

Artificial Gauge Fields in Ultracold Atomic Ensembles

by

Logan W. Cooke

A thesis submitted in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

Department of Physics
University of Alberta

©Logan W. Cooke, 2024

Abstract

Gauge theory provides a simple and robust way in which to describe the underlying symmetries of nature, and the ability to empirically test such theories is of vital importance. Quantum simulators have found an important application in recent decades, in the generation of artificial gauge fields. These systems allow researchers to create various fields which would otherwise only exist in the most exotic systems. In this thesis, we will discuss techniques with which to realize various artificial gauge fields in spinor Bose-Einstein condensates (BECs), which are coherently manipulated by external radiation. We derive the affect of radio-frequency (RF) and laser fields on atomic spins in a BEC, and then later describe how they may be used to generate both Abelian and non-Abelian gauge fields. We discuss techniques for many-body numerical simulation of BECs in the presence of artificial magnetic fields, created through laser dressing, and extend the algorithm to run more efficiently on a graphics processing unit (GPU). Next, we present a technique for using Floquet engineered RF fields to produce degeneracies in an otherwise nondegenerate spin system; consequently, geometric phases arise which are well described by the presence of a non-Abelian gauge field. We also discuss the relevance of this technique to holonomic quantum computing (HQC). Altogether, we present several theoretical, numerical, and experimental procedures with which researchers may explore both Abelian and non-Abelian gauge theories in ultracold atomic systems.

Preface

The BEC apparatus was designed by Dr. Lindsay LeBlanc, and assembled by Dr. Lindsay LeBlanc, Greg Popowich, Taras Hrushevskiy, and Dr. Erhan Saglamyurek. The original procedure for obtaining BEC was optimized by Dr. Lindsay LeBlanc, Dr. Erhan Saglamyurek, and Taras Hrushevskiy. Upgrades to the imaging system were done by myself and Benjamin Smith. Raman beams were setup and installed, initially by myself, Benjamin Smith, and Taras Hrushevskiy, and again later by myself, Dr. Arina Tashchilina, and Joseph Lindon. Maintenance of the system over the years, including, MOT alignments, ODT alignments, and modifications to the laser system, were carried out by myself, Taras Hrushevskiy, Benjamin Smith, Joseph Lindon, Dr. Arina Tashchilina, and Tian Ooi. Several vacuum bakes were done by myself, Dr. Lindsay LeBlanc, Greg Popowich, Taras Hrushevskiy, Benjamin Smith, Joseph Lindon, and Dr. Arina Tashchilina.

The repump laser was replaced by Tian Ooi, Joseph Lindon, and Dr. Arina Tashchilina. The ODT laser was replaced twice by myself, Joseph Lindon, Dr. Arina Tashchilina, and Tian Ooi. A SAS locking setup for the Ti:Sapph laser was installed originally by Dr. Erhan Saglamyurek and Taras Hrushevskiy, and again later by myself, Dr. Erhan Saglamyurek, and Taras Hrushevskiy.

The BEC apparatus, including the electronics, vacuum system, and laser system, were recovered following two independent floods; first, by myself, Benjamin Smith, Joseph Lindon, and Dr. Arina Tashchilina, and second by myself, Dr. Arina Tashchilina, Tian Ooi, and Joseph Lindon.

The potassium laser setup was originally designed by Dr. Lindsay LeBlanc, and adapted by myself and Anindya Rastogi. The potassium lasers were characterized, and frequency stabilized, by myself and Anindya Rastogi. Anindya Rastogi designed and fabricated the beat-note locking electronics. The potassium oven components were designed and fabricated by myself, and installed by myself, Dr. Lindsay LeBlanc, Taras Hrushevskiy, Anindya Rastogi, and Greg Popowich. Later, a new design was implemented by myself, and Luca Galler. The potassium setup has since been finished by Nicholas Milson.

The work in Ch. 3 is predominantly covered in Ref. [1], mainly pertaining to the spinor BEC simulation software. The original mean-field simulation code was

written by Dr. Lindsay LeBlanc, and modified by myself and Benjamin Smith. The code was packaged, and accelerated on GPU hardware by myself. The spinor version of the code was written and GPU hardware accelerated by myself. Later, both versions were improved and modified, and fully packaged by Benjamin Smith. The grant to acquire the NVIDIA TITAN V GPU was written by myself, Dr. Lindsay LeBlanc, and Benjamin Smith. The benchmarking between devices was carried out by Benjamin Smith.

The experiments discussed in Ch. 6 is featured in Ref. [2]. The Raman Λ -Scheme was setup, and early experiments performed by myself, Joseph Lindon, and Dr. Arina Tashchilina. An additional radio-frequency coil was added to the BEC vacuum cell by Benjamin Smith, and the electronics were later modified by myself, Joseph Lindon, and Dr. Arina Tashchilina. The RF impedance-matching circuit was designed and fabricated by Joseph Lindon. The back-end software for remote programming of the Tektronix AWG5204 device was written by Joseph Lindon. The scripts for implementing the Floquet engineered RF pulses was written by myself. This functionality was later packaged by Joseph Lindon, and further modified by myself. The AWG was incorporated into the apparatus by myself, Joseph Lindon, Tian Ooi, and Dr. Arina Tashchilina. The microwave circuit was originally installed by Dr. Lindsay LeBlanc, Dr. Erhan Saglamyurek, and Taras Hrushevskiy. It was later modified by Benjamin Smith. The SG imaging technique was optimized by myself, Benjamin Smith, and Taras Hrushevskiy. The $F = 1$ imaging system was installed and tested by myself and Dr. Arina Tashchilina.

The Floquet basis derivations in Ch. 5 were carried out by Mason Protter, based on the original works. The experimental measurements were verified through numerics written in `Python` by myself, with additional independent calculations from Mason Protter. Later, the code was adopted to `Julia` through a package written by Mason Protter. Calculations of the holonomies, and Wilson loops, were done by myself, Mason Protter, and Dr. Arina Tashchilina, and later generalized by Mason Protter. The various analysis techniques were devised by myself, Dr. Lindsay LeBlanc, Mason Protter, and Dr. Arina Tashchilina. The numerical simulations featured here were written by myself, and the data processing was performed by myself.

To my loving parents, brother, and sister...

Acknowledgments

First and foremost, I want to thank my advisor, Dr. Lindsay LeBlanc, for the endless support throughout my studies. Thank you for all of your guidance, wisdom, patience, and for maintaining a friendly and forgiving atmosphere. Thank you for sending me to conferences, and for giving me opportunities to interact with other students and researchers, of all levels. Thank you for providing me with what are arguably the more interesting TA positions in the first few years, appointments in which I faced ample challenges that have helped me grow. Additionally, thank you for encouraging me to develop the many skills I picked up along the way; I greatly appreciate your value for personal and professional growth, and for your respect in the humanity of all who you interact with. Thank you, immensely, for the opportunity to work with and learn from you over these past years!

I'd also like to thank my committee members, Dr. Frank Hegmann, and Dr. Joseph Maciejko, for agreeing to oversee my studies here, and for all of the challenging questions.

In no particular order, I want to thank the various colleagues I've worked with during my studies. Anindya Rastogi, thank you for welcoming me into the lab, as your "partner in crime" working on the potassium laser setup. You made my first year extremely enjoyable, and since then, I've greatly appreciated the many insightful conversations we've had, in addition to all of the silly ones. Taras Hrushevskiy, thank you for being the bottomless well of technical knowledge that you are, always managing to find the single out-of-place thing breaking the BEC, and always agreeing to help in any way you could. I learned a great deal working with you, things that ultimately drove the success of our experiments. I also look back quite fondly on the many late nights we spent in front of the QSim computer, particularly while collecting single-photon data while watching the winter Olympics. Thank you! I also want to thank Benjamin Smith, for many interesting and thoughtful conversations over the years. Anytime there was something I just really couldn't wrap my head around, I knew I was safe discussing it with you, and that we could figure it out together. Thanks for your support, and patience.

Dr. Arina Tashchilina, thank you, so much, for your endless energy and excitement for physics. Our experiments would not have been possible without your keen

insight and support. Your mentorship helped shape me into the researcher I am today. Additionally, thank you for being so friendly, patient, and for making even the most depressing situations at least a little bit laughable. Joseph Lindon, thank you for all of your help on the Floquet experiments; you are incredibly talented, dependable, and patient. Thank you for everything, especially all of the major quality-of-life improvements you made to the system. Mason Protter, I sincerely appreciate all of the work you put in to the Floquet project, and for helping me understand the various aspects along the way. Thank you for your keen theoretical insight, and patience explaining things to me. I also want to thank you for introducing me to Julia, which perhaps you regretted after having to look through some of my code. In any case, thank you, for everything.

Thanks as well to Dr. Erhan Saglamyurek, for teaching me many things about optics, and for the opportunity to work with you on the quantum memory project. Thanks Tian Ooi, for all of your help in collecting data, and on the QSim control program. Nicholas Milson, thanks for helping to improve the system; despite the struggles, it was very enjoyable working with you. Thanks as well to the many others who I had an opportunity to work with over the years: Jacques Thibault, Andrei Tretiakov, Michelle Sullivan, Brittany Lu, Bahar Babaei, Anna Prus-Czarnecka, Kusum Meena, Dr. Andal Narayanan, Luca Galler, Timothy Lee, Myles Ruether, Scott Wilson, and too many others. Thank you all!

Last, I want to thank my parents, Lorelee and Randy, for your immense and unconditional love and support; none of what I've accomplished would be possible without you, and everything you've done for me. Thanks to my brother Justin, for your genuine interest in what I do, and for being my absolute best friend, and to my sister, Ashley, for always being so thoughtful and caring. And to my partner, Danielle, thank you for putting up with me all these years, and for always brightening my day. I am deeply thankful for your belief in me, and for listening to me rant on about physics. Finally, thank you to all of my friends and family, of which I'm lucky enough that there are far too many to name; a special mention goes to Neil, Adam, and Zaheen, thanks to you all for being there for me.

Contents

List of Tables	xi
List of Figures	xii
List of Abbreviations	xv
List of Publications	xvii
1 Introduction	1
2 Experimental Methods in Ultracold Ensembles	5
2.1 BEC Production	6
2.1.1 The Ovens	6
2.1.2 2D and 3D MOT	8
2.1.3 Forced RF-Evaporation in an MT	13
2.1.4 Evaporation in an ODT	14
2.2 Dressed Atomic Ensembles	16
2.2.1 RF Dressing	16
2.2.2 Raman Dressing	21
2.3 Spin-Ensemble Measurement	27
2.3.1 Stern-Gerlach TOF Measurements	27
3 Simulating Gauge-Fields in BECs	31
3.1 Producing Abelian Gauge Fields in BECs	32
3.1.1 Spin-Momentum Coupling by Raman Dressing	33
3.1.2 Interpretation as Artificial Gauge Field	36
3.2 Bose-Einstein Condensates	41
3.2.1 Gross-Pitaevskii Equation	41
3.2.2 Thomas-Fermi Approximation	44
3.2.3 Spinors	47
3.3 Numerical Methods	47
3.3.1 Time Splitting Spectral Method	48

3.3.2	GPU Acceleration	52
3.3.3	Hardware Considerations	54
3.3.4	Results	54
3.4	Artificial Magnetic Fields in Spinor BECs	57
4	Non-Abelian Geometric Phase	64
4.1	Geometric Phase	65
4.1.1	Holonomies	66
4.1.2	Some Intuition	68
4.1.3	Abelian vs. Non-Abelian	69
4.1.4	Wilson Loops	71
4.2	Geometric Phases as Artificial Gauge Fields	74
4.3	Geometric Quantum Computing	76
4.3.1	Holonomic Quantum Computing	78
4.3.2	Adiabatic and Topological Quantum Computing	82
5	Floquet-Engineered Non-Abelian Gauge Fields	85
5.1	Floquet Engineered Holonomies	86
5.1.1	Micromotion	86
5.1.2	Bloch's Theorem	89
5.1.3	Fourier Components	91
5.2	Simulated Evolution	94
5.3	Detuning	101
5.3.1	Micromotion	102
5.3.2	Fourier Components	103
5.3.3	Detuned Simulations	105
5.4	Wilson Loops	107
5.4.1	Detuning	110
5.5	Summary	111
6	Periodically Driven Ultracold Ensembles	114
6.1	Experimental Procedure	115
6.1.1	State Preparation	116
6.1.2	Measurement Procedure	120
6.2	Engineering the Floquet Hamiltonian	121
6.2.1	Modulated RF-Dressing	121
6.2.2	Holonomies	123
6.3	Results	129
6.3.1	Time Evolution	130
6.3.2	Phase Control	133

6.3.3	Fidelity of Holonomies	137
6.4	Discussion	142
6.4.1	Detuning Sensitivity	142
6.4.2	Other Sources of Error	144
6.4.3	Floquet-Engineered Artificial Gauge Fields	147
6.4.4	Floquet-Engineered HQC	149
6.5	Conclusion	150
7	Conclusion	152
	References	156
A	Quadratic Zeeman Effect	173
B	Rotating Basis Transformations of Spin Operators	174
C	Spin-State Tomography	176
C.0.1	Spin-1/2	177
C.0.2	Spin-1	178
D	Floquet-Engineered Λ-Scheme	182
E	Computing Holonomies	186

List of Tables

3.1	Relevant specifications of GPUs and CPUs used in benchmarking hardware-accelerated GPE solver.	55
5.1	Parameterized loops and corresponding holonomies, characterized in Floquet-engineered holonomy experiments.	96
6.1	Statistics on the fidelities of measured holonomies.	140

