asymmetric structure of the FID signal interacts poorly with the envelope function used in our Fourier transform, which we describe in our companion work [69]. To bypass this issue, we adjoin a time reversed copy of the FID to the FID itself to create a symmetric signal to which we apply the Fourier transform. In the case of perfect pulsing, this causes no issues. However, in the case of varied pulsing, this mirroring process also mirrors the overall shift the spins experience from the z

coupling, resulting in two identical peaks mirrored over the Larmor frequency. Not using the mirroring results in artificial narrowing of the spectral line, so we present this dual-peak spectrum with the understanding that it is the superposition of two copies of the actual peak mirrored about the Larmor frequency.

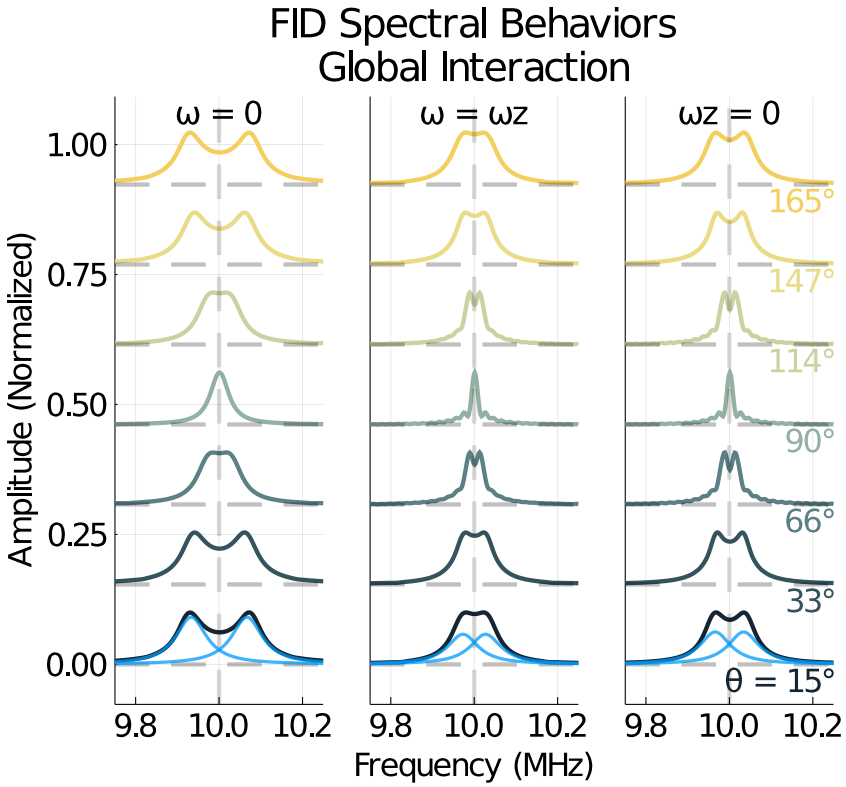

Figure 4.7: Variety of spectral behaviors observed in the FID for different flip angles for (ω, ω_z) equal to $(0, 3.15\Gamma)$, $(3.15\Gamma, 3.15\Gamma)$, and $(3.15\Gamma, 0)$. The spectra for $\theta = 15°$ are overlaid with the two identical, mirrored Lorentzian peaks whose sum fits the dual peak, indicating the effect of the mirroring.

4.2 Local Interactions

In the case of in-plane coupling only, the introduction of the localization has several effects on the interaction. For a very short range interaction, where the spins only interact strongly with a small number of neighboring spins, the dynamics of the

spins become highly dependent on local fluctuations in frequency and magnetization, dampening the effect of the interaction. Although phase locking is still observed, the degree of coherence the interaction can induce is limited by the local fluctuations, which ultimately dampen the effect and typically attenuate any additional peaks. Echoes of an intermediate correlation length typically behave similarly to those of nearly-global interactions, with the notable exception of the post-echo phase locking mentioned previously.

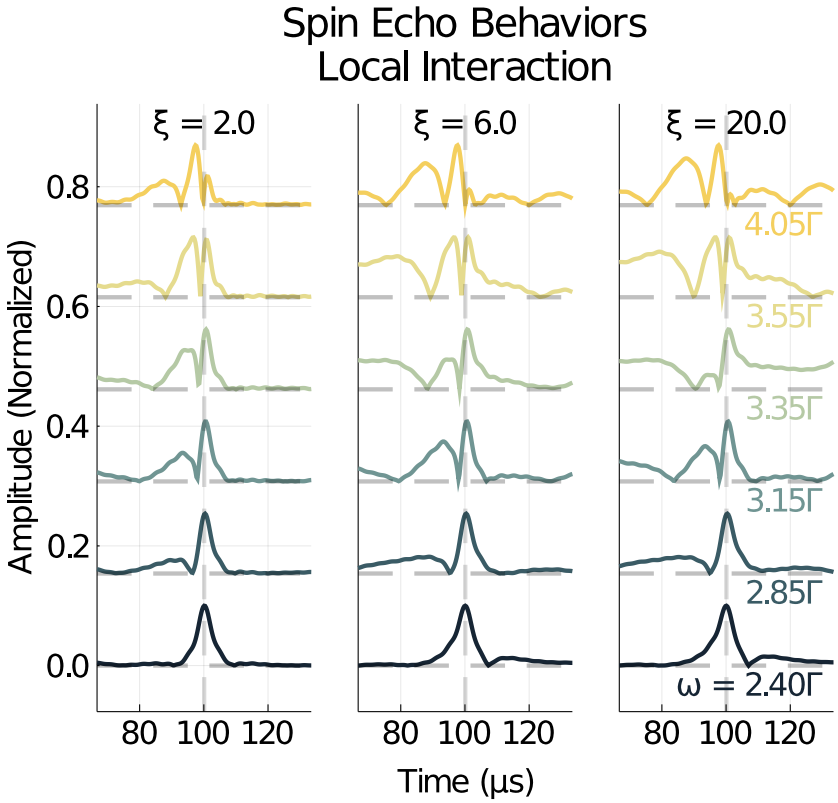

Figure 4.8: Variety of echo behaviors for local interactions for varying ξ with $\omega_z = 0$. Echo amplitudes have been normalized to improve visibility.

Figure 4.8 displays the variety of echo behaviors for several key ξ ranges. The echo amplitudes have been normalized to accent the different behaviors, but the stronger ω echoes have generally smaller amplitudes. These echoes were run on a 400×400 lattice with $n = 160,000$. The $\xi = 20.0$ results are nearly identical to the global interaction

results, indicating the strong convergence toward the global limit when the interaction becomes sufficiently long range. We can also immediately see the dampening effect of local fluctuations for short range interactions, with the most extreme behaviors only emerging at very strong ω. Additionally, at these short interaction ranges the fluctuations in local interaction strengths are sufficient to totally dampen the post-echo effects. However, at ranges as short as $\xi = 6.0$, we see the echo behavior begin to converge towards the long range/global limit, with only the most exotic behaviors, such as the post-echo phase locking, not appearing at the intermediate interaction lengths. The behavior for other values of ξ is consistent with these primary regions, with echoes in the $2.0 < \xi < 6.0$ rapidly shifting towards the $\xi = 6.0$ behavior followed by a slower asymptotic convergence towards the global behavior. We can see the shifting of the echo behavior by examining the first moment of the echo, calculated over the range shown in Figure 4.8.

$$\mu_1 = \sum (t_i - 2\tau) S(t_i) \tag{4.1}$$

for S the signal magnitude, given by the absolute value of the ensemble planar magnetization.

In Figure 4.9, we can clearly see that by $\xi \approx 5.0$, most of the effects of very short range couplings (in the case of $\alpha_z = 0$) have disappeared, and the echoes begin to asymptotically approach the global limit. The exception to this trend is the presence of the post echo phase locking, which we can see as a narrow band of near-zero first moment emerging in the long ξ limit near $\omega = 3.3\Gamma$.

Although the degree of localization for a wholly in-plane interaction does not have a significant effect on the echo behavior except in comparing the extremes, the introduction of out-of-plane coupling has a significant impact, even in the case of perfect pulsing. The introduction of the z coupling results in the local fluctuations in

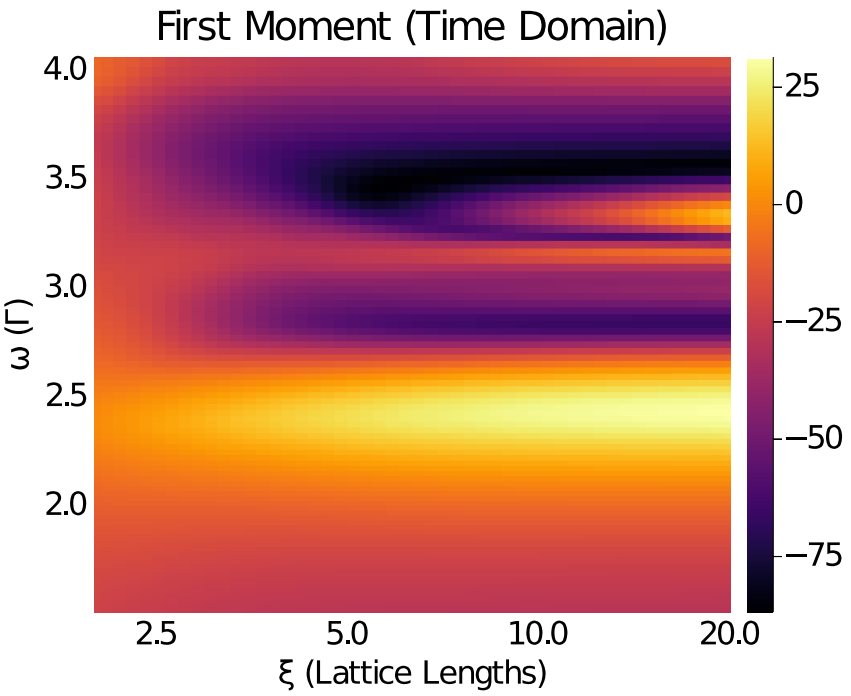

Figure 4.9: First moment μ_1 for echoes with $\omega = 1.5\Gamma$ up to $\omega = 4.05\Gamma$ and $\xi = 2.0$ to $\xi = 20.0$. A strongly negative first moment indicates a large pre-echo peak, and a highly positive first moment indicates a large amount of post-echo signal.