List of Figures

2.1	Schematic of the BEC apparatus, with various components labeled. . .	7
2.2	Level diagram for RF-dressing within an arbitrary F manifold, resulting in coupling between adjacent m_F levels.	18
2.3	Example Rabi-oscillations in $F = 1$ manifold, solved for numerically. . .	20
2.4	Level diagram for Raman coupling, showing the M -scheme and Λ -scheme.	22
2.5	Example Stern-Gerlach images, showing a full Rabi-oscillation in an RF coupled ultracold ensemble.	28
3.1	Atomic dispersions with Raman-dressed spin-momentum coupling, for various values of Raman Rabi frequency and detuning.	37
3.2	Dispersions of Raman dressed states, with relative spin populations indicated along bands to show the mixture of spins with varied momentum.	38
3.3	Artificial vector potential as it varies with detuning, indicating how it may be engineered in experiments.	40
3.4	Comparison of TF approximation for BEC density against numerical simulations, for different interaction volumes, along with sample numerical density profile.	46
3.5	Benchmark calculations of spinor GPE code, as run on the various devices summarized in Tab. 3.1.	56
3.6	Spinor BEC density distributions, in the weak-coupling regime with a small detuning gradient, exemplifying the behavior of the Raman coupled GPE.	58
3.7	Raman-coupled spinor ground states in the strong coupling regime, showing the clear emergence of quantized vortices.	60
3.8	Closer look at Raman-coupled BEC in the strong-coupling regime, where the phase of the spinor wavefunctions reveals the interference pattern produced by the overlapping waves to be quantized vortices. . .	60

3.9	Raman-coupled spinor BEC ground state in the strong coupling regime, with an artificial magnetic field strong enough to permit vortex nucleation in the higher density regions of the spins.	61
3.10	Raman-coupled spinor BEC with both strong coupling and detuning gradients, resulting in less miscible spins with multiple quantized vortices.	62
5.1	Various loops in parameter space, which were ultimately implemented in experiment (Sec. 6.1).	96
5.2	Numerical simulations of Floquet-engineered holonomies, showing state trajectories on the Bloch sphere for several loops.	97
5.3	Numerical simulations of Floquet-engineered holonomies, showing how spin populations vary in time, for $F = 1/2, 1, 2$	99
5.4	Loops in parameter space with more complicated structure, and numerical simulations of the resulting spin evolution.	100
5.5	Numerical simulations of $\hat{\Gamma}_A^\Delta(\ell_1)$ on the Bloch sphere, demonstrating the impact of each individual component of the detuning Δ	106
5.6	Numerical simulations of spin populations for detuned holonomy $\hat{\Gamma}_A^\Delta(\ell)$ with detunings Δ sampled from a Gaussian distribution.	108
5.7	Impact of each component of the detuning Δ on the trace-commutator, for spins $F = 1/2, 1, 2$	112
6.1	Experiment pulse sequence for Floquet-engineered holonomies.	116
6.2	Diagram depicting the pulse sequence for preparing atoms in any of the $F = 1$ magnetic sublevels.	119
6.3	Some sample RF-waveforms, programmed into an AWG to effect Floquet-driven holonomies in ultracold ensembles.	125
6.4	Comparison of sensitivity of standard unmodulated RF-pulses and holonomy $\hat{\Gamma}_A(\ell_1)$ to detuning, motivating our resonance calibration technique.	128
6.5	Measurements of the time evolution of spin populations, for loops ℓ_1 and ℓ_3 in the $F = 1$ manifold.	131
6.6	Measurements of the time evolution of spin populations, for loops ℓ_1 and ℓ_6 in the $F = 2$ manifold.	132
6.7	Phase-sensitive measurements of holonomic transformations in $F = 1$ spin manifold, demonstrating their path dependence.	135
6.8	Phase-sensitive measurements of holonomic transformations in $F = 2$ spin manifold, demonstrating their path dependence.	136
6.9	Density distributions of fidelities of measured holonomies, with and without detuning considered.	141

6.10	Histogram of detuning values from fits to measured holonomies. . . .	142
D.1	Sample Rabi-oscillations under a Raman Λ -scheme, showing significant decoherence and atom loss with increasing pulse duration. . . .	185

List of Abbreviations

- AOM** acousto-optical modulator. [12](#), [182](#), [183](#)
- AQC** adiabatic quantum computing. [78](#), [82](#), [83](#)
- AWG** arbitrary waveform generator. [115](#), [117](#), [118](#), [120](#), [123–125](#), [127](#), [129](#), [137](#), [183](#)
- BEC** Bose-Einstein condensate. [2](#), [3](#), [5–7](#), [15](#), [21](#), [28](#), [29](#), [31](#), [32](#), [41](#), [43–48](#), [50–52](#), [54](#), [58](#), [61](#), [62](#), [75](#), [76](#), [116](#), [118](#), [124](#), [145](#), [152](#), [153](#), [176](#), [182](#)
- CC** compute capability. [54](#), [55](#)
- CPU** central processing unit. [48](#), [51](#), [53–56](#)
- FFT** fast Fourier transform. [48–53](#), [56](#)
- GPE** Gross-Pitaevskii equation. [31](#), [32](#), [43](#), [44](#), [46–52](#), [55–57](#), [62](#), [153](#)
- GPU** graphics processing unit. [32](#), [48](#), [52–56](#), [62](#), [153](#)
- GQC** geometric quantum computing. [77](#), [78](#), [82–84](#)
- HQC** holonomic quantum computing. [78–85](#), [115](#), [123](#), [139](#), [142](#), [149–151](#), [153](#), [154](#)
- MOT** magneto-optical trap. [6–13](#), [27](#), [29](#)
- MT** magnetic trap. [13](#), [15](#), [27](#), [28](#)
- NAHQC** non-adiabatic holonomic quantum computing. [81](#), [82](#)
- NMR** nuclear magnetic resonance. [82](#), [154](#)
- ODT** optical-dipole trap. [14](#), [15](#), [29](#), [44](#), [48](#), [115](#), [116](#), [120](#), [123](#), [126](#), [145](#), [184](#)
- PID** proportional-integral-derivative. [11](#), [12](#), [183](#)

QC quantum computing. 64, 65, 70, 75, 77–79, 81, 83, 84, 115, 149, 153, 181

QIP quantum information processing. 65, 76, 85, 133, 139, 142, 150, 153

RF radio-frequency. 7, 11, 13–16, 18, 21, 23–32, 35, 114–116, 118–130, 133–137, 143, 148, 152, 154, 180, 182, 184

SG Stern-Gerlach. 27–29, 120, 146, 152

TDSE time-dependent Schrödinger equation. 44, 50, 65, 77, 87, 90, 94

TF Thomas-Fermi. 31, 44–46, 48, 50, 51, 57, 146

TOF time-of-flight. 27–29, 62, 116, 120, 145, 146, 152

TQC topological quantum computing. 78, 83, 84

TSSM time-splitting spectral method. 48–51, 53, 56

UHV ultrahigh vacuum. 6–8, 10

List of Publications

Saglamyurek, E., Hrushevskiy, T., Cooke, L., Rastogi, A. & LeBlanc, L. J. Single-photon-level light storage in cold atoms using the Autler-Townes splitting protocol. *Physical Review Research* **1**, 022004 (2019)

Saglamyurek, E., Hrushevskiy, T., Rastogi, A., Cooke, L. W., Smith, B. D. & LeBlanc, L. J. Storing short single-photon-level optical pulses in Bose–Einstein condensates for high-performance quantum memory. *New Journal of Physics* **23**, 043028 (2021)

Smith, B. D., Cooke, L. W. & LeBlanc, L. J. GPU-accelerated solutions of the nonlinear Schrödinger equation for simulating 2D spinor BECs. *Computer Physics Communications* **275**, 108314 (2022)

Lindon, J., Tashchilina, A., Cooke, L. W. & LeBlanc, L. J. Complete unitary qutrit control in ultracold atoms. *Physical Review Applied* **19**, 034089 (2023)

Cooke, L. W., Tashchilina, A., Protter, M., Lindon, J., Ooi, T., Marsiglio, F., Maciejko, J. & LeBlanc, L. J. Investigation of Floquet engineered non-Abelian geometric phase for holonomic quantum computing. *Physical Review Research*. (arXiv preprint: 2307.12957) (2024). Forthcoming

CHAPTER 1

Introduction

PHYSICISTS likely encounter gauge fields for the first time in their undergraduate studies, in the context of electromagnetism. After familiarizing themselves with Maxwell's equations to some extent, they are introduced with the concept that there is some redundancy seemingly built in, where the electric and magnetic fields may be expressed through a new set of fields, called the vector and scalar potential. The important insight here is that the electric and magnetic fields are uniquely determined by these new *gauge fields*, but the converse is not true. There are an infinite set of gauge fields that describe the same electric and magnetic fields, so choosing any such pair of them therefore describes the same physics. This is akin to choosing from a set of different style measurement devices, or gauges, which all record the same observations once utilized¹; it is thus known as *gauge freedom*. This is usually presented as an interesting, and potentially useful, mathematical tool, which can be used to simplify calculations in certain circumstances. To this end, these gauge fields conveniently cannot be measured, as the electric and magnetic fields are the *physical* fields. Such a viewpoint, unfortunately, undercuts the significant importance that gauge theory has in modern physics.

Gauge theories are formulated through group theory², which provides a useful framework from which we may understand the properties of gauge fields [3]. Electromagnetism, for instance, is described by a $U(1)$ Abelian gauge theory, where the $U(1)$ group is Abelian because its elements, complex numbers of the form $e^{i\phi}$ with $\phi \in \mathbb{R}$, commute; the values of ϕ are called the *generators* of the group, which is the role that the gauge fields play. Even with such seemingly simple fields as in classical electromagnetism, there are a plethora of interesting implications, not least of which

¹Note that we are not suggesting that this is why it is called gauge theory, rather, we are just providing a potential way to think about it.

²This is a branch of mathematics that this author was unfortunately told explicitly to stay away from by a certain calculus professor during undergraduate studies. Now, we find ourselves desperately learning all we can about it.

is the propagation of electromagnetic waves, *i.e.* light. Interestingly, in the context of quantum physics, one does not need to try very hard to find the emergence of gauge theories. For instance, imposing local gauge-invariance on the wavefunction of a free electron naturally leads to electromagnetic gauge potentials appearing in the Schrödinger equation [4]; light, is therefore a consequence of local gauge invariance.

Beyond Abelian gauge theories, there are also non-Abelian ones that are of immense importance in modern physics. Non-Abelian groups are comprised of elements that do not commute; generally this is because the group elements are matrices. The consequences of this are significant, and may manifest in several ways. For instance, in particle physics, quarks and gluons are described through an $SU(3)$ gauge theory, which is non-Abelian. The non-commuting elements of this group appear in the Lagrangian, which ultimately results in interacting gluons; the interaction strengths are related to the commutators, which are called *structure factors*. In contrast, no such interaction term appears for photons since they follow an Abelian gauge theory, and hence the commutators vanish [3]. Additionally, unitary evolution in quantum mechanics is also well described through groups. The $N \times N$ unitary matrices are members of the group $U(N)$, which is non-Abelian; this is just a statement that the order in which we apply unitary transformations matters, since these matrices do not in general commute. While invoking gauge theory in this context is generally not required, in this work we will find that it is extremely useful in classifying the types of transformations that are enabled by some Hamiltonian of interest.

Given the significant importance of gauge fields in modern physics, they ought to be a subject of experimental study. We may of course look at the dynamics of charges and currents to understand those behind electromagnetism, but once we foray into the realm of high-energy gauge theories we may find that the experiments are much more difficult. This is, therefore, a branch of study that quantum simulation is well equipped to offer value in. Quantum simulation, briefly put, aims to make highly controllable quantum systems behave in a manner that is equivalent to some target system of interest [5]. Typically, the target system is one that we do not have access to, or in some cases may not even exist in nature. Due to our high-degree of control over the quantum simulator, we are able to realize the physics of these systems anyways, at least under some constraints (lacking gauge invariance for instance). This practice usually amounts to Hamiltonian engineering; if the system Hamiltonian is made equivalent to that of some other system, the dynamics are directly analogous. Even though the objects described by one such Hamiltonian may differ greatly from those of the other, it is fair to interpret them as equivalent.

In this thesis, we will discuss the quantum simulation of gauge fields in ultracold neutral atoms. Specifically, we will introduce several ways in which we may manipulate a [Bose-Einstein condensate \(BEC\)](#) of ^{87}Rb atoms, so as to produce *ar-*

tificial gauge fields, that is, gauge fields that do not technically exist, but act as an appropriate way of describing the system dynamics. In each case there is a more direct way of describing the dynamics, namely through the typical interactions of light with atoms, but the existence of an equivalent description through gauge fields allows us to study such objects in a way that can be generalized to *real* situations in which those fields emerge. Conversely, we will find that viewing the problem through the lens of gauge theory provides some unique insights, which aid in predicting the behavior of many-body systems of interacting particles, such as a [BEC](#). Altogether, This work paves a way forward in generating artificial gauge fields that themselves are objects of immense interest to researchers, as well as tools with which to study many-body systems.

We begin in [Ch. 2](#) by describing the experimental apparatus in some detail, along with the procedure used to produce ultracold ensembles of ^{87}Rb . This includes the laser cooling process, and evaporative cooling in both a magnetic trap and optical dipole trap. We then describe two ways to use external fields to drive transitions between the magnetic sublevels of atoms, with either external oscillating magnetic fields, or laser fields. This is the primary means used to manipulate the ultracold ensemble in the remaining situations discussed in this thesis. Finally, we close this chapter by discussing the spin-sensitive measurement technique employed.