magnetization and frequency to induce fluctuations in the spin precession frequencies, resulting in irreversible decoherence. For highly local interactions, even a small degree of z coupling can result in total destruction of all signal, including that signal which results from the phase locking. As the interaction's correlation length increases, the echoes become more resilient to the dissipative effects of the out-of-plane coupling, as the ensemble averaging reduces the magnitude of local fluctuations and the strength of the z coupling must be correspondingly larger to result in a large degree of dissipation. The exact degree of dissipation induced by the out-of-plane coupling is additionally influenced by the strength of the planar coupling, as stronger planar coupling results in larger fluctuations in the local magnetization and larger z components of the expectation values of individual spins. Even for highly local interactions, if the planar component of the interaction is sufficiently weak the dissipative effect of the z coupling becomes less prominent.

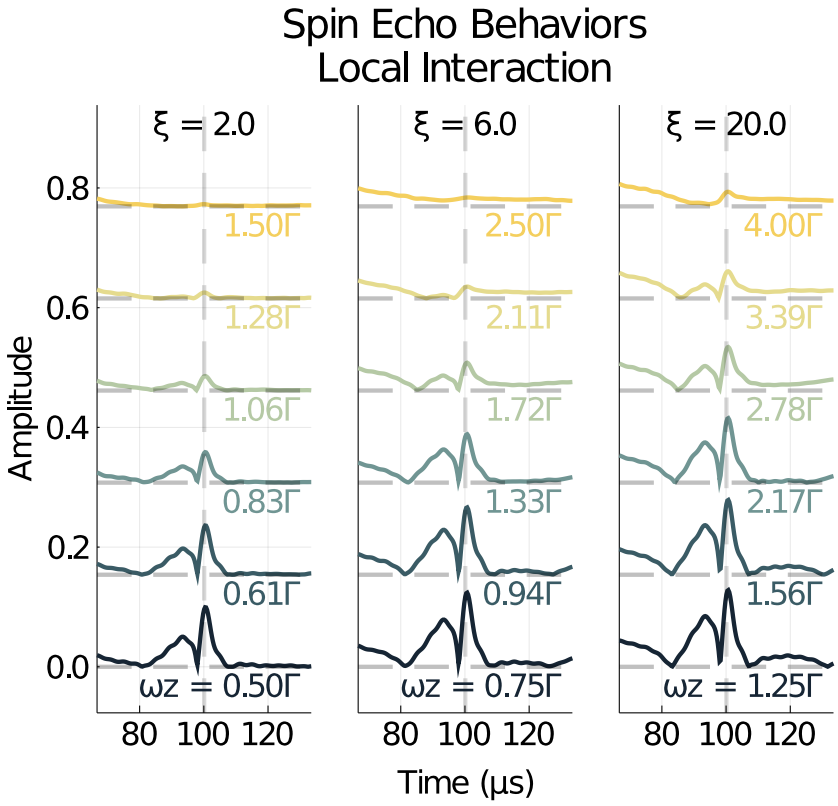

Figure 4.10: Echo behavior for varying ω_z and ξ with $\omega = 3.15\Gamma$. Echo amplitudes are not normalized to accentuate the attenuation from the introduction of ω_z.

We can see from Figure 4.10 that the introduction of even relatively weak ω_z relative to ω for small ξ can result in strong attenuation of the echo. Although a purely in-plane interaction also results in echo attenuation, the effect of the out-of-plane coupling is far stronger and can result in almost total destruction of the echo even at weaker interaction strengths. The same effect can be seen for echoes with longer range ξ, but the dissipative effect is reduced and stronger ω_z is required to cause a similar level of dissipation, as we see for the $\xi = 6.0, 20.0$ echoes in Figure 4.10. For a global interaction the local fluctuations disappear entirely and the presence of non-zero ω_z has no effect unless the spins are flipped at an angle $\theta \neq \pi/2$.

In the spectral domain, we know from the global discussion that we expect to see two primary features: narrowing of the peak from phase locking and hole burning.

In the case of perfect pulsing, the spin echo spectra in Figure 4.11 show notable differences between the highly local interactions and those which are longer range in ways that are consistent with our understanding of the temporal behavior. The echoes for $\xi = 2.0$ show no narrow central peaks because the large degree of dissipation prevents the re-emergence of any locking behavior prior to or during the echo. Rather, it shows strong hole burning effects at all values ω. Echoes with $\xi = 20.0$ show behavior almost identical to the global interaction, and those with $\xi = 6.0$ show a mixing of the two behaviors, with the intermediate ω values beginning to indicate the emergence of a narrow central peak as the phase locking becomes more prominent in the system.

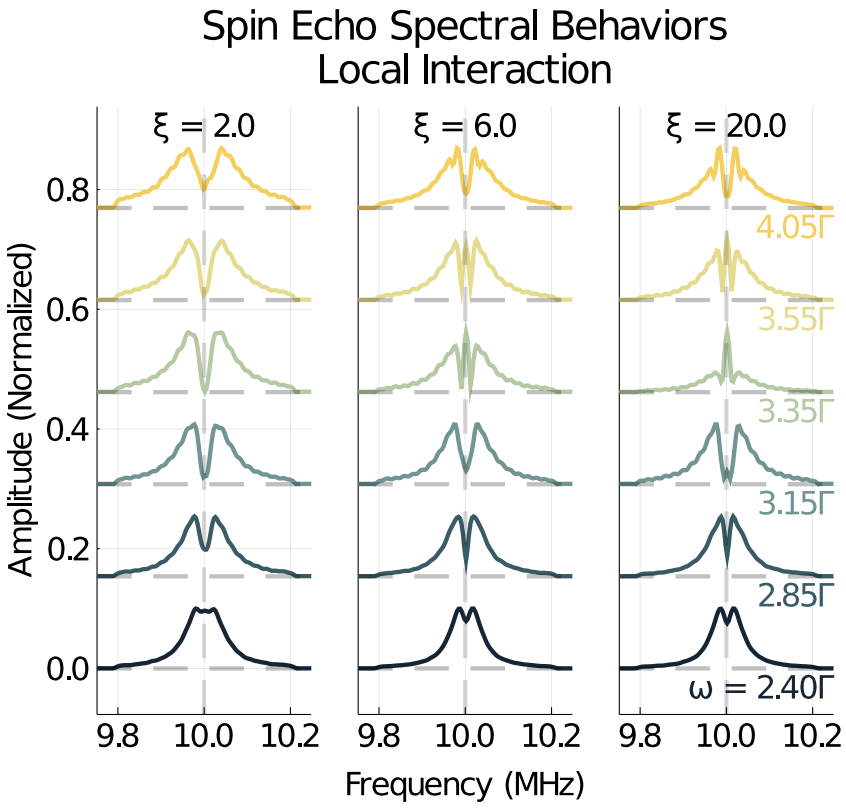

Figure 4.11: Spectral behavior for varying ξ and ω with $\omega_z = 0.0\Gamma$.

Examining the FID for varying ξ offers little more insight. All cases indicate relatively consistent narrowing of the spectrum as expected from the global case. During

the FID, the more local interactions have had little time to dephase on account of local fluctuations, and the phase locking behavior is still relatively strong, preventing any notable differences in the spectral behavior.

The introduction of z coupling introduces additional effects which are difficult to identify in the echo but become more prominent in the spectral domain. In the spectra of the spin echo, the introduction of the z coupling introduces a large amount of attentuation, especially for the more local interactions. Additionally, the dissipation from the z coupling destroys the hole burning effect, as the dramatic reduction in echo amplitude prevents the interaction driven shift at the echo that is responsible for the effect. Instead, we observe the emergence of a narrow peak about the Larmor frequency more indicative of the phase locking behavior seen in the FID. This peak likely arises as a result of the post-pulse decay, which becomes more prominent relative to the echo when the z coupling induced dissipation is very strong. Lacking any pulsing, the FID shows the same narrowing across all localizations observed in the global interaction.

Varying the flip angles for local interactions has many of the same effects we see in the global case. The echo behavior is largely consistent with what we already observe in the global case with the expected variations based on ξ. The intermediate flip angles introduce phase locking behavior in the presence of the planar interaction which competes with the dissipation from the highly local interactions. The presence of the out-of-plane coupling additionally has a similar effect, introducing dissipation and disrupting the balance between the phase locking and hole burning effects.

In the presence of varied pulsing, it is far more informative to examine the spectral behavior. One of the most important spectral features, and the one which gives some of the most significant insight into the interaction's structure, is the pulse dependent shift. However, the magnitude of ω_z to induce a noticeable shift in the spectral domain is far larger than the values of ω_z which can induce strong dissipation in

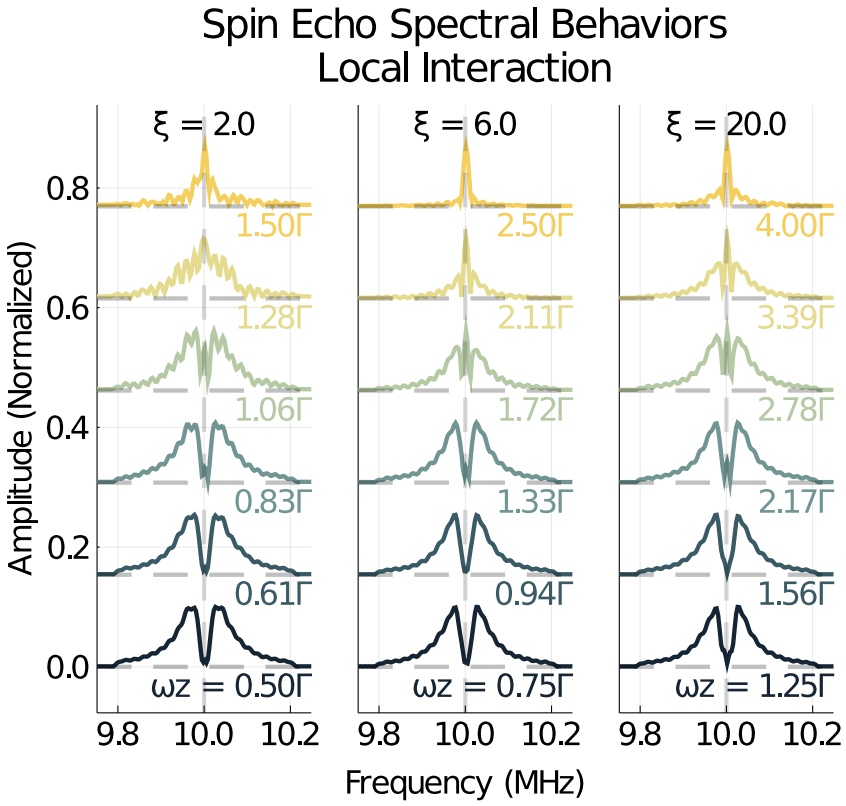

Figure 4.12: Spectral behavior for varying ξ and ω_z with $\omega = 3.15\Gamma$.

the echoes. The dissipation is weaker for smaller flip angles, as might be expected given the strength of the dissipation scales with the strength of the planar interaction and the smaller flip angles effectively reduce the strength of the planar interaction. However, even in the presence of this strong dissipation, valuable information can be gathered from the spectral domain.

The observed spectral behaviors for $\omega_z = 0$ and $\omega = 0$ are essentially identical to the global case– a pulse dependent shift that does not fundamentally alter the hole burning or phase locking effects already seen for that range of ξ. The introduction of both in-plane coupling and out-of-plane coupling also results in behavior very similar to the global case, with the addition of strong attenuation for small ξ on account of the large values of ω_z needed to induce a noticeable shift. At these values of ω_z, the highly local echoes are almost entirely attenuated and the observed spectral behavior

is not a reliable indicator of the dynamics.

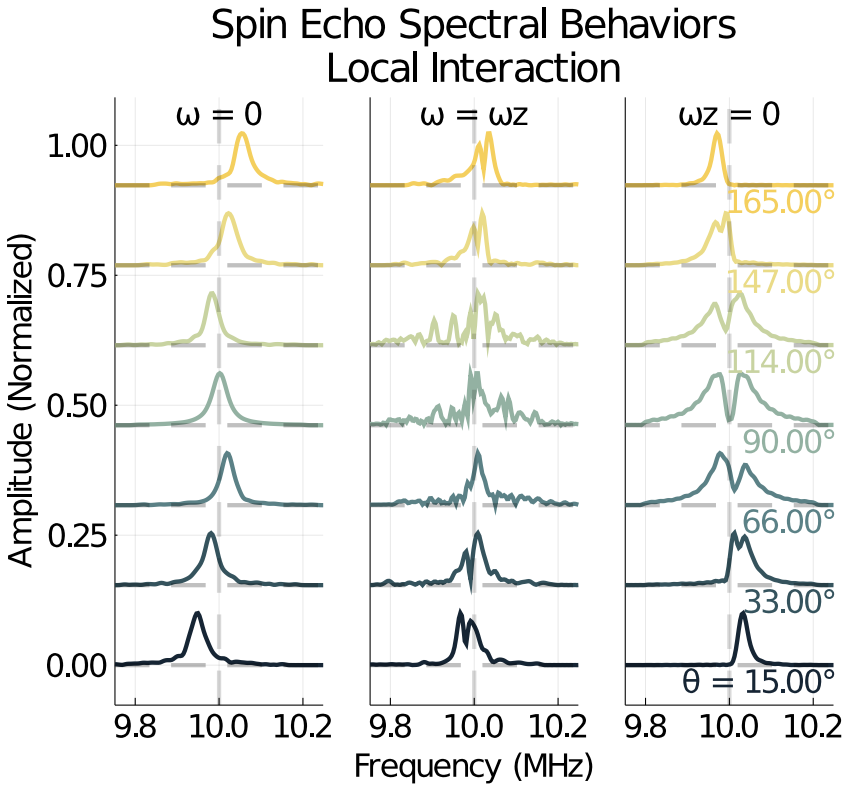

Figure 4.13: Variety of spectral behaviors observed in the spin echo for different flip angles with $\xi = 2.0$ and (ω, ω_z) equal to $(0, 3.15\Gamma)$, $(3.15\Gamma, 3.15\Gamma)$, and $(3.15\Gamma, 0)$.

4.2.1 A Quantum Treatment

Although the geometric interpretation gives a good intuitive understanding of the spin dynamics, we can also derive these effects using some basic quantum mechanics and our understanding of the behavior of the spins. Consider the following spin-½ Hamiltonian

$$H = \frac{1}{2}\begin{pmatrix} -B & -\bar{A} \\ -A & B \end{pmatrix} \tag{4.2}$$

For $A = A_x + iA_y$. This is equivalent to the instantaneous Hamiltonian in first order average Hamiltonian theory (or for $\alpha_x = \alpha_y$) for one of our spins. Treating the Hamiltonian as static, we can find its eigenvalues and eigenvectors

$$\lambda_\pm = \pm\frac{\Lambda}{2}\,, \quad |v_\pm\rangle = \frac{1}{N_\pm}\begin{pmatrix} {}^{(B\mp\Lambda)}/_{A} \\ 1 \end{pmatrix} \tag{4.3}$$

for $\Lambda = \sqrt{|A|^2 + B^2}$. In the following calculation, the normalization constant will cancel so we leave it unspecified. We can calculate the precession frequency of a given spin by finding the time derivatives of the expectation values of the spin. The spin expectation values move on a sphere of radius s and by finding the x and y components of the instantaneous velocity of this point on the sphere we can effectively find the precession frequency. To do so, we use a state in the z basis

$$|\psi\rangle = \alpha\,|+\rangle + e^{i\varphi}\beta\,|\,-\,\rangle \tag{4.4}$$

Since the state can only be specified up to an overall phase, we can assume without loss of generality that α and β are real. The expectation values along each Cartesian axis can be given in terms of α, β and φ

$$\langle I_x\rangle = \alpha\beta\cos\varphi\,, \quad \langle I_y\rangle = \alpha\beta\sin\varphi\,, \quad \langle I_z\rangle = \frac{\alpha^2 - \beta^2}{2} \tag{4.5}$$

By converting this state to the eigenbasis, time evolving, and then returning to the z basis we can find the time dependent state

$$\begin{aligned} |\psi(t)\rangle = &\left[\alpha\cos{}^{\Lambda t}/_{2} + i\beta\sin{}^{\Lambda t}/_{2}\left(\frac{\alpha B}{\Lambda\beta} + \frac{\bar{A}e^{i\varphi}}{\Lambda}\right)\right]|+\rangle \\ &+ e^{i\varphi}\left[\beta\cos{}^{\Lambda t}/_{2} + i\alpha\sin{}^{\Lambda t}/_{2}\left(\frac{e^{-i\varphi}A}{\Lambda} - \frac{\beta B}{\alpha\Lambda}\right)\right]|\,-\,\rangle \end{aligned} \tag{4.6}$$

We note that in the case $A = 0$, this state reduces to the standard state under Larmor precession. In our system, the Hamiltonian is time dependent through the time dependence of the local magnetization (here, A). This time evolved state is only applicable in the $t \to 0$ limit, i.e. in the instantaneous limit where H can be treated as static. From here, we must calculate the expectation values along each Cartesian axis and their time derivatives at $t = 0$. After a fair bit of arithmetic, we arrive at

$$\begin{aligned} \left.\frac{d\langle I_x\rangle}{dt}\right|_{t=0} &= \langle I_y\rangle B - \langle I_z\rangle A_y \\ \left.\frac{d\langle I_y\rangle}{dt}\right|_{t=0} &= -\langle I_x\rangle B + \langle I_z\rangle A_x \\ \left.\frac{d\langle Z\rangle}{dt}\right|_{t=0} &= A_y\langle I_x\rangle - A_x\langle I_y\rangle \end{aligned} \tag{4.7}$$

which are simply the Bloch equations derived from the expectation values of our spin state. This result justifies our claim of the equivalence of the geometric and exact approach in the case of spin-$1/2$. With these derivatives, we can immediately see how the observed spectral behavior emerges from our Hamiltonian. Throughout our echo, the net magnetization (assuming ninety degree pulsing) is along either the y or negative y axis and thus $A_x = 0$. Immediately following the $\pi/2$ pulse, every spin is aligned with the y axis. In the rotating frame, those spins which have $B < 0$, i.e. those with $\Delta\nu < 0$, oscillate towards $\langle X\rangle > 0$ and rotate below the plane, and those with positive $\Delta\nu$ similarly rotate above the plane. In doing so, they gain a non-zero $\langle Z\rangle$ which opposes precession towards the x axis, reducing their effective precession frequency. This effect results in the spectral narrowing we observe during the FID.