In [Ch. 3](#) we discuss a particular implementation of Abelian artificial gauge fields, using the laser manipulation scheme discussed in the previous chapter. This scheme motivates a theoretical investigation into the effects of such gauge fields on [BECs](#) where the interactions are important, leading to the development of a high-performance computing package for such a system. We introduce the relevant mean-field theory of interacting bosons, followed by an algorithm that may be used to solve for spinor [BEC](#) ground states and dynamics, when subjected to various external fields. The algorithm is accelerated on graphics hardware, and then applied to determine several interesting ground states under the effects of the artificial gauge fields discussed prior.

Later, in [Ch. 4](#) we discuss a different way in which gauge theory manifests itself in quantum systems, namely through geometric phase. We introduce the concept of geometric phase in systems undergoing adiabatic evolution. This leads to the emergence of either Abelian or non-Abelian gauge fields as a way of describing the resulting state transformations. We discuss some properties of each, in addition to a means of distinguishing them from each other in a gauge invariant manner, through the Wilson loop. We then describe how geometric phases may be interpreted as artificial gauge fields, discussing some of the insights we might hope to learn from them, in addition to what types of phases we hope to some day engineer. The application of geometric phase as a means of quantum control is then discussed, in

relation to several quantum computing approaches that rely on it.

Producing non-Abelian gauge fields has historically been challenging, so in Ch. 5 a technique for surpassing some such issues is introduced, using Floquet engineering. The periodic modulation of Hamiltonians may lead to the generation of non-Abelian geometric phases, which are described by the presence of artificial gauge fields in the parameter space of the Hamiltonian. We derive the emergence of such fields, and demonstrate the resulting transformations numerically. A fairly general source of error is then considered, the impact of which is also demonstrated numerically. Chapter 6 goes on to describe the experimental demonstration of this technique in the ground state spin manifolds of an ultracold ensemble of ^{87}Rb . We discuss how the necessary periodic driving is accomplished by modulated radio-frequency coupling between atomic spins. We then describe the procedure for state preparation and measurement, before showing the results, compared with numerical simulations. A detailed discussion of the experiment reveals some of the shortcomings, in addition to the major successes of this preliminary investigation.

Altogether, the work presented in this thesis establishes some powerful tools and techniques now at the disposal of contemporary researchers, for generating and studying artificial gauge fields in the lab. The insights gained from such quantum simulation experiments, in the context of gauge theory, is of tremendous importance to our current understanding of the universe in which we live.

CHAPTER 2

Experimental Methods in Ultracold Ensembles

REALIZING artificial gauge fields in ultracold atomic ensembles requires several key methods, experimentally speaking. First, and most obviously, one must be capable of efficiently producing the ultracold ensembles. Following this, one requires a means with which to manipulate them. Finally, a means to measure the states or relevant observables is required. In this section, we aim to address each of these requirements, to a level of detail that is sufficient for the chapters ahead.

We begin in Sec. 2.1 by giving a general overview of the apparatus that we have used, maintained, and augmented over the duration of this thesis work. The apparatus, which was originally assembled and optimized for producing ^{87}Rb BECs as part of the MSc thesis *Hrushevskiy* 2017 [6], has seen several changes over time. In addition to the typical maintenance and periodic replacements or adjustments of existing infrastructure, it has been modified to also incorporate ^{39}K into the system. This involved the construction of an additional laser system, and the fabrication of several more vacuum parts. While the addition of ^{39}K was a primary motivation for some of the work discussed in Ch. 3, it has yet to be fully integrated into the apparatus; as such, here we will focus on the parts relevant to the cooling of ^{87}Rb to quantum degeneracy. The section does not aim to provide a rigorous description of each component or cooling process used, but rather a general understanding of the various things that impact the work that follows.

Section 2.2.1 then focuses on the primary means we have to manipulate ultracold ensembles, namely with external radiation. We derive a Hamiltonian which describes the coupling of internal spin states by an external oscillating magnetic field. A similar result is obtained when considering transitions driven in the ensemble by a bichromatic laser field. Each of these techniques is then applied in later chapters to produce various artificial gauge fields that act on the spin systems.

Last, we discuss a way to make projective measurements in the canonical spin basis. Combining some of the techniques described in the former sections (Sec. 2.1 and Sec. 2.2.1), we are able to measure spin states through standard absorption imaging. This will become especially relevant in Ch. 6, but also has applications to the work in Ch. 3. With this, we will have covered all of the necessary ingredients in the quantum simulation experiments that follow: state preparation, manipulation, and measurement.

2.1 BEC Production

Since their synthesis in the 1990s [7, 8], BECs have become a hallmark success of quantum physics, and precision control technologies. BECs are macroscopic quantum objects which are almost perfectly pure [9], and may be precisely and coherently manipulated in various ways (See Sec. 2.2). Due to many advancements in their production process, BECs are now even being produced in space, including during the rocket launch stage [10]. That said, for the efficient production of large degenerate ensembles, the entire process must be well calibrated. Here we give a general overview of the BEC apparatus used in the work throughout this thesis, and the cooling procedure used to produce degenerate ensembles of ^{87}Rb ; a complete description of the apparatus used and the processes mentioned here may be found in Hrushevskiy 2017 [6].

2.1.1 The Ovens

The ultrahigh vacuum (UHV) system used here consists of three major stages: the source of ^{87}Rb (the *ovens*), 2D magneto-optical trap (MOT) stage, and the 3D MOT chamber (the *science chamber*). A schematic of the apparatus, with labels indicating these regions along with various other components, is shown in Fig. 2.1. Sources of ^{87}Rb , which come in glass ampoules, are held in the UHV system on a custom built ampoule stand. After sealing the system and baking, a flexible coupling with a metal rod mounted within is used to break the glass ampoule. The UHV components containing the source are wrapped in resistive heating tape; warming these elements to 45°C increases the vapour pressure of ^{87}Rb in the system enough to provide the number of atoms needed. A laser beam tuned above an atomic resonance (called *blue-detuning*, due to the shorter wavelength) is transmitted through a viewport on the back of the ovens, to push atoms along into later stages, eventually loading the 3D MOT; we refer to this beam as the *push beam*.

Due to the relatively high pressure in the ovens, there is a custom copper aper-




Figure 2.1: Schematic of the BEC UHV system, indicating the various key regions. The ovens contain sources of both ^{87}Rb and ^{39}K , which may be accelerated towards the other regions by the push beam through a viewport at the end of the apparatus. Past the cold cup is the 2D MOT region, where atoms are pre-cooled, and focused through a differential pumping tube towards the 3D MOT, where they are trapped and cooled to quantum degeneracy. In addition to the gradient coils around the 3D MOT, there are bias coils along each axis, and radio-frequency (RF) coils along z and x . Figure is adapted from Ref. [6].

ture mounted inside, prior to the 2D MOT stage, with an ion pump above. This aperture, called the *cold cup*, is thermally coupled to a Peltier cooler and water cooling line outside the UHV system, and is kept at 16 °C at all times. Since the surrounding UHV components are either room temperature or warmer, the cold cup acts as the most likely surface to which ^{87}Rb atoms adsorb; it effectively works as a non-evaporable getter, placed between the ovens and 2D MOT stage, which should remain at the lowest pressure possible. The aperture allows the push beam to pass through, carrying atoms with it.

The ion pump mounted above the cold cup, which maintains a low pressure in the ovens, pumps more ^{87}Rb than it is likely designed to do, given its vicinity to them. Metallic ^{87}Rb is deposited on the capacitor plates in large enough quantities that the capacitance changes, which alters the pressure reading of the pump. Over time the pressure reading from the pump increases, despite no evidence that the true pressure within has increased. Eventually it may pass a threshold for which the pump will turn off, to protect itself; if this happens, the pressure in the system will raise dramatically, perhaps to the point that it would have to be unsealed and re-baked. For this reason, we added a set of aluminum baffles within the fitting leading into the ion pump, in an effort to stem the flow of ^{87}Rb into the pump. So far, it seems to have prolonged the lifetime of the pump, but more time is needed to be certain of this.

Similar to the ion pumps, the glass push-beam viewport is also susceptible to ^{87}Rb deposition, which etches away at the glass-metal seal. Eventually, this cracks the glass, causing a leak. For this reason, the window is always kept 5 °C higher than all the surrounding parts, to make it less likely for ^{87}Rb to adsorb there. Unfortunately this does not seem to be enough to protect it. As such, we machined and installed a stainless-steel aperture just in front of the window, so that the push beam may still pass through, but the reduced vacuum conductance reduces the ^{87}Rb buildup. This measure has been successful in improving the lifetime of the viewport.

2.1.2 2D and 3D MOT

Once atoms make it past the cold cup they enter the 2D MOT stage. We will first introduce the general principle behind a MOT, and then discuss the details of the 2D and 3D MOT stages.

Principles of a MOT

A MOT accomplishes two tasks simultaneously: it significantly reduces the average speed of atoms (the temperature) through the *optical molasses* process, and it applies a trapping force that constrains atoms to the minimum of an applied

magnetic field gradient [11, 12]. The setup consists of pairs of counter-propagating laser beams which are *detuned* to a frequency below an atomic resonance (called *red-detuning* due to the longer wavelength). In a MOT, they also have opposite circular polarizations.

The optical molasses protocol can be broken down into two dominant cooling processes: Doppler cooling, and sub-Doppler cooling. The former relies on the Doppler shift of atoms bringing one laser within a pair closer to resonance with the relevant transition; this creates an imbalance of the scattering forces between the beams. In the case of red-detuning, an atom is more likely to scatter light from the laser it is moving towards. As such, on average over many absorption and spontaneous emission cycles, atoms lose momentum, which is dissipated through the spontaneous emission. The net force on atoms from Doppler cooling is proportional to their speed, and in an opposing direction. For this reason, the Doppler cooling model is the same as that for a classical particle moving through a viscous fluid, hence the name optical molasses. In this case, the laser fields play the role of the viscous fluid.

Doppler cooling is the dominant force on higher-velocity atoms, but once they approach the Doppler cooling limit [11, 12] other processes take over. In the configuration described above with orthogonal circular polarizations, the relevant process is called motion-induced orientation cooling [13, 14]. In all sub-Doppler cooling processes, the structure of the atomic ground states as they vary in the laser fields become relevant. The counter propagating laser fields create a standing wave with a linear polarization everywhere, but with an angle that rotates along the beam axis. Due to the local electric field changing in space, as atoms move along the standing wave their stationary states change. Atoms moving relatively fast will not adiabatically follow the local field, and hence non-adiabatic effects dominate, which result in imbalances in population between the ground state magnetic sublevels. As in Doppler cooling, this ultimately results in an imbalance of scattering forces on atoms due to each laser; in this case, atoms are once again more likely to scatter light from the laser they move towards, resulting in a cooling effect.

The optical molasses protocol acts to cool atoms dramatically. While the Doppler cooling limit for ^{87}Rb is $146\ \mu\text{K}$ [15], through sub-Doppler processes the temperature achieved is often much colder. The fundamental temperature limit in laser cooling is the recoil limit, which is determined by the velocity of atoms after a single spontaneous emission event from the lasers; for ^{87}Rb this temperature is $362\ \text{nK}$ [15]. Despite this, atoms in an optical molasses continue to expand outside the influence of the laser fields, reducing their efficiency.

A MOT prevents the expansion of atoms within the beams by the addition of a magnetic field gradient, which splits the excited state manifolds through the

Zeeman effect (See Sec. 2.2.1). In the presence of a gradient with a global minima, the spatial dependence of the Zeeman splitting, along with the circularly polarized lasers, once again results in an imbalance in scattering forces between lasers. In this case, if the circular polarizations are chosen properly (according to the direction of the magnetic field on either side of the minima), atoms preferentially scatter light from the laser they are closest to, thereby pushing them towards the field minima (regardless of their velocity). A MOT therefore cools atoms through the aforementioned optical molasses processes, and confines them to an applied gradient minimum thus preventing their expansion outside the cooling beams. Additionally, MOTs have significantly higher capture velocities, as compared to standard optical molasses setups, allowing them to be loaded from hotter sources, typically resulting in a much higher number of trapped atoms [12].

2D MOT

The 2D MOT stage consists of a rectangular glass cell, with two pairs of counter-propagating circularly-polarized laser beams transmitted through in the directions orthogonal to the push beam path. The beams are made highly elliptical with cylindrical lenses, so that they extend over a large length of the cell. Additionally, there are two pairs of copper coils arranged in anti-Helmholtz configurations, to provide the necessary magnetic field gradient. As atoms move along the push beam towards the 3D MOT stage, the 2D MOT beams cool them in the orthogonal directions to this path, focusing them into a narrow atomic “beam”. Individually controlling the currents in each coil also allows us to roughly align the atomic beam to the 3D MOT position.

Just before the 3D MOT chamber there is a differential pumping tube in the UHV system, reducing the vacuum conductance between the two stages; additionally, there is a second ion pump just above the tube. The pre-cooling and focusing accomplished by the 2D MOT is essential to efficiently loading the 3D MOT, by guiding atoms through the narrow tube, in addition to keeping the pressure in the 3D MOT chamber low. Only atoms with sufficiently low transverse velocities make it through the tube, and onwards to the 3D MOT. With the push beam and 2D MOT operating optimally, we achieve a factor of 7-10 times more atoms in the 3D MOT than without them. This is a considerable increase in the overall atom number, which is essential for the latter evaporative cooling stages that take place after the 3D MOT processes.