Near the echo time, the dynamics change. The interaction only shifts the x component of the planar precession frequency, and therefore only strongly effects those spins nearly aligned with the y axis. As the echo approaches, most spins are highly dispersed about the z axis, but those spins with very small $\Delta\nu$ remain relatively in line with the y axis. Because these spins oscillate very slowly, they only gain a small

z component and tend to remain in a slow orbit closely aligned with their initial axis. However, because of the π pulse, their values of $\langle Z \rangle$ are opposite to what their natural precession would create during the FID. Thus, when the echo occurs and the strength of the interaction (i.e. A_y) becomes strong again, these spins experience the opposite effect as during the FID. The interaction creates a shift in their precession frequency which aligns with their natural frequency, shifting their value of $\Delta\nu$ away from zero, resulting in the hole burning effect. The effect scales with the value of $\Delta\nu$, as the value of $\Delta\nu$ determines how close a given spin remains to the y axis at the time of the echo. The visibility of the effect depends on the likelihood that all spins near a value of $\Delta\nu$ remain relatively close spatially, which only occurs for small values of $\Delta\nu$. Those spins with larger Larmor shifts are more broadly dispersed about the z axis and their individual shifts are too varied to accumulate and create a noticeable effect in the spectrum.

4.2.2 A Note on Dissipation

Before turning to more technical results, we briefly touch on the effects of dissipation on the results we discuss above. Including the dissipation terms from the Linbladian, either from dephasing or relaxation from thermal effects, will generally be destructive of the phase locking effects but will leave the hole burning effects relatively unaffected, assuming the strength of the dissipation is not excessive.

The phase locking depends on a relatively delicate balance of the spin states and interaction strength, as we have seen. With the introduction of dissipation which will cause spins to break away from the locked state, not only reducing the number of phase locked spins but also reducing the strength of the interaction and its ability to keep those remaining spins locked, we expect to see the phase locking behavior become more rare and the hole burning effects of the echo to become more prominent. In essence, additional dissipation will increase the point at which the echo and spectral

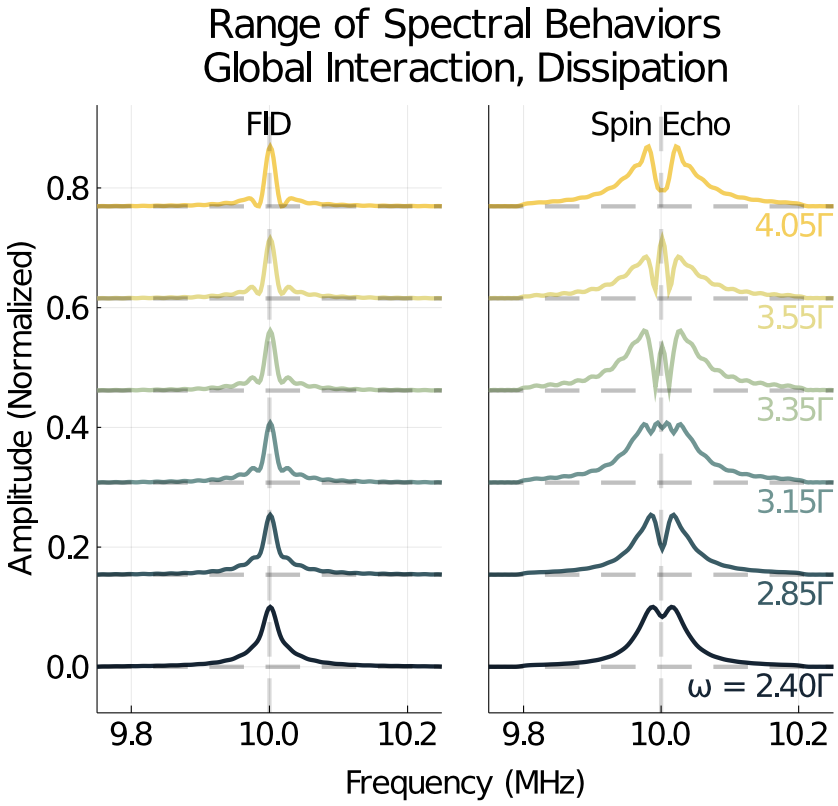

Figure 4.14: Variety of spectral behaviors for a global interaction for a spectrum obtained from the FID and the spin echo in the presence of phase decoherence from dissipative effects ($\Gamma_3 = 0.001$).

behavior approaches the global limit, and longer range interactions will behave more similarly to the very local interactions, with the except of the strong dissipation from out-of-plane coupling which will remain characteristic of the highly localized regime. We see this precise behavior if Figure 4.14, where the strong interactions in the range $\omega = 3.35 - 3.55\Gamma$ show a significant reduction in the strength of the central peak in the echo spectrum indicative of phase locking. The echoes show the same behavior, with the post echo phase locking effect disappearing in the presence of dissipation.

4.3 Γ as a Universal Scale

As we saw in Section 4.1, much of the behavior of the echoes is driven by the relationship between the Larmor shift $\Delta\nu$ and the interaction strength. Because of this close relationship, echoes with different linewidths but the same interaction strength ω will not display the same dynamics. Rather, echoes with a broader spectral line will require a correspondingly higher interaction strength to result in the same dynamics. Additionally, for a very large Γ, the FID and echo decay very rapidly, and the interaction, which relies on the phase coherence of the spins for its strength, is only active very briefly. Thus, even for a relatively strong ω, if the time scale set by Γ is too short, the interaction will not be "active" for any notable length of time, and the echo will only be lightly perturbed. In terms of the phenomenological behavior of the echo, the exact value of ω is not important– rather, the effective value of ω, set by its scale relative to Γ, determines the behavior of the interaction.

This scaling effect results in an approximate equivalence across different values of Γ as long as the scale of ω relative to Γ remains fixed. Additionally, because Γ sets the time scale of the interaction, a rescaling of τ by Γ is necessary. All together, for values of ω, τ which retain the same scale relative to Γ, the spin echo will produce highly similar phenomenological behavior up to some variation on account of local fluctuations from frequency sampling and other numerical effects.

Importantly, we note that this observed equivalence across Γ is not a result of a universal re-scaling all parameters in the simulation "cancelling out". The necessary re-scaling of ω is not merely a re-scaling by the ratio of the linewidths. Scaling ω by the ratio of the linewidths does not produce identical behavior– rather, there are slight offsets and shifts in the scaling that break any numerical equivalence in the comparison, making this result a feature of the dynamics of the echoes and not a feature of the numerics involved in calculating them.

To demonstrate this behavior, we generate echoes for seven different linewidths

from 25kHz up to 100kHz. We additionally check $\Gamma = 10$kHz, but we discuss this case separately. For each linewidth, we simulate echoes with values of ω from $\approx 0.8\Gamma$ up to 4Γ and echo times τ ranging from $1/\Gamma$ up to $10/\Gamma$, although there is some fluctuation in the exact values of ω for different linewidths. We set the time step by $\Gamma dt = 1/10$. By scaling both τ and dt with Γ, all echoes will have the same number of time steps, removing the need for interpolation and the associated potential for additional error.

The error for a given echo is calculated as the average error across the echo relative to the average echo at that interaction strength. For each pair ω, τ, we calculate an average echo averaging the five echoes of different Γ with that same value of ω, τ relative to their value Γ. We then calculate the error of an echo by

$$\varepsilon = \frac{1}{L} \sum_{i=1}^{L} \left| S(t_i) - S_{ref}(t_i) \right| \tag{4.8}$$

The value of ε, averaged over the give Γ values, is displayed in Figure 4.15 vs τ and ω as a fraction of $1/\Gamma$ and Γ, respectively.

From Figure 4.15, we can see not only is the error very small for most of the parameter space examined, but it is also primarily defined by the value of ω and not by τ, although high values of τ do somewhat reduce the ω necessary to accrue more error. We see the error only becomes notable (on the order of 1-10% on average) when ω approaches $3.5\times$ the value of Γ. This regime represents interaction strengths where the local fluctuations from the frequency sampling grow very strong at this sample size. Figure 4.16 shows an overlay of the echoes for all seven values Γ as well as the average echo for $\tau = 2.7/\Gamma$ and $\omega \approx 3\Gamma$. Despite each echo having a different echo time τ in absolute units, the behavior of the echoes is phenomenologically very similar.