3D MOT

As with the 2D MOT, the 3D MOT cell is rectangular and made entirely of glass for maximum optical access. The cell is shrouded in several sets of magnetic coils: one pair of Helmholtz coils along each Cartesian axis to provide tuneable bias fields, two sets of RF coils for driving RF transitions between atomic spin states (Sec. 2.2.1), and a large set of water-cooled coils in anti-Helmholtz configuration to create the gradient used in the MOT and for magnetic trapping. The optics for the MOT beams are mounted on a custom rigging which fits around the cell, directing the light towards the center of the gradient coils. The 3D MOT has three pairs of counter-propagating beams with opposing circular polarizations, to cool and trap atoms in all dimensions.

There are several requirements of the laser system in order to implement optical molasses cooling in ^{87}Rb . For alkali atoms, it is common to use the D2 line for cooling, which is the transition between $S_{1/2}$ and $P_{3/2}$ manifolds. This transition is at a wavelength of about 780 nm, for which there are commercial diode lasers available. In particular, we use the common $|F = 2\rangle \rightarrow |F' = 3\rangle$ transition, with F the total atomic angular momentum quantum number. This transition is useful because, due to selection rules [12], atoms are forbidden from falling into the other ground state hyperfine level, $F = 1$; this is said to be a *closed-cycle* transition for this reason.

The laser frequency must be tuned precisely; the theoretically optimum red-detuning for a MOT is 3Γ , where Γ is the natural linewidth of the atomic excited state used in the cooling cycle [12]. All transitions within both the D1 and D2 lines have a natural linewidth of 6.1 MHz. As such, the laser frequency must be stable to within approximately one linewidth for efficient and consistent cooling. Despite the fact that we purposefully use a closed-cycle transition, being red-detuned from the excited state hyperfine level means the lasers are still relatively close to resonance with the neighboring level below; through off-resonant scattering with this lower level, atoms may fall out of the cooling cycle by spontaneously emitting into the lower ground state hyperfine manifold, $F = 1$. The ground state hyperfine splitting in ^{87}Rb is 6.8 GHz, which is much larger than the laser linewidth. As such, there is an additional laser resonant with the $|F = 1\rangle \rightarrow |F' = 2\rangle$ D2 transition to pump atoms that have fallen out of the cooling cycle back in, which we denote the *repump* laser.

Locking both lasers to the correct frequency, and keeping them stable to within 300 kHz is accomplished through saturated absorption spectroscopy [12]. This is a conventionally used technique in which the spectrum from a warm vapour cell is used in a **proportional-integral-derivative (PID)** feedback circuit to frequency stabilize

lasers to the relevant transition. The technique allows us to resolve individual atomic transitions up to the natural linewidth, allowing us to not only select the right wavelength, but to stabilize them to the desired precision as well.

We use saturated absorption spectroscopy to lock the repump laser to the indicated repump transition, and then use a *beat-note* locking system to lock the cooling laser ≈ 6.8 GHz away. This is done by interfering the repump and cooling beams on a photodiode, and using [PID](#) feedback to match their beat frequency to a reference signal. The particular lock frequency may be controlled remotely by changing the reference signal, and is used throughout the cooling procedure at various points described below to tune the specific laser frequency. Once a laser is set to the correct resonant frequency ω , [acousto-optical modulators \(AOMs\)](#) are used to precisely tune the frequency of each particular beam with respect to one another, within a range of $\omega \pm 160$ MHz.

To load atoms into the science chamber and cool them, the push beam and 2D [MOT](#) are enabled during the entire Doppler cooling stage of the 3D [MOT](#), which lasts approximately 15 s. We use close to the maximum available power in each of the four 2D and six 3D [MOT](#) beams, which is about 20 mW per beam (the 2D [MOT](#) typically requires more, since it has large elliptical beams). During this time, the large coils are used to provide a magnetic field gradient. We are typically able to collect about 10^9 atoms at a temperature of about 500 μK [6]. This is higher than even the Doppler cooling limit, since the presence of magnetic fields reduces the efficiency of sub-Doppler cooling processes, and the trapping forces in a [MOT](#) also apply to atoms which are already cool thus increasing the overall expected temperature. The trade off, though, is significantly more atoms that may be held for longer.

Following the [MOT](#) cycle, which is optimized for Doppler cooling processes and trapping, we further reduce the ensemble temperature by optimizing the sub-Doppler cooling processes. The gradient coils are turned off, and the bias coils are tuned such that they approximately cancel ambient magnetic fields. The sub-Doppler processes depend on the standing wave wavelength; longer (shorter) wavelengths cool faster (slower) atoms more efficiently. As such, the cooling laser frequency is linearly scanned from a lower frequency to higher, which first cools the faster atoms, gradually optimizing for the slowest. This results in a much higher cooling efficiency overall. The entire process takes 20 ms, and captures approximately all of the [MOT](#) atoms, reducing their temperature to about 50 μK [6].

2.1.3 Forced RF-Evaporation in an MT

Following the 3D MOT, we load atoms into a magnetic trap (MT). Magnetic fields split atomic energies through the Zeeman effect by an amount

$$V = g_F \mu_B m_F B, \quad (2.1)$$

according to the magnetic sublevels m_F , where g_F is the total angular momentum g -factor, μ_B the Bohr magneton, and B the applied magnetic field amplitude (see Sec. 2.2.1 for more details). When there is a magnetic field gradient, the energies depend on space, and hence a trapping (or anti-trapping) force may result, depending on the sign of $g_F m_F$; for $g_F m_F > 0$ atoms seek the magnetic field minimum, and for $g_F m_F < 0$ the maximum. Atoms in the MOT fill all m_F levels of the $F = 2$ hyperfine ground state, which has $g_F = 1/2$. The large coils in our apparatus produce a quadrupole gradient, and so atoms in the $m_F = 1, 2$ states are trapped, while the $m_F = -1, -2$ states are anti-trapped, and $m_F = 0$ is approximately unaffected.

To trap as many atoms from the MOT as possible, we must prepare the atoms in the ground state $m_F = 2$ sublevel, which experiences the strongest trapping potential. This is done with optical pumping [11–13]; a laser addressing the $D2$ $F = 2 \rightarrow F' = 2$ transition with circular polarization is applied for about 1 ms, affecting repeated σ^+ transitions in the atoms, pumping them all into the $m_F = 2$ ground state.

In the MOT, the applied gradient is fairly weak, only about 11 G/cm in our case [6]; this is because the trapping effect in a MOT comes from the interactions with the lasers, not the magnetic gradient itself. This gradient is produced with only 26 A flowing through the coil, at 1 V. After optical pumping, we ramp this gradient up to 27 G/cm [6], which may be accomplished with 65 A of current at about 2.5 V. This gradient is sufficient to hold the $m_F = 2$ atoms against gravity, but not the $m_F = 1$ atoms. The ensemble is held in this trap for 100 ms, allowing the few atoms which remain in other m_F levels to fall away from the trap.

Following this stage, the trap is compressed further by ramping up the current significantly. We apply a gradient of 176 G/cm [6], with 430 A at 15 V. This drastically increases the atomic density to the point that rate of collisions within the gas is enough start the forced RF-evaporation.

High-velocity atoms have sufficient kinetic energy to climb the trapping potential, moving further out from the trap center. Due to the energy splitting from the Zeeman effect, the energy difference between the various m_F levels thus depends on the distance away from the trap center. An RF coil, is used to drive transitions directly between adjacent m_F levels within the $F = 2$ ground state manifold, by

producing an oscillating magnetic field on resonance with the Zeeman splitting (See Sec. 2.2.1 for details on this coupling mechanism). In this way, we selectively remove atoms in a particular energy class by coupling them to the anti-trapped m_F levels, just by driving an RF coil at a chosen frequency.

To perform evaporation in this way, we linearly ramp the carrier frequency of an RF coil aligned along the z -axis of the system (Fig. 2.1) from 25 MHz to 1 MHz, over about 6.5 s. The high velocity atoms are removed from the trap by the RF radiation, leaving those remaining to collide and re-thermalize, thus reducing the overall temperature of the ensemble. As the frequency of the RF drive is reduced, the temperature is reduced. At the end of the RF evaporation stage, roughly 10^6 atoms remain, with a temperature of about $10 \mu\text{K}$ [6].

2.1.4 Evaporation in an ODT

The final stage of the cooling procedure is evaporation in an optical-dipole trap (ODT). Atoms are trapped at the intersection of two orthogonal laser fields through the ac Stark shift [11, 12]. The ground state atomic energies are shifted in the presence of electric fields (See Sec. 2.2.2), and so atoms may be trapped in the intensity gradients of focused lasers. When the lasers are red-detuned from an atomic transition, atoms are attracted to the high-intensity regions of the beam. If the chosen laser is too close to resonance, while it may produce a strong trapping force for even modest powers, the scattering from off-resonant absorption processes produces significant heating effects, reducing the trapping lifetime. For this reason, it is common to use a highly detuned laser for optical trapping, where the lower trapping potential is offset by high laser powers.

In our case, we use an intensity-stabilized 1064 nm laser, far detuned from both the D1 and D2 lines, with approximately 2.2 W of power in each of the two trapping beams. These beams are derived from the same fiber laser with a 10 W total output. The ability for these beams to trap atoms may be quantified in terms of the *trap depth*, which is often given in units of temperature. It is defined by the minimum energy atoms would need in order to escape the trap potential. For an ODT derived by lasers with total power P and beam waist w_0 , the trap depth is [16],

$$V_{\text{ODT}} = \frac{\hbar\Gamma^2}{2} \frac{P}{\pi w_0^2 I_0 \Delta}, \quad (2.2)$$

where Γ is the natural linewidth of the addressed transition. The saturation intensity I_0 is given by,

$$I_0 = \frac{\pi^2 \hbar c \Gamma}{3\lambda^3}, \quad (2.3)$$

where λ is the trapping laser wavelength. Finally, Δ is the effective laser detuning, given by,

$$\frac{1}{\Delta} = \frac{1}{\Delta_1} + \frac{2}{\Delta_2}, \quad (2.4)$$

with Δ_1, Δ_2 the detuning from the D1, D2 lines respectively.

The **ODT** beams in our system are focused to beam waists of about $70 \mu\text{m}$. They are aligned to a region several micrometers below the **MT** center, and made to intersect. The saturation intensity of ^{87}Rb is $I_0 = 16.7 \text{ W/m}^2$ for the D2 line [16], and the effective detuning of the 1064 nm laser is $\Delta = 33.4 \text{ THz}$. This results in a trap depth (Eq. 2.2) of approximately $V_{\text{ODT}}/k_B = 65 \mu\text{K}$, far above the final temperature of the atoms within the **MT**.

To load the **ODT**, the beams are ramped linearly to full power over about 4 s at the start of the **RF** evaporation stage. Following the **RF** frequency ramp, the **MT** is decompressed by linearly ramping down the coil current to zero in 0.9 s, transferring atoms into the **ODT**.

Once loaded, atoms are permitted to evaporate from the **ODT** similar to the forced **RF** evaporation described before (Sec. 2.1.3). In this case, the **ODT** beam power is gradually reduced, lowering the trap depth. The higher velocity atoms have sufficient kinetic energy to escape the trap, leaving the lower velocity atoms to collide and re-thermalize. The **ODT** evaporation stage is done in two stages; the first is an exponential ramp from full power to about 20% power, followed by a linear ramp to about 13% power [6]. The entire process takes between 2-4 s.

Following the exponential power ramp, atoms are still too hot to have Bose-condensed. It is during the linear ramp where the ensemble temperature becomes low enough to observe signs of quantum degeneracy. In both stages, atoms are imaged at long time-of-flight values, about 20 ms, in order to let the thermal background expand away from atoms which have condensed.

The key signature of atoms macroscopically filling their ground state is a bimodal density distribution. Thermal atoms follow a Gaussian density distribution, while a **BEC** approximately follows an inverted parabola (see Sec. 3.2.2) [9], which is more narrow than the thermal distribution. As the final power of the **ODT** beams is lowered, the emergence of a narrow parabolic peak above a Gaussian background may be observed, until eventually the Gaussian background is negligible, resulting in a **BEC** of $> 95\%$ purity. The **BECs** we produce typically have about 10^4 to 10^5 atoms. The temperature of the remaining thermal background is less than 40 nK. Atoms may be held in the **ODT** for > 5 s without significant losses, which permits a vast array of potential experiments with varying duration.

2.2 Dressed Atomic Ensembles

In most ultracold atom experiments, the key means through which the atomic Hamiltonian is engineered is by external radiation, be it a laser or other field. Ultracold atoms are prepared in ultrahigh vacuum chambers, typically suspended either by a magnetic field or laser field (as in an optical dipole trap, see Sec. 2.1). Therefore, essentially the only way to manipulate them is through external fields, with ultracold mixtures being the exception [17, 18]. As such, it is important we understand how these ensembles interact with radiation.

In this section we do not aim to provide a general description, but will instead focus on the relevant case of alkali atoms being driven by oscillating fields near an atomic resonance, colloquially referred to as *dressing*. We concern ourselves with transitions within a single hyperfine manifold F , between the sublevels $m_F = 0, \pm 1, \dots \pm F$. In particular we will first look at magnetic dipole transitions, driven through RF magnetic fields. We then find that electric dipole transitions from lasers, detuned from excited state resonances, may be treated in a similar way (with different selection rules, of course), and look at the case of bichromatic lasers driving *Raman* transitions.

2.2.1 RF Dressing

Consider the Hamiltonian for an atom with magnetic moment $\hat{\boldsymbol{\mu}}$ in an external magnetic field \mathbf{B} ,

$$\hat{H} = -\hat{\boldsymbol{\mu}} \cdot \mathbf{B}. \quad (2.5)$$

The total magnetic moment of the atom contains contributions from the electronic orbital angular momentum, spin, and the nuclear magnetic moment. We will assume that magnetic fields are weak, such that $\langle \hat{H} \rangle$ does not exceed the atomic hyperfine splitting; we may therefore use the projection theorem to project it along the total angular momentum $\hat{\mathbf{F}}$ [12, 19], yielding,

$$\hat{H} = g_F \mu_B \mathbf{B} \cdot \hat{\mathbf{F}}, \quad (2.6)$$

with μ_B the Bohr magneton, and g_F the total angular momentum g -factor. This Hamiltonian therefore describes the effect of the magnetic field \mathbf{B} on magnetic sublevels m_F within the hyperfine manifold F . Note that the angular momentum operators, such as $\hat{\mathbf{F}}$, have units of \hbar .

For time-independent fields, Eq. 2.6 is the Hamiltonian describing the linear Zeeman effect. Assuming the magnetic field is $\mathbf{B} = B_z \mathbf{e}_z$ (using the conventional notation in atomic physics for the unit vector along direction \mathbf{k} as \mathbf{e}_k), the Zeeman

energy for an atom in state $|m_F\rangle$ is,

$$\hbar\omega_Z = g_F\mu_B m_F B_z. \quad (2.7)$$

The linear Zeeman effect is used in experiment to split spins in energy through an applied magnetic field, or *bias* field, to create the resonance conditions for the time-dependent fields that follow. For alkali atoms the ground-state hyperfine splitting is about 1 GHz, in the microwave regime, and the Zeeman shifts typically seen in experiment are at radio-frequencies of about 1 MHz. The higher-order correction to the linear splitting, known as the quadratic Zeeman effect, is briefly covered in App. A.

We will now look at the effects of a time-dependent magnetic field, driving transitions between magnetic sublevels within a single F manifold. A bias field is applied along the z -axis, inducing a splitting $\hbar\omega_Z$. In addition to the static field, we apply a time-dependent magnetic field, $\mathbf{B} = \mathbf{B}(t)$, with a frequency close to resonance with the level-splitting. An oscillating magnetic field is applied with amplitude $B(t)$ and frequency ω_{RF} , which has the form,

$$\mathbf{B}_{\text{RF}}(t) = B(t) \sin[\omega_{\text{RF}}t + \phi(t)] \mathbf{e}_x, \quad (2.8)$$

where $\phi(t)$ is a time-dependent phase. We have allowed both the amplitude and phase of this field to vary in time, for generality. Substituting this into Eq. 2.6, together with the bias field providing the splitting $\hbar\omega_Z$, we obtain the Hamiltonian,

$$\hat{H}_{\text{RF}} = \Omega(t) \sin[\omega_{\text{RF}}t + \phi(t)] \hat{F}_x + \omega_Z \hat{F}_z, \quad (2.9)$$

where

$$\Omega(t) = \frac{\langle m_F | g_F \mu_B B(t) | m_F \pm 1 \rangle}{\hbar}, \quad (2.10)$$

is the Rabi frequency, defined through the coupling matrix element for adjacent m_F levels. A schematic of this coupling scheme for an arbitrary F manifold is shown in Fig. 2.2, where the coupling is between adjacent levels.

Typically, the coupling strength $\Omega \ll \omega_{\text{RF}}, \omega_Z$, and so it is convenient to transform into a basis rotating with the frequency of the driving field, called the *rotating frame*. More specifically, since we are allowing the driving field's phase $\phi(t)$ to vary in time, we will choose a basis rotating at frequency $\omega_{\text{RF}} + \partial_t \gamma(t)$, where $\gamma(t)$ is an arbitrary phase which may include some or all frequency modulation terms in $\phi(t)$; this gives us the freedom to choose a basis in which the terms in $\phi(t)$ show up as either frequency or phase shifts, which will be useful later in Ch. 5. The Hamiltonian

mation [20], ignoring terms which rotate at $2\omega_{\text{RF}}$. This leaves

$$\hat{\mathcal{H}}_{\text{R}} \simeq \frac{i\Omega}{4} \left[e^{i(\gamma-\phi)} \hat{F}_+ - e^{-i(\gamma-\phi)} \hat{F}_- \right] + \delta \hat{F}_z, \quad (2.16)$$

$$\delta = \omega_{\text{Z}} - (\omega_{\text{RF}} + \partial_t \gamma). \quad (2.17)$$

Rewriting this in terms of the spin operators again, using that $\hat{F}_{\pm} = \hat{F}_x \pm i\hat{F}_y$, we find,

$$\begin{aligned} \hat{\mathcal{H}}_{\text{R}} &= \frac{i\Omega}{4} \left[e^{i(\gamma-\phi)} (\hat{F}_x + i\hat{F}_y) - e^{-i(\gamma-\phi)} (\hat{F}_x - i\hat{F}_y) \right] + \delta \hat{F}_z, \\ &= \frac{\Omega}{2} \left\{ \frac{i}{2} [e^{i(\gamma-\phi)} - e^{-i(\gamma-\phi)}] \hat{F}_x - \frac{1}{2} [e^{i(\gamma-\phi)} + e^{-i(\gamma-\phi)}] \hat{F}_y \right\} + \delta \hat{F}_z. \end{aligned} \quad (2.18)$$

Finally, simplifying the complex exponentials and re-inserting the explicit time-dependencies, we have,

$$\hat{\mathcal{H}}_{\text{R}}(t) = \frac{\Omega(t)}{2} \left\{ \sin[\phi(t) - \gamma(t)] \hat{F}_x - \cos[\phi(t) - \gamma(t)] \hat{F}_y \right\} + \delta(t) \hat{F}_z. \quad (2.19)$$

The Hamiltonian (Eq. 2.19) describes the coupling of magnetic sublevels within the hyperfine manifold F , from a magnetic field with time-dependent amplitude and phase (frequency). Due to the linear nature of the Zeeman splitting, this Hamiltonian is expressed in terms of the spin algebra, and therefore generates transformations in $SU(2)$. The field amplitude $\Omega(t)$ controls the coupling strength, as it multiplies the \hat{F}_x and \hat{F}_y terms, while the phase $\phi(t)$ controls the phase of the coupling, tuning between \hat{F}_x and \hat{F}_y . The frequency of the field with respect to the Zeeman resonance, called the detuning $\delta(t)$, controls the dressed-state energies, through the \hat{F}_z term. We can see several example numerical simulations for different choices of these parameters in Fig. 2.3, demonstrating how this Hamiltonian may be used to generate interesting dynamics within the spin manifold. This includes an example state-preparation scheme, adiabatic rapid passage, in which an adiabatic change in the detuning through resonance may be used to efficiently transfer population from one state to another, as shown in Fig. 2.3(c).

The arbitrary phase term $\gamma(t)$ in Eq. 2.19 gives us some flexibility in choosing our basis. For instance, in the case where we have frequency modulation, we may choose $\gamma(t)$ appropriately so that the modulation appears in the Hamiltonian (Eq. 2.19) either as a frequency detuning in $\delta(t)$, or as a time-dependent phase; both cases are equivalent, but in certain scenarios one may be a more convenient representation than the other.

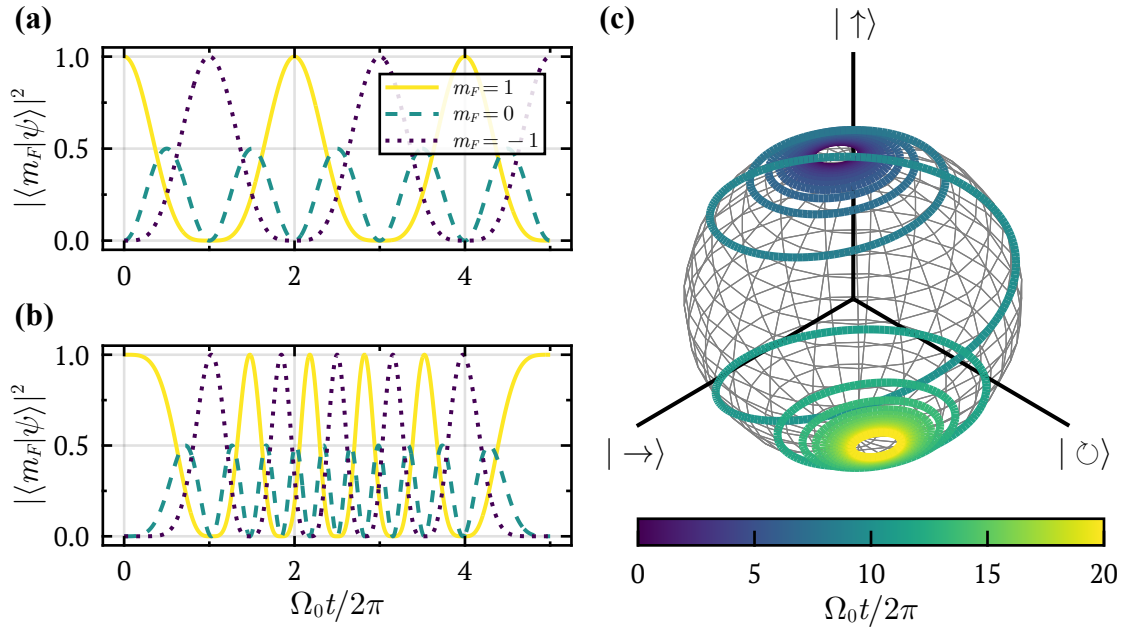


Figure 2.3: Example numerical solutions to Eq. 2.19, showing how the populations in each m_F level evolve over time. (a) Rabi oscillations in an $F = 1$ manifold, with $\Omega(t) = \Omega_0$, $\phi(t) = 0$, and $\delta(t) = 0$. Evolution of the spins is shown in the \hat{F}_z eigenbasis, over the duration $T = 10\pi/\Omega_0$. (b) The effect of amplitude modulation, with $\Omega(t) = \pi\Omega_0 \sin(\Omega_0 t/10)$, $\phi(t) = 0$, and $\delta(t) = 0$. The coupling strength is initially off, brought to a maximum of $\pi\Omega_0$ at $t = T/2$, and is then slowly reduced again until it is completely off at time T . (c) Simulation of adiabatic rapid passage [21–23], plotted on the Bloch-sphere for an $F = 1/2$ manifold. The coupling strength and phase are kept constant, $\Omega(t) = \Omega_0$, and $\phi(t) = 0$, while the detuning is varied as $\delta(t) = \delta_0 \sinh[(t - T/2)/30]$, with $\delta_0 = 10\Omega_0$. In this case, the detuning is therefore slowly varied from an initially large (with respect to Ω_0) positive value to a large negative value. The effect is that the eigenstates adiabatically follow the Hamiltonian, mapping $|m_F = +1/2\rangle \rightarrow |m_F = -1/2\rangle$, and vice versa. This demonstrates an effective way prepare different states with high fidelity and robustness to noise.

2.2.2 Raman Dressing

Similar to the case of RF-dressing above, we may use laser fields to generate a similar Hamiltonian to that in Eq. 2.19. With RF-dressing, time-dependent magnetic fields were used to directly couple m_F levels through a magnetic-dipole transition. Here, lasers induce electric dipole transitions, and due to selection rules may not directly couple the magnetic sublevels [12]. Instead, an effective coupling between m_F levels may be produced through intermediate excited states, by the absorption of light from one laser, and *stimulated* emission by the other; this is therefore a coherent process. Two or more laser fields with appropriate polarizations may couple m_F levels through a mutual excited state; if the lasers are detuned far from resonance with the excited state transition, the upper level may be adiabatically eliminated [20], allowing us to neglect them (although, losses due to spontaneous emission may be significant and are not properly considered in this reduced model [24]). Transitions of this nature are often referred to as two-photon Raman transitions, as the coherence between ground state levels is produced by the absorption and stimulated emission of two laser modes, though quantization of the fields is not necessary to describe the process.

In contrast to the RF-dressing discussed earlier, Raman dressing has several distinct differences. Electric-dipole transitions are much stronger than magnetic-dipole transitions, and therefore much larger Rabi frequencies (Eq. 2.10) may be obtained, with even modest laser powers. Furthermore, due to various combinations of beam polarizations, couplings between different sets of m_F levels may be engineered. Two common coupling schemes are shown in Fig. 2.4. The M -scheme, Fig. 2.4(a), uses one beam to induce either σ^+ or σ^- transitions, and another to induce π transitions, in order to couple adjacent m_F levels. Similarly, the Λ -scheme, Fig. 2.4(b), uses two beams with opposite circular polarizations to couple an m_F level with a $m_F \pm 2$ level.

Perhaps the most important distinction between RF-dressing and Raman dressing is the significant momentum transferred to atoms from laser fields; this can be safely ignored in the former, but must be accounted for in the latter. To illustrate this, consider the recoil energy $E_R = \hbar^2 k^2 / 2m$ imparted on an atom through a single scattering event with light having wavenumber k . To get some rough numbers, we assume a mass of $m = 10^{-26}$ kg, which is close to the mass of a Li atom, the lightest alkali to be Bose-condensed. For a 1 MHz RF carrier $k \approx 0.02 \text{ m}^{-1}$ and so $E_R/h \approx 10^{-12}$ Hz. Conversely, for a laser with wavelength $\lambda = 1000 \text{ nm}$, $k \approx 6 \times 10^6 \text{ m}^{-1}$, which yields $E_R/h \approx 100 \text{ kHz}$. The relevant energy scale of a BEC with which to make a comparison is the chemical potential μ , which is typically of the order $\mu/h \approx 1 \text{ kHz}$. As such, we can see that the Raman transitions impart significant mo-




Figure 2.4: Example Raman coupling schemes for an $F = 1$ spin manifold, with ground state Zeeman splitting ω_Z . The m_F levels are coupled through an intermediate excited state manifold F' , which has been adiabatically eliminated due to the large detuning Δ of the Raman lasers. Each scheme consists of two laser fields with strengths Ω_A and Ω_B , respectively. (a) Raman M -scheme, in which the polarization of the two laser fields are such that one induces σ^+ transitions, and the other π transitions. The result is a uniform coupling in the entire ground state manifold. (b) Raman Λ -scheme, in which the polarizations of the two lasers are such that they effect σ^\pm transitions. Due to selection rules, the $m_F = 0$ level is completely decoupled; the result is a pseudospin-1/2 system, comprised of the $m_F = \pm 1$ levels.

momentum to the atoms, which must be accounted for here. This momentum transfer will be useful in generating artificial gauge fields in Ch. 3.