Once the linewidth begins to diverge dramatically from those against which it is being compared, the nearly-linear relationship of ω begins to diverge. In the case

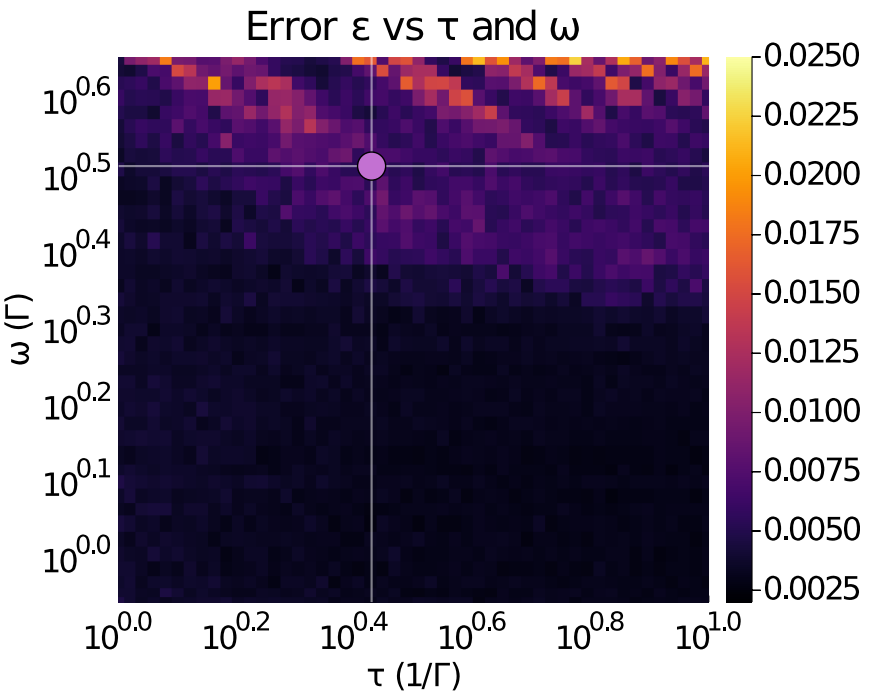

Figure 4.15: Error ε relative to the average echo shape for a given τ, ω, averaged over Γ, as a function of τ and ω. The purple marker indicates the value of ω and τ that corresponds to Figure 4.16.

of $\Gamma = 10$kHz, the behavior of the echoes at similar ω was noticeably different. This breakdown is emblematic of the imperfect nature of ω, in that it represents a constant, maximal value of the interaction and does not capture the true dynamics of the interaction. The approximation of the scaling only holds so long as the linewidths, relative to $\Gamma = 50$kHz, are within a range of $1/2 - 2\times$, which represents a range of linewidths for which a standard single shot spin echo would be used.

Because of the complicated relationship between ω, α, and ξ, the equivalence of the behavior of the echoes for different Γ is difficult to quantify. Given that ω is not an exact parameter of the model, strict scaling of ω does not guarantee equivalent behavior, although echoes with similar behavior will have similar values of ω relative to Γ. However, relatively small changes in ω (e.g. a shift of a few kHz) can result in noticeably different behavior. Ultimately, relative to the behaviors set by some reference Γ (here, we use 50kHz), there is a linear relationship between the reference

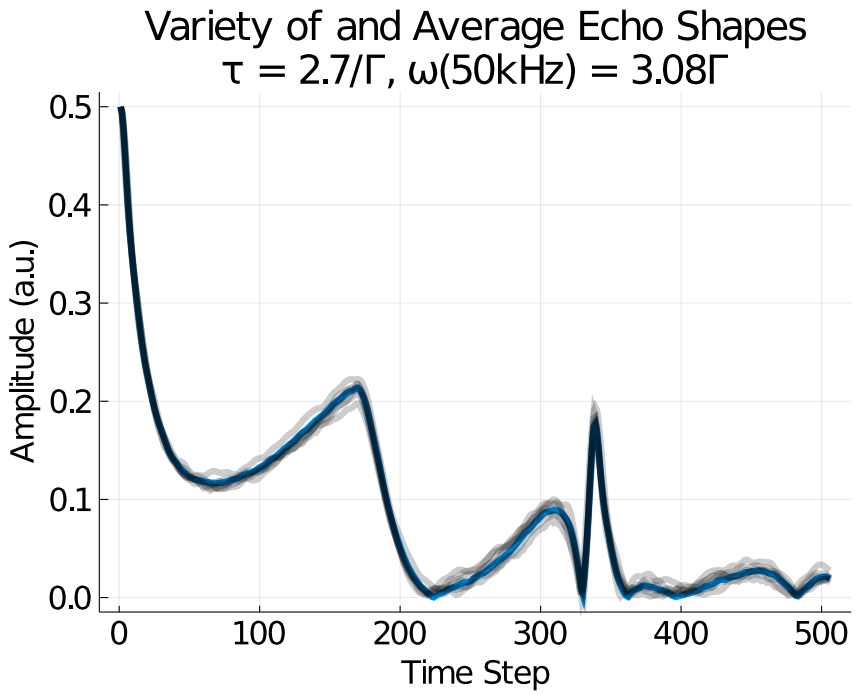

Figure 4.16: Overlay of echo shapes for $\omega \approx 3\Gamma$ and $\tau = {}^{2.7}/_{\Gamma}$. The average echo shape is displayed in blue. Each echo has a different τ in absolute units, but the same number of time points. The echoes are therefore plotted against time step and not time to enable a closer comparison.

ω (i.e. that of the 50kHz echo) and the ω of equivalent behavior for some other Γ. However, the details of that linear relationship depend on Γ.

In order to quantify the drift in ω values with Γ, we can use some basic fitting to generate a function $\Omega(\omega, \Gamma)$. This function takes the ω of a given behavior for $\Gamma = 50$kHz and shifts it appropriately to give the proper ω to get equivalent behavior at some other Γ. The form of $\Omega(\omega, \Gamma)$ is linear, with coefficients depending on the value of Γ.

$$\Omega(\omega, \Gamma) = m(\Gamma)\omega + b(\Gamma) \tag{4.9}$$

Although the relationship between ωs at different Γs is very nearly a direct scaling (i.e. $\omega_f = ({}^{\Gamma_f}/_{\Gamma_0})\omega_0$), the small shifts in both slope and intercept of Equation 4.9 play important roles in the dynamics of the observed echoes due to the scale of

the parameters at play and therefore such direct scaling is insufficient to capture equivalent behavior. The coefficients m and b can be fit relative to Γ, as well. The slope $m(\Gamma)$ depends linearly on Γ

$$m(\Gamma) = m_1\Gamma + m_2 \tag{4.10}$$

but the intercept $b(\Gamma)$ has a complicated functional relationship to Γ and is difficult to precisely fit. It can be fit reasonably well with a power law fit

$$b(\Gamma) = b_1\Gamma^{b_2} + b_3 \tag{4.11}$$

The fit coefficients m_i and b_i are presented in Table 4.1 for reference. All fits have R^2 values greater than 0.999. A function to calculate the appropriate shift is included in the repository containing the simulation code [83].

Coefficient	Value	Margin of Error
m_1	1.92×10^{-5}	5.33×10^{-9}
m_2	0.0369	4.19×10^{-4}
b_1	0.00404	3.70×10^{-5}
b_2	-3.72×10^{-9}	3.58×10^{-10}
b_3	1.29	0.0080

Table 4.1: Fit Coefficients for $b(\Gamma)$, $m(\Gamma)$.

The higher error in the fit coefficients of $b(\Gamma)$ arise from the difficulty in determining an accurate functional form for the data. A power law fitting was ultimately found be well fitting and give lower uncertainty in the coefficients than other functional forms.

We can use this fitting to predict the appropriate value of ω to recreate the phenomenological behavior of the behavior seen at a known ω at $\Gamma = 50kHz$. Because of the approximate nature of the fitting, the observed behaviors show some variation, but they retain strong phenomenological similarity, to the point where only very small

tuning of ω is required to acquire more accurate behavior. To demonstrate this, we generate nine random linewidths between 25kHZ and 100kHZ and generate echoes using the relationship $\Omega(\omega, \Gamma)$ to set their interaction parameters. We plot these echoes, as well as the $\Gamma = 50$kHz echoes used as reference, in Figure 4.17.

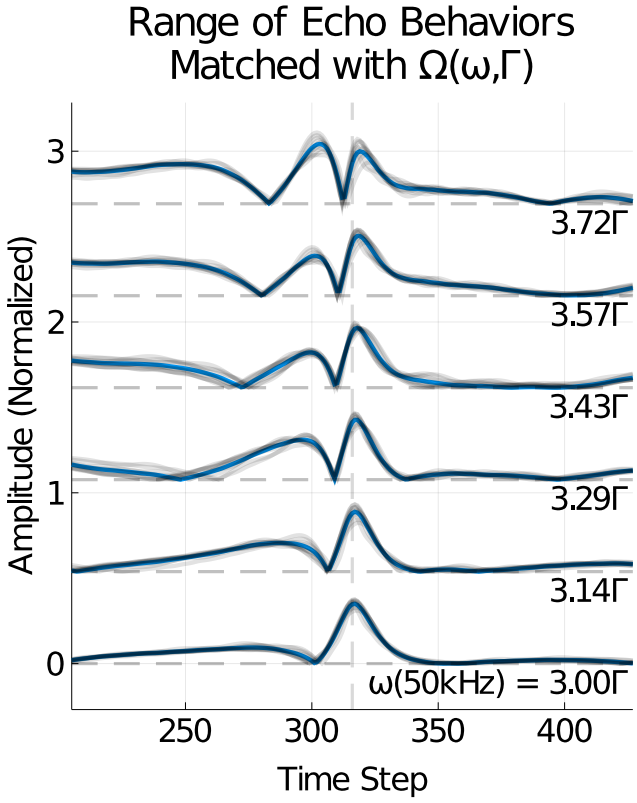

Figure 4.17: Overlay of echo shapes for ω predicted by $\Omega(\omega, \Gamma)$ relative to $\Gamma = 50$kHz. The average echo shape is displayed in blue. Each echo has a different τ in μs, but the same number of time points. Echoes use $n = 400 \times 400$ and $\xi = 6.0$.

We can use this feature of the model to generalize results from a single linewidth Γ systems with other linewidths. These results not only generalize the dynamics of the model and set a universal scale for its parameters, but they also significantly reduce the computation time necessary to examine a broad spectrum of behaviors, as scanning over different linewidths is no longer necessary to examine the dynamics of the model.