Light Shifts as Effective Magnetic Fields

We begin by assuming that atoms in hyperfine manifold F are dressed by lasers which are detuned from any excited states, such that the excited levels may be adiabatically eliminated [20, 25], as described above. The atomic energies are altered by the incident electric fields through the scalar and vector light shifts, according to,

$$\hat{H}_{\text{LS}} = \hat{u}_s (\mathbf{E}^* \cdot \mathbf{E}) + \frac{i u_v (\mathbf{E}^* \times \mathbf{E})}{\hbar} \cdot \hat{\mathbf{J}}, \quad (2.20)$$

where $\hat{u}_s = u_s \hat{\mathbf{1}}$ and $u_v \propto u_s$ are scalar and vector light shifts, respectively; they are proportional to the ac polarizability of the atoms, which depends on the fine structure [25]. We will refrain from providing a discussion of these coefficients, but a detailed calculation may be found in Goldman et al. 2014 [25]; it is worth noting that u_s , and u_v have units of $\text{C}^2 \text{m}^2 \text{J}^{-1}$, so each of the terms in \hat{H}_{LS} have the correct units of energy.

The scalar light shift does not depend on the atomic state, while the vector light

shift is given in terms of the total electronic angular momentum $\hat{\mathbf{J}}$. We will assume instead that the vector shift acts on the total atomic angular momentum $\hat{\mathbf{F}}$, ignoring the extra contributions from the nucleus $\hat{\mathbf{I}}$ [26]. In this way, the vector light shift takes the form of the Zeeman Hamiltonian, Eq. 2.6, with the effective magnetic field,

$$\mathbf{B}_{\text{eff}} = \frac{i u_{\text{v}} (\mathbf{E}^* \times \mathbf{E})}{\mu_{\text{B}} g_J}, \quad (2.21)$$

where g_J is the total electronic angular momentum g -factor. In the presence of a bias magnetic field $\mathbf{B} = B_0 \mathbf{e}_z$, the Raman coupling Hamiltonian is therefore,

$$\hat{H}_{\text{L}} = \hat{u}_{\text{s}} (\mathbf{E}^* \cdot \mathbf{E}) + \frac{\mu_{\text{B}} g_J}{\hbar} (\mathbf{B} + \mathbf{B}_{\text{eff}}) \cdot \hat{\mathbf{F}}. \quad (2.22)$$

Bichromatic Laser Fields

Suppose we have two external laser fields with frequencies ω and $\omega + \omega_{\Delta}$, which we denote as Raman A and B, respectively. The total electric field produced by these beams is,

$$\mathbf{E} = \mathbf{E}_{\text{A}} e^{-i\omega t} + \mathbf{E}_{\text{B}} e^{-i[(\omega + \omega_{\Delta})t + \phi(t)]}, \quad (2.23)$$

where the Raman B laser has a relative phase $\phi(t)$ with respect to the Raman A beam; this may be accomplished by deriving both beams from the same laser, and frequency modulating them, as is common. The amplitude and polarization of the beams are contained within $\mathbf{E}_{\text{A}}(\mathbf{r}, t)$ and $\mathbf{E}_{\text{B}}(\mathbf{r}, t)$, which we will introduce later.

The effective magnetic field produced by the two lasers according to Eq. 2.21 is,

$$\mathbf{B}_{\text{eff}} = \mathbf{B}_{\delta} + \mathbf{B}_{\Omega} e^{i(\omega_{\Delta} t + \phi)} + \mathbf{B}_{\Omega}^* e^{-i(\omega_{\Delta} t + \phi)}, \quad (2.24)$$

where,

$$\mathbf{B}_{\delta} = \frac{i u_{\text{v}}}{\mu_{\text{B}} g_J} (\mathbf{E}_{\text{A}}^* \times \mathbf{E}_{\text{A}} + \mathbf{E}_{\text{B}}^* \times \mathbf{E}_{\text{B}}), \quad (2.25)$$

$$\mathbf{B}_{\Omega} = \frac{i u_{\text{v}}}{\mu_{\text{B}} g_J} (\mathbf{E}_{\text{B}}^* \times \mathbf{E}_{\text{A}}). \quad (2.26)$$

We will assume that the frequency difference between the lasers, ω_{Δ} , is close to the m_F level splitting, ω_Z , provided by the bias field. As in the case of RF-dressing, it is convenient to remove the dependence of ω_{Δ} by changing into the rotating frame, and applying the rotating wave approximation. As before (Eq. 2.11), the Hamiltonian \hat{H}_{L} (Eq. 2.22) transforms according to,

$$\hat{\mathcal{H}}_{\text{R}} = \hat{U}^{\dagger} \hat{H}_{\text{L}} \hat{U} - i \hbar \hat{U}^{\dagger} \partial_t \hat{U}, \quad (2.27)$$

with the unitary,

$$\hat{U}(t) = \exp \left\{ -\frac{i[\omega_\Delta t + \gamma(t)]}{\hbar} \hat{F}_z \right\}, \quad (2.28)$$

with $\gamma(t)$ an arbitrary phase, as in the case of RF-dressing (Sec. 2.2.1).

We will once again take advantage of the raising and lowering operators, \hat{F}_\pm , and \hat{F}_z , which transform as (App. B)

$$\hat{U}^\dagger \hat{F}_z \hat{U} = \hat{F}_z, \quad (2.29)$$

$$\hat{U}^\dagger \hat{F}_\pm \hat{U} = e^{\pm i(\omega_\Delta t + \gamma)} \hat{F}_\pm. \quad (2.30)$$

Rewriting \hat{H}_L in terms of these operators, we have,

$$\hat{H}_L = \hat{u}_s (\mathbf{E}^* \cdot \mathbf{E}) + \frac{\mu_{BGF}}{2\hbar} (\mathbf{B} + \mathbf{B}_{\text{eff}}) \cdot \left[(\mathbf{e}_x - i\mathbf{e}_y) \hat{F}_+ + (\mathbf{e}_x + i\mathbf{e}_y) \hat{F}_- + \mathbf{e}_z \hat{F}_z \right]. \quad (2.31)$$

Computing the transformation (Eq. 2.11) yields,

$$\begin{aligned} \hat{\mathcal{H}}_R = \hat{u}_s (\mathbf{E}^* \cdot \mathbf{E}) + \frac{\mu_{BGF}}{2\hbar} (\mathbf{B} + \mathbf{B}_{\text{eff}}) \cdot & \left[(\mathbf{e}_x - i\mathbf{e}_y) e^{i(\omega_\Delta t + \gamma)} \hat{F}_+ \right. \\ & \left. + (\mathbf{e}_x + i\mathbf{e}_y) e^{-i(\omega_\Delta t + \gamma)} \hat{F}_- \right] + \left[\frac{\mu_{BGF}}{\hbar} (\mathbf{B} + \mathbf{B}_{\text{eff}}) \cdot \mathbf{e}_z - (\omega_\Delta + \partial_t \gamma) \right] \hat{F}_z. \end{aligned} \quad (2.32)$$

Before applying the rotating wave approximation, we must substitute in the expressions for \mathbf{E} and \mathbf{B}_{eff} [20], which gives,

$$\begin{aligned} \hat{\mathcal{H}}_R = \hat{u}_s (\mathbf{E}_A^* \cdot \mathbf{E}_A + \mathbf{E}_B^* \cdot \mathbf{E}_B + \mathbf{E}_A^* \cdot \mathbf{E}_B e^{-i(\omega_\Delta t + \phi)} + \mathbf{E}_B^* \cdot \mathbf{E}_A e^{i(\omega_\Delta t + \phi)}) \\ + \frac{\mu_{BGF}}{2\hbar} \left[\mathbf{B}_\delta \cdot (\mathbf{e}_x - i\mathbf{e}_y) e^{i(\omega_\Delta t + \gamma)} \hat{F}_+ + \mathbf{B}_\delta \cdot (\mathbf{e}_x + i\mathbf{e}_y) e^{-i(\omega_\Delta t + \gamma)} \hat{F}_- \right. \\ + \mathbf{B}_\Omega \cdot (\mathbf{e}_x - i\mathbf{e}_y) e^{2i\omega_\Delta t} e^{i(\phi + \gamma)} \hat{F}_+ + \mathbf{B}_\Omega \cdot (\mathbf{e}_x + i\mathbf{e}_y) e^{i(\phi - \gamma)} \hat{F}_- \\ \left. + \mathbf{B}_\Omega^* \cdot (\mathbf{e}_x - i\mathbf{e}_y) e^{-i(\phi - \gamma)} \hat{F}_+ + \mathbf{B}_\Omega^* \cdot (\mathbf{e}_x + i\mathbf{e}_y) e^{-2i\omega_\Delta t} e^{-i(\phi + \gamma)} \hat{F}_- \right] \\ + \left\{ \frac{\mu_{BGF}}{\hbar} [B_0 + \mathbf{B}_\delta \cdot \mathbf{e}_z + \mathbf{B}_\Omega \cdot \mathbf{e}_z e^{i(\omega_\Delta t + \phi)} + \mathbf{B}_\Omega^* \cdot \mathbf{e}_z e^{-i(\omega_\Delta t + \phi)}] - (\omega_\Delta + \partial_t \gamma) \right\} \hat{F}_z. \end{aligned} \quad (2.33)$$

We then apply the rotating wave approximation, ignoring all terms which rotate with frequency ω_Δ , $2\omega_\Delta$, which leaves,

$$\begin{aligned} \hat{\mathcal{H}}_R \simeq \hat{u}_s (\mathbf{E}_A^* \cdot \mathbf{E}_A + \mathbf{E}_B^* \cdot \mathbf{E}_B) + \frac{\mu_{BGF}}{2\hbar} \left[+\mathbf{B}_\Omega \cdot (\mathbf{e}_x + i\mathbf{e}_y) e^{i(\phi - \gamma)} \hat{F}_- \right. \\ \left. + \mathbf{B}_\Omega^* \cdot (\mathbf{e}_x - i\mathbf{e}_y) e^{-i(\phi - \gamma)} \hat{F}_+ \right] + \left[\frac{\mu_{BGF}}{\hbar} (B_0 + \mathbf{B}_\delta \cdot \mathbf{e}_z) - (\omega_\Delta + \partial_t \gamma) \right] \hat{F}_z. \end{aligned} \quad (2.34)$$

Similar to the RF-dressing case, the detuning (Eq. 2.17) is then defined as,

$$\delta(t) = \frac{\mu_B g_F}{\hbar} B_0 - [\omega_\Delta + \partial_t \gamma(t)], \quad (2.35)$$

and the scalar light shift,

$$\hat{\mathcal{V}}_s(\mathbf{r}, t) = \hat{u}_s (\mathbf{E}_A^* \cdot \mathbf{E}_A + \mathbf{E}_B^* \cdot \mathbf{E}_B). \quad (2.36)$$

We may then write $\hat{\mathcal{H}}_R$ in a more compact form as,

$$\hat{\mathcal{H}}_R = \hat{\mathcal{V}}_s + \boldsymbol{\Omega} \cdot \hat{\mathbf{F}}, \quad (2.37)$$

with,

$$\boldsymbol{\Omega} \cdot \hat{\mathbf{F}} = \frac{\Omega^*}{2} \hat{F}_+ + \frac{\Omega}{2} \hat{F}_- + \Omega_z \hat{F}_z, \quad (2.38)$$

$$\Omega = \frac{\mu_B g_F}{\hbar} \mathbf{B}_\Omega \cdot (\mathbf{e}_x + i\mathbf{e}_y) e^{i(\phi - \gamma)}, \quad (2.39)$$

$$\Omega_z = \frac{\mu_B g_F}{\hbar} (\mathbf{B}_\delta \cdot \mathbf{e}_z) - \delta. \quad (2.40)$$

Here we've expressed $\boldsymbol{\Omega}$ in the *circular* basis, that is, in terms of the circular polarizations of \mathbf{B}_{eff} , $\mathbf{e}_x \pm i\mathbf{e}_y$, and the component oscillating along the bias field axis, \mathbf{e}_z . Looking at the operators accompanying the components of \mathbf{B}_{eff} , we can see that the circular parts effect raising and lowering operations on the m_F states, \hat{F}_+ and \hat{F}_- , respectively. The component of \mathbf{B}_{eff} parallel to the bias field affects the m_F level splitting, through \hat{F}_z .

While the circular basis gives us some intuition for how the various components of \mathbf{B}_{eff} may couple m_F levels, as in the RF-dressing case, it is generally more convenient to work in the conventional spin, or *Cartesian* basis. Rewriting $\hat{\mathcal{H}}_R$ (Eq. 2.37) in this basis, we have,

$$\begin{aligned} \hat{\mathcal{H}}_R &= \hat{\mathcal{V}}_s + \frac{\Omega + \Omega^*}{2} \hat{F}_x - i \frac{\Omega - \Omega^*}{2} \hat{F}_y + \Omega_z \hat{F}_z, \\ &= \hat{\mathcal{V}}_s + \Re(\Omega) \hat{F}_x + \Im(\Omega) \hat{F}_y + \Omega_z \hat{F}_z, \end{aligned} \quad (2.41)$$

with $\Re(\tilde{z})$ the real and $\Im(\tilde{z})$ the imaginary parts of \tilde{z} , respectively. From this, we may therefore express the various terms of the Hamiltonian (Eq. 2.37) in whichever basis.

M-Scheme Configuration

To complete the analogue between Raman and RF dressing, we will look at a particular spatial configuration of Raman lasers to arrive at a Hamiltonian which is

equivalent to that of the RF-dressing case, Eq. 2.19. This configuration is a variant of the M -scheme depicted in Fig. 2.4(a). The spatial parts of the electric fields for Raman A and B are,

$$\mathbf{E}_A(\mathbf{r}, t) = \tilde{E}_A(t)e^{ik_L x}\mathbf{e}_z, \quad (2.42)$$

$$\mathbf{E}_B(\mathbf{r}, t) = \tilde{E}_B(t)e^{ik_L x}\mathbf{e}_y, \quad (2.43)$$

where we have chosen both beams to be parallel along the x -direction, with orthogonal linear polarizations. Raman A drives π -transitions, with its field oriented along the quantization axis defined by the bias field \mathbf{B}_0 . Raman B drives both σ^\pm transitions, with its linear polarization orthogonal to the quantization axis. Since both beams are parallel, after the scattering of light from one beam and stimulated emission into the other by atoms, there is no net momentum transfer to the atoms; this will result in a coupling scheme similar to in the RF-dressing case, where the momentum was negligible. In a later chapter (Ch. 3) we will look at a different configuration in which the momentum transferred to atoms through this process is the essential ingredient in generating artificial gauge fields.

To compute the Hamiltonian $\hat{\mathcal{H}}_R$ (Eq. 2.37), we start with the effective magnetic fields,

$$\mathbf{B}_\Omega = \frac{i u_\nu}{\mu_B g_J} \tilde{E}_A \tilde{E}_B \mathbf{e}_x, \quad (2.44)$$

$$\mathbf{B}_\delta = 0. \quad (2.45)$$

From this, we have,

$$\Omega = \frac{i u_\nu g_F}{\hbar g_J} \tilde{E}_A \tilde{E}_B e^{i(\phi-\gamma)}, \quad (2.46)$$

$$\Omega_z = -\delta. \quad (2.47)$$

Next, the scalar light shift term is,

$$\hat{\mathcal{V}}_s = \hat{u}_s \left(\tilde{E}_A^2 + \tilde{E}_B^2 \right). \quad (2.48)$$

The Hamiltonian is therefore,

$$\begin{aligned} \hat{\mathcal{H}}_R(t) = \hat{u}_s \left(\tilde{E}_A^2(t) + \tilde{E}_B^2(t) \right) + \frac{u_\nu g_F}{\hbar g_J} \tilde{E}_A(t) \tilde{E}_B(t) \left\{ \sin[\phi(t) - \gamma(t)] \hat{F}_x \right. \\ \left. + \cos[\phi(t) - \gamma(t)] \hat{F}_y \right\} - \delta(t) \hat{F}_z, \quad (2.49) \end{aligned}$$

which is equivalent to the RF-dressed Hamiltonian (Eq. 2.19) up to a phase. We may

therefore use two laser fields to affect the same dynamics as an RF field, but with potentially much larger effective Rabi frequencies due to the relative higher strength of electric dipole transitions, when compared to magnetic dipole transitions.

2.3 Spin-Ensemble Measurement

After using either RF (Sec. 2.2.1) or Raman (Sec. 2.2.2) dressing to couple m_F levels, we need a means to perform measurements within the spin manifold. Depending on the specific experiment we are trying to perform, there are various other requirements we may have, such as resolving the spatial structure of the spins (See Ch. 3). Here we will focus on a method for performing projective measurements in the spin manifold, discerning the relative populations between spins in time-of-flight (TOF) measurements. Together with the spin tomography described in App. C, we are able to perform a series of such projective measurements in order to fully reconstruct an arbitrary spin wavefunction.

2.3.1 Stern-Gerlach TOF Measurements

A useful resource in spin-sensitive measurement, which we have already taken advantage of for the purposes of the MOT (Sec. 2.1.2) and MT (Sec. 2.1.3), is the effect of magnetic field gradients on the m_F levels, which we wish to resolve. A magnetic field gradient may be applied during TOF, resulting in a spin-dependent force on atoms through the linear Zeeman effect (Eq. 2.6). Over time, the spins spatially separate, and are individually imaged; this is a manifestation of the Stern-Gerlach (SG) effect, in ultracold atoms.

More specifically, consider a magnetic field gradient along the z -axis during TOF, $\mathbf{B} = B(z)\mathbf{e}_z$. Using the linear Zeeman effect once again, Eq. 2.6, the classical force on the m_F levels is,

$$\begin{aligned}\mathbf{F}_B &= -g_F m_F \mu_B \nabla B, \\ &= -g_F m_F \mu_B \partial_z B \mathbf{e}_z,\end{aligned}\tag{2.50}$$

which points along the z -axis according to the sign of $g_F m_F$. As before, this applies individually to each atom in the ensemble; we are ignoring many-body effects here.

In our system, we use the main gradient coils used in the MOT and MT stages (Sec. 2.1) which applies a gradient in the vertical direction. Therefore, as atoms fall in TOF they spatially separate. By controlling both the duration of the fall and the strength of the gradient, we may capture each spin within a single absorption image. For a set of sample images, see Fig. 2.5.




Figure 2.5: Series of SG-TOF images, demonstrating an RF-coupled BEC in the $F = 1$ manifold (See Sec. 2.2.1). The m_F levels are separated in space as they fall in a magnetic field gradient (Eq. 2.50). Gravity points downwards in the imaging plane. Each image is a destructive measurement of both the state and the BEC itself, so each image in the series consists of a completely new BEC produced, and dropped in TOF after an RF pulse with a different duration; together they show a full Rabi oscillation over about $50 \mu\text{s}$. The “hole” appearing in the middle of several BEC images are an imaging defect, due to the diffraction of the imaging laser around the BEC with high optical depth.

Imaging atoms with a TOF greater than about 25 ms is difficult, as they reach the edge of the imaging beam in this time; as such, the SG gradient must be applied in a time shorter than this. Conversely, due to the inherent difficulties in quickly switching on a large amount of current, we are unable to apply a gradient in an amount of time shorter than about 12 ms. Despite these restrictions, we find that during a 20 ms TOF, a current pulse roughly 10 ms in duration, with an amplitude comparable to that used in the uncompressed MT (27 G/cm, see Sec. 2.1.3), the m_F components separate enough that there is no overlap between the expanded density distributions for either the $F = 1$ or $F = 2$ manifolds. The components may therefore be individually fit, discerning the atom numbers and thus the relative populations.

This technique amounts to a projective measurement in the eigenbasis of \hat{F}_z , where the relative populations give the amplitudes for each eigenstate. Oftentimes, it is also important to determine the relative phases of the wavefunction, for which quantum state tomography is required (App C). Measurements must be made in other bases for this to be accomplished, so the SG technique has to be adapted to more than just projections into the z -basis.

The simplest way to imagine this would be to generate a field gradient along other axes; a gradient along x -would, in principle, correspond to a projection onto the eigenbasis of \hat{F}_x . However, carefully controlling gradients in this way would either require additional coils along each projection axis, or the ability to displace

the center of the main coil gradient consistently; both approaches provide unique challenges. The former would require many coils along non-trivial directions, since spin tomography in manifolds $F > 1/2$ require many different basis measurements (App. C). The alignment of these coils would need to be carefully calibrated for, and could change with each new optimization of the BEC due to periodic alterations of the MOT and ODT alignment. Similarly, the latter could be accomplished by the appropriate application of bias fields to shift the minimum of the gradient, but this would be sensitive to the coil currents, and background fields. Both approaches would also suffer from the fact that the BEC falls in TOF through the field, so when the gradient is oriented at some arbitrary angle the fields sampled could alter the results.

To overcome these challenges, we use RF pulses to change the measurement basis prior to TOF. Using the RF-dressed Hamiltonian, Eq. 2.19, we may apply a short pulse that maps the eigenbasis of the operator \hat{F}_v to those of \hat{F}_z , where v is the basis in which we want to measure.

To see this we first look at the simplest case, when $F = 1/2$. We may apply resonant RF pulses with varied phase and duration (pulse area) to produce the desired transformations. From the Hamiltonian, Eq. 2.19, we thus set the amplitude of the pulses to Ω , detuning $\delta = 0$, and phase reference $\gamma = 0$.

Looking first at the eigenstates of \hat{F}_x , expressed in the z -basis they are,

$$|\rightarrow\rangle = \frac{1}{\sqrt{2}} (|\uparrow\rangle + |\downarrow\rangle), \quad (2.51)$$

$$|\leftarrow\rangle = \frac{1}{\sqrt{2}} (|\uparrow\rangle - |\downarrow\rangle), \quad (2.52)$$

where $|\uparrow\rangle$, and $|\downarrow\rangle$ are the eigenstates of \hat{F}_z . To measure in this basis, an RF pulse with phase $\phi = 0$ and duration $T = \pi/\Omega$ yields the evolution operator $\hat{U}_{\text{RF}}^x = \exp(-i\pi\hat{F}_y/2\hbar)$. This pulse therefore transforms the x -eigenstates as,

$$\hat{U}_{\text{RF}}^x |\rightarrow\rangle = |\downarrow\rangle, \quad (2.53)$$

$$\hat{U}_{\text{RF}}^x |\leftarrow\rangle = |\uparrow\rangle. \quad (2.54)$$

Hence, the populations in each of these eigenstates map to those of \hat{F}_z . Following a SG-TOF measurement, the resulting spin populations therefore correspond to those in the x -basis. Similarly, the eigenstates of \hat{F}_y expressed in the canonical basis are,

$$|\odot\rangle = \frac{1}{\sqrt{2}} (|\uparrow\rangle + i|\downarrow\rangle), \quad (2.55)$$

$$|\ominus\rangle = \frac{1}{\sqrt{2}} (|\uparrow\rangle - i|\downarrow\rangle). \quad (2.56)$$

Applying an RF pulse with phase $\phi = \pi/2$ and duration $T = \pi/\Omega$ produces the evolution operator $\hat{\mathcal{U}}_{\text{RF}}^y = \exp\left(i\pi\hat{F}_x/2\hbar\right)$. The basis therefore transforms as,

$$\hat{\mathcal{U}}_{\text{RF}}^y |\circ\rangle = i |\downarrow\rangle, \quad (2.57)$$

$$\hat{\mathcal{U}}_{\text{RF}}^y |\ominus\rangle = |\uparrow\rangle, \quad (2.58)$$

which coincide with the eigenbasis of \hat{F}_z up to a global phase. Hence, the populations in each of the y -eigenstates map to the z -eigenstates. For spin state tomography in an $F = 1/2$ manifold, these are the only measurements necessary to full reconstruct an arbitrary spin state (App. C).

For higher spin systems the same procedure holds, in principle, but discerning a set of informationally complete basis measurements is more difficult. In general, pulses with varied phases and pulse areas are required. The RF pulses only generate transformations in $SU(2)$, which limits the available bases; the ability to perform transformations in $SU(2F + 1)$ would likely make this process easier, but such coupling schemes are significantly more complicated, and an active area of research [27].

CHAPTER 3

Simulating Gauge-Fields in BECs

THE eager experimentalist is always in search of new tools to put in their toolbox, though not every addition must be a literal tool or new device. The apparatus, and [RF](#) and laser-dressing methods described in the previous chapter, [Ch. 2](#), have introduced some very powerful tools already at our disposal. Access to macroscopic quantum systems such as [BECs](#) is highly sought after, and our apparatus is capable of manipulating such a system in various different ways. With the aim of studying interesting artificial gauge fields in ultracold atoms, however, we are always in search of new techniques, methods of quantum control, or theoretical insights. In this chapter, we will discuss two such additions to our metaphorical toolbox, which are demonstrated together.

We begin in [Sec. 3.1](#) by describing how the laser-dressing technique described in [Sec. 2.2.2](#) may be used to generate Abelian artificial gauge fields in ultracold ensembles [[28–30](#)]. We begin by adopting a particular spatial and polarization configuration of the Raman lasers, in which the momentum transferred to atoms by two coherent scattering events becomes important. The Hamiltonian is well represented in a new basis, where the spin and momentum degrees of freedom are coupled. The resulting spectra is well described by two bands, consisting of mixtures between the spin components at different momenta. We go on to describe in [Sec. 3.1.2](#) how in certain parameter regimes, the dynamics of the neutral atoms is analogous to that of a charged particle under the influence of Abelian electromagnetic gauge fields.

With the aim of implementing such a scheme in an ultracold ensemble, in [Sec. 3.2](#) we discuss the mean-field properties of [BECs](#), particularly through the [Gross-Pitaevskii equation \(GPE\)](#), which serves as the nonlinear Schrödinger equation for interacting gases. We derive the time-independent [GPE](#) by minimizing the mean-field energy functional, demonstrating how its solutions are the ground states of such a system. We then describe the [Thomas-Fermi \(TF\)](#) approximation, which allows us to calculate several important properties of the many-body ground states

with relative simplicity, and reasonable accuracy. We end this section by briefly discussing the generalization of these results to spinor BECs, which will be necessary to consider in the study of Raman-dressed artificial gauge fields.

In Sec. 3.3, we then describe a numerical technique for finding solutions to the GPE, both for ground states and real-time dynamics. After describing the algorithm, as implemented in a custom Python package [1], we discuss its extension to spinor BECs with Raman dressing, enabling us to simulate the effects of the artificial gauge fields described in the prior sections. Simulations of this nature are computationally expensive, so we finish this section by describing how the program was accelerated on a graphics processing unit (GPU), resulting in significant increases to performance. We discuss some nuances of hardware acceleration, and show some results comparing the performance across several available devices [1].