4.4 Application to Experiment

Having explored the dynamics of the interaction across a variety of interaction ranges ξ and strengths ω, we can begin to see the applications of this model to a real system. In order to observe any significant asymmetry in the echo, an interaction strength $\omega > 2.5\Gamma$ is typically required. The strong multi-peak behavior only begins to emerge when $\omega > 3\Gamma$. For extremely small ξ, such behaviors would only occur for extremely strong coupling. At $\xi = 2.0$ and $\Gamma = 50$kHz, the required coupling strength $\alpha \approx 13$kHz. These strong required couplings, combined with the dissipative effects of α_z, the fact that the mean field approximation is far weaker at these small ξ, and the increasing relevance of standard NMR terms like the J coupling, makes the observation of our observed behavior at small ξ very unlikely. In the intermediate range, however, when the echo and spectral behavior still diverges from the global limit, the mean field approximation remains strong, nearest- and next nearest-neighbor terms do not dominate, and the dissipative effects are mitigated at the interaction strengths necessary to induce the notable pulse-dependent shifts. Additionally, the required interaction strength is significantly reduced on account of the larger number of spins which interact. For $\xi = 6 - 8$, a coupling of only around $\alpha = 1$kHz is necessary. As ξ approaches the long range regime $\xi \to 15 - 20$, the coupling needs only be on the order of 100Hz. We find these regimes more likely, given the weak initial strength of the nuclear spin coupling and the fact that the dynamics are driven by long range correlations which result in a dramatic increase in the effective interaction strength. In a scenario where such behavior is experimentally observed, we expect it is driven by long range correlations, where $\xi > 10.0$ and α need only be a fraction of the linewidth, anywhere from 1% to 0.1% of Γ depending on the range of the interaction. In the case of the pulse dependent shifts observed in FeSe, a comparison to our results can yield a preliminary approximation of the type of interaction necessary to induce this behavior [4]. Without knowledge of the exact power necessary to result in a $\pi/2$ pulse,

our ability to approximate the shift's pulse dependence is limited. The relative shifts and significant increase in noise between alignments with the [100] axis relative to the [110] axis would be indicative of an anisotropic coupling with intermediate range and a [100] component notably stronger than the [110] coupling. We emphasize this prediction represents a rough approximation and demonstration of the model's application. The presence of multiple peaks in the FID spectrum for the [110] alignment is not in agreement with the predictions of our model, and our predictions merely suggest that a closer examination of this behavior may be merited.

With this understanding of how the spectral and echo behavior varies with regard to the interaction's strength and range, as well as how it varies for different pulsing and different relative values of in-plane and out-of-plane coupling, we can see how a knowledge of such behavior can be used to examine echoes for which the parameters are not known. As we mentioned before, the most valuable tool in identifying the interaction's strength is the pulse dependent shift. By adjusting the orientation of the crystal axes of the sample relative to B_0 and measuring magnitude of the pulse dependent shift, one can extract the differences between different pairings of the interaction strengths. The ability to measure the relative strengths of the interaction along its different axes using the pulse dependent shift for different orientations gives a potent and novel tool for the extraction of information about the anisotropy of the spin susceptibility. Combined with additional measurements of the anisotropy of the hyperfine interactions, this approach can be used to extract previously unavailable information about the electronic structure and give valuable insight into the structure of strongly correlated systems where such behaviors emerge.

Chapter 5

Conclusions and Future Research

Nuclear Magnetic Resonance represents a powerful tool for the study of strongly correlated electrons and the exotic dynamics they exhibit. Its sensitivity to the spin suscepbility through the Knight shift and the spin-lattice relaxation time T_1 place it somewhat uniquely in a position to probe specifically information about the spin state of the electrons. In circumstances where the spin state of strongly correlated electronic modes, such as in unconventional superconductivity, this ability can prove invaluable. Indeed, NMR has been used widely to probe for the existence of unconventional superconductivity, although to varying degrees of success [17–21, 29, 32–36]. These unconventional states are highly susceptible to a wide array of experimental parameters, which makes them difficult to probe without disturbing or destroying the state itself.

In the search for evidence of unconventional phases, NMR has revealed a wide array of unusual signatures, most notably in a pulse dependence of the Knight shift. Although in some situations, such as in Sr_2RuO_4, this dependence likely emerges as a result of heating effects, in others such as FeSe, where the pulse power is tightly constrained to avoid the influence of heating, the effect may be more experimentally significant [4]. Understanding the origin of this effect can give insight into the dynam-

ics of strongly correlated electrons and how they influence the nuclear spin system and also help design new methods of analysis of these systems.

We developed a simple model of a strongly correlated electronic mode interacting with the nuclear ensemble in order to examine the ways such a mode would appear in the echo behavior and spectral domain. We used a mean field approximation with a localization function to create an effective range of the interactions. Using the rotating frame, we examined the explicit time dependence and used the Magnus expansion to quantify the ways in which the rapidly oscillating terms affect the Hamiltonian's dynamics. We modelled the system in Liouville space in order to include dissipative effects which would be prominent in a real system and can have a notable effect on the induced coherence from strong interactions of the model.

In order to simulate the system, we created a custom simulation package. Very few simulations exist which are suitable for our model, or for studying dynamic spin systems in general. Those that do are unsuitable for our model and for NMR more generally [64]. We built several iterations of the simulation, with the final version being written totally in CUDA. We examined the structure of the simulation and checked the accuracy of its custom algorithms against well known and well trusted versions used in Julia. We bench-marked the accuracy of the simulation using Julia, getting extremely strong agreement up to the limits of precision possible in data transfer between CUDA and Julia. Performance in the CUDA simulation is dramatically improved, showing two orders of magnitude of improvement over a purely CPU simulation and an almost 25 times speed-up over a hybrid simulation. Finally, we charted out the convergence of the model with regards to the ensemble size and time step to set limits on the values we can use for these parameters in order to keep error at an acceptable level.

We examined the behavior of the model in both the time domain and spectral domain, with a focus on the spectral behavior. We showed how this type of interaction

can induce multi-peak behavior in the echo and how it results in a pulse dependent shift remarkably similar to that observed in candidates for exotic superconductors [4]. Despite using a simple model, our results provide strong evidence that the presence of such a mode may be responsible for the emergence of these behaviors. These results do not demonstrate definitive proof nor refute current findings, but rather they provide compelling evidence for the possibility of an alternate explanation. In the context of these findings, and the ways such a strongly correlated mode can induce these behaviors, a careful re-examination of the pulse dependence of the Knight shift in these exotic materials may be merited, particularly in the context of a new experimental protocol based on our results: varying the pulse to study a range of flip angles from 15° to 165° and adjusting crystal axes to examine the shift in each configuration. In the presence of a mode similar to the one we model, such a protocol could be used not only to extract the relative strength of the interaction based on the magnitude of the shift but also the anisotropy of the electronic susceptibility. In this way, we laid out the framework for a new experimental protocol which can be used in the search for strongly correlated electronic modes.

Despite the already strong results gained from the studying of this model using our custom simulation package, there is much still to be done. From a physical perspective, a thorough study of the effects of dissipation is the natural next step. Using the Lindblad framework we have already built, the inclusion of thermal effects or the introduction of dephasing can be added and the ways these features change the behavior of the model can be studied. The simulation as is can be used to model a wide variety of behaviors in the mean field limit, and the introduction of new elements such as antiferromagnetic coupling can be easily included. New localization schemes or different types of coupling to the mode can be introduced, such as cross-couplings. The implementation of new pulse schemes and experimental protocols could be used to see how the interaction's effects change under different experimental settings.

From a computational side, the simulation in many ways is still in its infancy. Many of the solvers and algorithms are very basic and only suitable for low dimensional systems. An easy source of improvement is implementing a more sophisticated and better scaling solver than the Gaussian Elimination process currently used. The use of atomic functions to improve processes that currently run in serial would result in additional gains in efficiency. More nuanced parallelization schemes could be introduced, splitting computations for a single spin up across multiple threads to improve performance. Multi-GPU parallelization would break the ensemble size barrier introduced by the limited number of threads available to the simulation currently, although this would also require additional memory management that may decrease the gains in efficiency.

From a broader perspective, the use of CUDA to accelerate simulation of dynamic spin systems still has a large amount of room to grow. CUDA has many built-in highly optimized large matrix operations and solvers which could potentially be applied to model sets of spins interacting in the product space with a high degree of efficiency. With appropriate spectral averaging techniques, simulations of clusters of spins interacting with an electronic mode in the full product space could give yet another tool to examine the effects of strongly correlated electrons in NMR and examine the limits of the mean field approaches. These approaches could be combined with more sophisticated approaches to the Magnus expansion, much like outlined in [64, 91]. The limits of a fully CUDA driven approach to dynamic simulations of spin systems have yet to be fully realized.

Bibliography

[1] Peter Fulde and Richard A. Ferrell. Superconductivity in a strong spin-exchange field. *Phys. Rev.*, 135:A550–A563, Aug 1964.

[2] Henri Alloul. NMR in strongly correlated materials. *Scholarpedia*, 10(1):30632, 2015.

[3] Yuji Matsuda and Hiroshi Shimahara. Fulde–ferrell–larkin–ovchinnikov state in heavy fermion superconductors. *Journal of the Physical Society of Japan*, 76(5):051005, may 2007.

[4] I. Vinograd, S. P. Edwards, Z. Wang, T. Kissikov, J. K. Byland, J. R. Badger, V. Taufour, and N. J. Curro. Inhomogeneous knight shift in vortex cores of superconducting fese. *Phys. Rev. B*, 104:014502, Jul 2021.

[5] W. D. Knight. Nuclear magnetic resonance shift in metals. *Phys. Rev.*, 76:1259–1260, Oct 1949.

[6] L. C. Hebel and C. P. Slichter. Nuclear spin relaxation in normal and superconducting aluminum. *Phys. Rev.*, 113:1504–1519, Mar 1959.

[7] Anton Potocnik. Superconductivity as observed by magnetic resonance. 2010.

[8] Y. Maeno, H. Hashimoto, K. Yoshida, S. Nishizaki, T. Fujita, J. G. Bednorz, and F. Lichtenberg. Superconductivity in a layered perovskite without copper. *Nature*, 372(6506):532–534, Dec 1994.

[9] T M Rice and M Sigrist. Sr2ruo4: an electronic analogue of 3he? *Journal of Physics: Condensed Matter*, 7(47):L643–L648, nov 1995.

[10] Takuji Nomura and Kosaku Yamada. Perturbation theory of spin-triplet superconductivity for sr 2ruo 4. *Journal of the Physical Society of Japan*, 69(11):3678–3688, 2000.

[11] Q. H. Wang, C. Platt, Y. Yang, C. Honerkamp, F. C. Zhang, W. Hanke, T. M. Rice, and R. Thomale. Theory of superconductivity in a three-orbital model of sr2ruo4. *EPL (Europhysics Letters)*, 104(1):17013, oct 2013.

[12] Andrew Mackenzie, Thomas Scaffidi, Clifford Hicks, and Yoshiteru Maeno. Even odder after twenty-three years: The superconducting order parameter puzzle of sr2ruo4. *npj Quantum Materials*, 2, 06 2017.

[13] Yoshiteru Maeno, Shunichiro Kittaka, Takuji Nomura, Shingo Yonezawa, and Kenji Ishida. Evaluation of spin-triplet superconductivity in sr2ruo4. *Journal of the Physical Society of Japan*, 81(1):011009, 2012.

[14] Catherine Kallin. Chiral p-wave order in sr2ruo4. *Reports on Progress in Physics*, 75(4):042501, mar 2012.

[15] Ying Liu and Zhi-Qiang Mao. Unconventional superconductivity in sr2ruo4. *Physica C: Superconductivity and its Applications*, 514:339–353, 2015. Superconducting Materials: Conventional, Unconventional and Undetermined.

[16] K. D. Nelson, Z. Q. Mao, Y. Maeno, and Y. Liu. Odd-parity superconductivity in sr2ruo4. *Science*, 306(5699):1151–1154, 2004.

[17] H Mukuda, K Ishida, Y Kitaoka, K Asayama, Z. q Mao, Y Mori, and Y Maeno. 17o nmr probe of spin fluctuations in triplet superconductor sr2ruo4. *Physica B: Condensed Matter*, 259-261:944–946, 1999.

[18] K. Ishida, Y. Kitaoka, K. Asayama, S. Ikeda, S. Nishizaki, Y. Maeno, K. Yoshida, and T. Fujita. Anisotropic pairing in superconducting sr_2ruo_4 : ru nmr and nqr studies. *Phys. Rev. B*, 56:R505–R508, Jul 1997.

[19] Kenji Ishida, Hiroshi Murakawa, Hidekazu Mukuda, Yoshio Kitaoka, Z.Q. Mao, and Yoshiteru Maeno. Nmr and nqr studies on superconducting sr2ruo4. *Journal of Physics and Chemistry of Solids*, 69(12):3108–3114, 2008. SNS2007.

[20] Kenji Ishida, Masahiro Manago, Katsuki Kinjo, and Yoshiteru Maeno. Reduction of the 17o knight shift in the superconducting state and the heat-up effect by NMR pulses on sr2ruo4. *Journal of the Physical Society of Japan*, 89(3):034712, mar 2020.

[21] A. Pustogow, Yongkang Luo, A. Chronister, Y.-S Su, Dmitry Sokolov, F. Jerzembeck, A. Mackenzie, C. Hicks, N. Kikugawa, S. Raghu, E. Bauer, and S. Brown. Constraints on the superconducting order parameter in sr2ruo4 from oxygen-17 nuclear magnetic resonance. *Nature*, 574:1–4, 10 2019.

[22] A. Bianchi, R. Movshovich, C. Capan, P. G. Pagliuso, and J. L. Sarrao. Possible fulde-ferrell-larkin-ovchinnikov superconducting state in $cecoin_5$. *Phys. Rev. Lett.*, 91:187004, Oct 2003.

[23] Shugo Ikeda, Hiroaki Shishido, Miho Nakashima, Rikio Settai, Dai Aoki, Yoshinori Haga, Hisatomo Harima, Yuji Aoki, Takahiro Namiki, Hideyuki Sato, and Yoshichika Onuki. Unconventional superconductivity in cecoin5 studied by the specific heat and magnetization measurements. *Journal of the Physical Society of Japan*, 70(8):2248–2251, 2001.

[24] C Petrovic, P G Pagliuso, M F Hundley, R Movshovich, J L Sarrao, J D Thompson, Z Fisk, and P Monthoux. Heavy-fermion superconductivity in cecoin5 at 2.3 k. *Journal of Physics: Condensed Matter*, 13(17):L337–L342, apr 2001.

[25] R. Movshovich, M. Jaime, J. D. Thompson, C. Petrovic, Z. Fisk, P. G. Pagliuso, and J. L. Sarrao. Unconventional superconductivity in CeIrIn_5 and cecoin$_5$: Specific heat and thermal conductivity studies. *Phys. Rev. Lett.*, 86:5152–5155, May 2001.

[26] N. J. Curro, B. Simovic, P. C. Hammel, P. G. Pagliuso, J. L. Sarrao, J. D. Thompson, and G. B. Martins. Anomalous nmr magnetic shifts in cecoin$_5$. *Phys. Rev. B*, 64:180514, Oct 2001.

[27] M. Kenzelmann, S. Gerber, N. Egetenmeyer, J. L. Gavilano, Th. Strässle, A. D. Bianchi, E. Ressouche, R. Movshovich, E. D. Bauer, J. L. Sarrao, and J. D. Thompson. Evidence for a magnetically driven superconducting q phase of cecoin$_5$. *Phys. Rev. Lett.*, 104:127001, Mar 2010.

[28] K. Izawa, H. Yamaguchi, Yuji Matsuda, H. Shishido, R. Settai, and Y. Onuki. Angular position of nodes in the superconducting gap of quasi-2d heavy-fermion superconductor cecoin$_5$. *Phys. Rev. Lett.*, 87:057002, Jul 2001.

[29] V.F. Mitrović, G. Koutroulakis, M.-A. Vachon, M. Horvatić, C. Berthier, G. Lapertot, and J. Flouquet. Phase diagram of cecoin5 in the vicinity of hc2 as determined by NMR. *Physica B: Condensed Matter*, 403(5-9):986–989, apr 2008.

[30] V. F. Mitrović, M. Horvatić, C. Berthier, G. Knebel, G. Lapertot, and J. Flouquet. Observation of spin susceptibility enhancement in the possible fulde-ferrell-larkin-ovchinnikov state of cecoin$_5$. *Phys. Rev. Lett.*, 97:117002, Sep 2006.

[31] B.-L. Young, R. R. Urbano, N. J. Curro, J. D. Thompson, J. L. Sarrao, A. B. Vorontsov, and M. J. Graf. Microscopic evidence for field-induced magnetism in cecoin$_5$. *Phys. Rev. Lett.*, 98:036402, Jan 2007.

[32] K. Kakuyanagi, M. Saitoh, K. Kumagai, S. Takashima, M. Nohara, H. Takagi, and Y. Matsuda. Texture in the superconducting order parameter of $cecoin_5$ revealed by nuclear magnetic resonance. *Phys. Rev. Lett.*, 94:047602, Feb 2005.

[33] K. Kumagai, K. Kakuyanagi, M. Saitoh, S. Takashima, M. Nohara, H. Takagi, and Y. Matsuda. Microscopic evidence of the fflo state in cecoin5 probed by nmr. *Physica B: Condensed Matter*, 378-380:347–350, 2006. Proceedings of the International Conference on Strongly Correlated Electron Systems.

[34] K. Kumagai, K. Kakuyanagi, M. Saitoh, T. Oyaizu, Y. Matsuda, S. Takanashi, M. Nohara, and H. Takagi. Nmr study of the novel superconducting state in cecoin5. *Journal of Magnetism and Magnetic Materials*, 310(2, Part 1):715–717, 2007. Proceedings of the 17th International Conference on Magnetism.

[35] Ken–ichi Kumagai, T. Oyaizu, Yuji Furukawa, Hiroaki Shishido, and Yuji Matsuda. Nmr study of a fflo state of cecoin5 in a perpendicular and a parallel field. *Physica B-condensed Matter*, 403:1144–1146, 2008.

[36] V. F. Mitrović, G. Koutroulakis, M. Klanjšek, M. Horvatić, C. Berthier, G. Knebel, G. Lapertot, and J. Flouquet. Comment on "texture in the superconducting order parameter of cecoin5 revealed by nuclear magnetic resonance". *Physical Review Letters*, 101(3), jul 2008.

[37] PERCH Solutions Ltd. Perch nmr software.

[38] Tim Claridge. Software review of mnova: Nmr data processing, analysis, and prediction software. *Journal of Chemical Information and Modeling*, 49(4):1136–1137, 2009. PMID: 19326937.

[39] Yuri Binev, Maria M. B. Marques, and João Aires-de Sousa. Prediction of 1h nmr coupling constants with associative neural networks trained for chemical

shifts. *Journal of Chemical Information and Modeling*, 47(6):2089–2097, 2007. PMID: 17956083.

[40] Charles D. Schwieters and G.Marius Clore. The vmd-xplor visualization package for nmr structure refinement. *Journal of Magnetic Resonance*, 149(2):239–244, 2001.

[41] Frédéric A. Perras, Cory M. Widdifield, and David L. Bryce. Quest—quadrupolar exact software: A fast graphical program for the exact simulation of nmr and nqr spectra for quadrupolar nuclei. *Solid State Nuclear Magnetic Resonance*, 45-46:36–44, 2012.

[42] Denimar Possa, Anderson C. Gaudio, and Jair C.C. Freitas. Numerical simulation of nqr/nmr: Applications in quantum computing. *Journal of Magnetic Resonance*, 209(2):250–260, 2011.

[43] Davide Candoli, Samuele Sanna, and Vesna F. Mitrović. Pulsee: A software for the quantum simulation of an extensive set of nuclear magnetic and quadrupole resonance observables, 2021.

[44] Mikhail Veshtort and Robert G. Griffin. Spinevolution: A powerful tool for the simulation of solid and liquid state nmr experiments. *Journal of Magnetic Resonance*, 178(2):248–282, 2006.

[45] Mads Bak, Jimmy T Rasmussen, and Niels Chr Nielsen. Simpson: A general simulation program for solid-state nmr spectroscopy. *Journal of Magnetic Resonance*, 147(2):296–330, 2000.

[46] H.J. Hogben, M. Krzystyniak, G.T.P. Charnock, P.J. Hore, and Ilya Kuprov. Spinach – a software library for simulation of spin dynamics in large spin systems. *Journal of Magnetic Resonance*, 208(2):179–194, 2011.

[47] Wyndham B. Blanton. Blochlib: a fast nmr c++ tool kit. *Journal of Magnetic Resonance*, 162(2):269–283, 2003.

[48] Christian Bengs and Malcolm H. Levitt. Spindynamica: Symbolic and numerical magnetic resonance in a mathematica environment. *Magnetic Resonance in Chemistry*, 56(6):374–414, 2018.

[49] J.R. Johansson, P.D. Nation, and Franco Nori. Qutip: An open-source python framework for the dynamics of open quantum systems. *Computer Physics Communications*, 183(8):1760–1772, 2012.

[50] Wolfram Research, Inc. Mathematica, Version 13.0.0. Champaign, IL, 2021.

[51] Guido Van Rossum and Fred L. Drake. *Python 3 Reference Manual*. CreateSpace, Scotts Valley, CA, 2009.

[52] H. J. Reich. Windnmr-pro.

[53] A. Allouche and G. Pouzard. Computer simulation of ft-nmr multiple pulse experiment. *Computer Physics Communications*, 54(1):171–176, 1989.

[54] Martin Weigel. Performance potential for simulating spin models on gpu. *Journal of Computational Physics*, 231(8):3064–3082, 2012.

[55] Kenneth Esler, Jeongnim Kim, David Ceperley, and Luke Shulenburger. Accelerating quantum monte carlo simulations of real materials on gpu clusters. *Computing in Science Engineering*, 14(1):40–51, 2012.

[56] Morad Abdelaziz. Gpu-opencl accelerated probabilistic power flow analysis using monte-carlo simulation. *Electric Power Systems Research*, 147:70–72, 2017.

[57] H. Wan Chan Tseung, J. Ma, and C. Beltran. A fast gpu-based monte carlo simulation of proton transport with detailed modeling of nonelastic interactions. *Medical Physics*, 42(6Part1):2967–2978, 2015.

[58] Benjamin Block, Peter Virnau, and Tobias Preis. Multi-gpu accelerated multi-spin monte carlo simulations of the 2d ising model. *Computer Physics Communications*, 181(9):1549–1556, 2010.

[59] Tobias Preis, Peter Virnau, Wolfgang Paul, and Johannes J. Schneider. Gpu accelerated monte carlo simulation of the 2d and 3d ising model. *Journal of Computational Physics*, 228(12):4468–4477, 2009.

[60] Axel D. Dente, Carlos S. Bederián, Pablo R. Zangara, and Horacio M. Pastawski. Gpu accelerated trotter-suzuki solver for quantum spin dynamics, 2013.

[61] Pui-Wai Ma, S.L. Dudarev, and C.H. Woo. Spilady: A parallel cpu and gpu code for spin–lattice magnetic molecular dynamics simulations. *Computer Physics Communications*, 207:350–361, 2016.

[62] Dan Abell, D. Meiser, D.P. Barber, Mengjie Bai, and Vahid Ranjbar. Gpu-accelerated spin dynamics and analysis for rhic. *IPAC 2013: Proceedings of the 4th International Particle Accelerator Conference*, pages 1037–1039, 01 2013.

[63] Paul Thomas Heistracher. *Atomistic Spin Dynamics*. PhD thesis, Vienna University of Technology, 2017.

[64] Alex Tritt, Joshua Morris, Joel Hochstetter, R. P. Anderson, James Saunderson, and L. D. Turner. Spinsim: a gpu optimized python package for simulating spin-half and spin-one quantum systems, 2022.

[65] Charles P. Slichter. *Principles of Magnetic Resonance*. Springer-Verlag, 1990.

[66] Georgios Koutroulakis. *Unravelling the mysteries of unconventional superconductivity with NMR: The curious case of CeCoIn5*. PhD thesis, Brown University, 2010.

[67] Tôru Moriya. The effect of electron-electron interaction on the nuclear spin relaxation in metals. *Journal of the Physical Society of Japan*, 18(4):516–520, 1963.

[68] J. Korringa. Nuclear magnetic relaxation and resonnance line shift in metals. *Physica*, 16(7):601–610, 1950.

[69] S. Carr, C. Snider, E. Barrett, D. E. Feldman, C. Ramanathan, J. B. Marston, and V. F. Mitrović. Signatures of spin-susceptibility anisotropy and electronic nematicity in nuclear magnetic resonance. unpublished, 2021.

[70] Kei Yosida. Magnetic properties of cu-mn alloys. *Phys. Rev.*, 106:893–898, Jun 1957.

[71] M. A. Ruderman and C. Kittel. Indirect exchange coupling of nuclear magnetic moments by conduction electrons. *Phys. Rev.*, 96:99–102, Oct 1954.

[72] Wilhelm Magnus. On the exponential solution of differential equations for a linear operator. *Communications on Pure and Applied Mathematics*, 7(4):649–673, 1954.

[73] Eugene Stephane Mananga and Thibault Charpentier. On the floquet–magnus expansion: Applications in solid-state nuclear magnetic resonance and physics. *Physics Reports*, 609:1–49, 2016. On the Floquet–Magnus expansion: Applications in solid-state nuclear magnetic resonance and physics.

[74] E. Fel'dman, Dmitri Feldman, and E. Kuznetsova. Floquet hamiltonian and entanglement in spin systems in periodic magnetic fields. *Applied Magnetic Resonance*, 48, 06 2017.

[75] Elena Kuznetsova, E. Fel'dman, and Dmitri Feldman. Paradoxes of the magnus expansion at investigations of an equilibrium magnetization and entanglement in the multi-pulse spin locking. *Uspekhi Fizicheskih Nauk*, 186, 01 2016.

[76] E.B. Fel'dman. On the convergence of the magnus expansion for spin systems in periodic magnetic fields. *Physics Letters A*, 104(9):479–481, 1984.

[77] L.L. Buishvili, M.G. Menabde, M.E. Gumberidze, and O.V. Sumenko. "magnus paradox" and averaging method. *Physica A: Statistical Mechanics and its Applications*, 170(1):143–149, 1990.

[78] Francisco M. Fernández. Convergence of the magnus expansion. *Phys. Rev. A*, 41:2311–2314, Mar 1990.

[79] Taiki Haga. Divergence of the floquet-magnus expansion in a periodically driven one-body system with energy localization. *Phys. Rev. E*, 100:062138, Dec 2019.

[80] Andreas Brinkmann. Introduction to average hamiltonian theory. i. basics. *Concepts in Magnetic Resonance Part A*, 45A(6):e21414, 2016.

[81] Morag Am-Shallem, Amikam Levy, Ido Schaefer, and Ronnie Kosloff. Three approaches for representing Lindblad dynamics by a matrix-vector notation. *arXiv e-prints*, page arXiv:1510.08634, October 2015.

[82] MATLAB. *version 7.10.0 (R2010a)*. The MathWorks Inc., Natick, Massachusetts, 2010.

[83] The repository for code used in this work is available at `https://github.com/charlessnider/Spin_Echo_Sim`.

[84] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.

[85] Tim Besard, Christophe Foket, and Bjorn De Sutter. Effective extensible programming: Unleashing Julia on GPUs. *IEEE Transactions on Parallel and Distributed Systems*, 2018.

[86] Tim Besard, Valentin Churavy, Alan Edelman, and Bjorn De Sutter. Rapid software prototyping for heterogeneous and distributed platforms. *Advances in Engineering Software*, 132:29–46, 2019.

[87] NVIDIA, Péter Vingelmann, and Frank H.P. Fitzek. Cuda, release: 10.2.89, 2020.

[88] Nicholas J. Higham. The scaling and squaring method for the matrix exponential revisited. *SIAM J. Matrix Anal. Appl.*, 26(4):1179–1193, April 2005.

[89] Cleve Moler and Charles Van Loan. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review*, 45(1):3–49, 2003.

[90] Robert Crovella. How many concurrent threads are running on my GeForce GTX 1080 Ti?

[91] N. Auer, L. Einkemmer, P. Kandolf, and A. Ostermann. Magnus integrators on multicore cpus and gpus. *Computer Physics Communications*, 228:115–122, 2018.