Finally, we end this chapter by combining the subjects of all the previous sections into a series of simulations. Specifically, we use the GPU accelerated spinor GPE solver to simulate the many-body ground states of a BEC, subjected to the Raman-dressing described in Sec. 3.1. The Raman parameters are chosen so as to produce an artificial magnetic field in the BEC, which produces circulating currents of atoms, resulting in interesting vortex structures [31]. We offer some intuition behind the ground state solutions for various choices of parameters, before providing some concluding remarks.

3.1 Producing Abelian Gauge Fields in BECs

In Sec. 2.2.2 we described a means to use bichromatic laser fields to dress the hyperfine manifolds of alkali-metal atoms, to induce coupling between spins. For co-propagating lasers, we showed that the resulting Hamiltonian, Eq. 2.49, is remarkably similar to that of RF-dressed spins, Eq. 2.19 (Sec. 2.2.1). In the case of Raman coupling, the momentum transferred to atoms from the laser fields cancels out; this is because the coherence between spins is produced by absorption of light from one laser, followed by *stimulated* emission into the other. For co-propagating beams, the momenta therefore cancel¹. If the beams are not co-propagating, then this momentum transfer will become important; in fact, it is useful in engineering various artificial gauge fields of interest, as described in the seminal work of *Lin* et al. 2009 [28], and the studies that followed [29, 32–34].

In this section we will describe how, for a particular configuration of Raman beams, we may produce simple artificial gauge fields in an ultracold ensemble, start-

¹Technically there is a difference between wavevectors of each laser, as they have different frequencies. Typically, this difference is on the order of about 1 MHz, and so this momentum is negligible as discussed in Sec. 2.2.2.

ing from the laser-dressed Hamiltonian, Eq. 2.37. Due to the momentum transfer from the Raman lasers, we will find a convenient basis in which to represent the full atomic Hamiltonian; the dynamics in this basis are well understood by the presence of an Abelian vector and scalar potential, as from classical electrodynamics. Through this coupling scheme, the physics of neutral atoms is made equivalent to charged particles moving in gauge potentials [28–30, 32–35].

3.1.1 Spin-Momentum Coupling by Raman Dressing

Consider the laser-dressed Hamiltonian, Eq. 2.37, which was previously derived in Sec. 2.2.2. This Hamiltonian describes an atom in the ground state hyperfine manifold F , subjected to a bias magnetic field $\mathbf{B} = B_0 \mathbf{e}_z$, and two laser fields defined through their electric fields. The lasers have a frequency difference that is similar to the Zeeman splitting induced by the bias field, allowing us to move into a rotating frame where the Hamiltonian is stationary². Starting from this result, we will now consider the case where the lasers are counter-propagating, with opposite linear polarizations. The spatial parts of the electric fields are,

$$\mathbf{E}_A(\mathbf{r}, t) = \tilde{E}_A(t) e^{ik_L x} \mathbf{e}_y, \quad (3.1)$$

$$\mathbf{E}_B(\mathbf{r}, t) = \tilde{E}_B(t) e^{-ik_L x} \mathbf{e}_z. \quad (3.2)$$

The first laser, Raman A, propagates along the positive x -axis with a linear polarization along y ; this laser will therefore induce σ^\pm transitions in the atoms. The second beam, Raman B, propagates along the negative x -axis with a linear polarization along z ; this field therefore induces π -transitions in atoms. Both beams are assumed to have the same wavevector, k_L . Together, these beams form a Raman M -scheme (see Fig. 2.4). From these, the effective magnetic fields are (Sec. 2.2.2),

$$\begin{aligned} \mathbf{B}_\Omega &= \frac{i u_v}{\mu_B g_J} (\mathbf{E}_B^* \times \mathbf{E}_A), \\ &= \frac{i u_v}{\mu_B g_J} \tilde{E}_A \tilde{E}_B e^{2ik_L x} \mathbf{e}_x, \end{aligned} \quad (3.3)$$

and $\mathbf{B}_\delta = 0$. The components of the vector $\boldsymbol{\Omega}$ are thus,

$$\begin{aligned} \Omega &= \frac{\mu_B g_F}{\hbar} \mathbf{B}_\Omega \cdot (\mathbf{e}_x + i \mathbf{e}_y) e^{i(\phi - \gamma)}, \\ &= -\frac{i u_v g_F}{\hbar g_J} \tilde{E}_A \tilde{E}_B e^{2ik_L x} e^{i(\phi - \gamma)}, \end{aligned} \quad (3.4)$$

²We have allowed the field amplitudes and phases to change in time still.

and $\Omega_z = -\delta$. We have retained the phase difference between the beams, ϕ , and the term γ that describes the freedom to choose the phase of the rotating basis.

The Raman coupling Hamiltonian, expressed in the circular basis of \hat{F}_\pm and \hat{F}_z , is therefore,

$$\hat{\mathcal{H}}_R = \hat{\mathcal{V}}_s + \frac{i u_v g_F}{2 \hbar g_J} \tilde{E}_A \tilde{E}_B \left[e^{-2i k_L x} e^{-i(\phi-\gamma)} \hat{F}_+ - e^{2i k_L x} e^{i(\phi-\gamma)} \hat{F}_- \right] - \delta \hat{F}_z, \quad (3.5)$$

with $\hat{\mathcal{V}}_s = \hat{u}_s \left(\tilde{E}_A^2 + \tilde{E}_B^2 \right)$ the scalar light shift from the beams. The $e^{\pm 2i k_L x}$ indicates a rotation in phase with motion along the x -axis, on a length scale defined by k_L . We may think of the Raman beams imprinting a phase grating on the spins, where the spin quantization axis varies along x . We will move into a basis that rotates with this quantization axis, through the unitary,

$$\hat{U}(\hat{x}) = \exp \left[-\frac{i}{\hbar} \left(2k_L \hat{x} - \frac{\pi}{2} \right) \hat{F}_z \right]. \quad (3.6)$$

Note that we have promoted the position coordinate to an operator \hat{x} ; previous to this, the position and momentum could be neglected, as all of the coupling dynamics were described by the spin degrees of freedom. Since this basis transformation depends on space, there will be a corresponding shift in the momentum of atoms, in a way that depends on spin. We will no longer be able to separate translational and spin degrees of freedom. As such, we will need to consider the kinetic and potential energies of the atoms as well.

For now we will continue with the Raman coupling Hamiltonian, which transforms as $\hat{\mathcal{H}}_R^\circ = \hat{U}^\dagger(\hat{x}) \hat{\mathcal{H}}_R \hat{U}(\hat{x})$, where the “ \circ ” superscript denotes the new basis rotating in space. As before, the operators \hat{F}_\pm and \hat{F}_z transform according to (App. B),

$$\hat{U}^\dagger(\hat{x}) \hat{F}_z \hat{U}(\hat{x}) = \hat{F}_z, \quad (3.7)$$

$$\hat{U}^\dagger(\hat{x}) \hat{F}_\pm \hat{U}(\hat{x}) = e^{\pm i(2k_L \hat{x} - \pi/2)} \hat{F}_\pm. \quad (3.8)$$

The Raman coupling Hamiltonian becomes,

$$\begin{aligned} \hat{\mathcal{H}}_R^\circ &= \hat{\mathcal{V}}_s + \frac{i u_v g_F}{2 \hbar g_J} \tilde{E}_A \tilde{E}_B \left[e^{-2i k_L x} e^{-i(\phi-\gamma)} e^{i(2k_L \hat{x} - \pi/2)} \hat{F}_+ \right. \\ &\quad \left. - e^{2i k_L x} e^{i(\phi-\gamma)} e^{-i(2k_L \hat{x} - \pi/2)} \hat{F}_- \right] - \delta \hat{F}_z, \\ &= \hat{\mathcal{V}}_s + \frac{u_v g_F}{2 \hbar g_J} \tilde{E}_A \tilde{E}_B \left[e^{-i(\phi-\gamma)} \hat{F}_+ + e^{i(\phi-\gamma)} \hat{F}_- \right] - \delta \hat{F}_z. \end{aligned} \quad (3.9)$$

Substituting in for \hat{F}_\pm , we find in the Cartesian spin basis,

$$\hat{\mathcal{H}}_R^\circ = \hat{\mathcal{V}}_s + \frac{u_v g_F}{\hbar g_J} \tilde{E}_A \tilde{E}_B \left[\cos(\Phi - \gamma) \hat{F}_x + \sin(\phi - \gamma) \hat{F}_y \right] - \delta \hat{F}_z. \quad (3.10)$$

Thus, in this basis, the Raman coupling once again looks similar to the case of co-propagating beams, or of RF-dressing (Eq. 2.19). However, as stated prior, due to the positional dependence of this basis, we must also consider the kinetic and potential energies of the atoms.

The full atomic Hamiltonian is,

$$\hat{H}_F = \frac{\hbar^2 \hat{\mathbf{k}}^2}{2m} \hat{\mathbf{1}} + \hat{V}(\hat{\mathbf{r}}) \hat{\mathbf{1}} + \hat{\mathcal{H}}_R. \quad (3.11)$$

The first term represents the kinetic energy, with $\hbar \hat{\mathbf{k}}$ the atomic momentum operator, and m the mass. We assume that atoms are confined in an external potential $\hat{V}(\hat{\mathbf{r}})$. Neither of these terms depend on spin, so they are $\propto \hat{\mathbf{1}}$. In the new basis defined by $\hat{U}(\hat{x})$, the Hamiltonian becomes $\hat{H}_F^\circ = \hat{U}^\dagger(\hat{x}) \hat{H}_F \hat{U}(\hat{x})$. We have already computed how the Raman coupling Hamiltonian transforms, and since the potential energy commutes with both \hat{x} and \hat{F}_z , it remains unchanged. For unitary transformations of the form $\hat{U}(\hat{\mathbf{r}}) = \exp(-i \mathbf{q} \cdot \hat{\mathbf{r}}/\hbar)$, we may make use of the identity,

$$\hat{U}^\dagger(\hat{\mathbf{r}}) \hat{H}(\hat{\mathbf{p}}) \hat{U}(\hat{\mathbf{r}}) = \hat{H}(\hat{\mathbf{p}} - \mathbf{q}) \quad (3.12)$$

for constant vector \mathbf{q} , and with $\hat{\mathbf{p}}$ the momentum operator. The kinetic energy therefore becomes,

$$\hat{U}^\dagger(\hat{x}) \frac{\hbar^2 \hat{\mathbf{k}}^2}{2m} \hat{U}(\hat{x}) = \frac{\hbar^2}{2m} \left(\hat{k}_x \hat{\mathbf{1}} - 2k_L \hat{F}_z/\hbar \right)^2 + \frac{\hbar^2}{2m} \left(\hat{k}_y^2 + \hat{k}_z^2 \right) \hat{\mathbf{1}}. \quad (3.13)$$

In this basis, the full Hamiltonian is therefore,

$$\begin{aligned} \hat{H}_F^\circ = \frac{\hbar^2}{2m} \left(\hat{k}_x \hat{\mathbf{1}} - 2k_L \hat{F}_z/\hbar \right)^2 + \frac{u_v g_F}{\hbar g_J} \tilde{E}_A \tilde{E}_B \left[\cos(\Phi - \gamma) \hat{F}_x + \sin(\phi - \gamma) \hat{F}_y \right] \\ - \delta \hat{F}_z + \hat{\mathcal{V}}_s + \frac{\hbar^2}{2m} \left(\hat{k}_y^2 + \hat{k}_z^2 \right) \hat{\mathbf{1}} + \hat{V}(\hat{\mathbf{r}}) \hat{\mathbf{1}}, \end{aligned} \quad (3.14)$$

In this configuration, \hat{k}_y and \hat{k}_z remain separable from spin, and the trapping potential is also unchanged. We can see in the above that there is a mixture of spin and momentum degrees of freedom, along the x -axis, due to the momentum imparted on atoms by the Raman lasers through the coherent scattering process. Note that, for the remainder of this thesis, we will refer to $\hat{\mathbf{k}}$ directly as momentum, since $\hat{\mathbf{k}} = \hat{\mathbf{p}}/\hbar$. Intuitively, an atom in some m_F level may change its state to $m_F \pm 1$ through the

absorption and spontaneous emission of light from the Raman laser fields, just as in the previous coupling schemes (Sec. 2.2); $\hat{\mathcal{H}}_R^\circ$ is the same as in the co-propagating case, describing the coupling of adjacent spins in the manifold. The difference here is that, by undergoing this process, the momentum of the atom must also change by $\pm 2\hbar k_L$ (depending on whether it increased or decreased its spin), due to the recoil from the two scattering events required to make such a change; k_L is called the *recoil* momentum for this reason. Therefore, to understand this Hamiltonian more completely we must consider momentum, as well as spin.

3.1.2 Interpretation as Artificial Gauge Field

Looking at the Raman-dressed Hamiltonian, Eq. 3.14, we can focus on the spin-momentum coupled terms, *i.e.* those that are not proportional to $\hat{\mathbb{1}}$. We will make the additional assumption that the phase difference between the Raman beams, $\phi = 0$, and similarly $\gamma = 0$; these parameters only act to define the phase reference of the coupling, which is unimportant here. Last, for simplicity, we will assume that the spin manifold $F = 1/2$; in this case, we can replace the spin matrices with the Pauli matrices by $\hat{F}_q = \hbar\hat{\sigma}_q/2$, $q \in \{x, y, z\}$. The Hamiltonian becomes,

$$\hat{\mathcal{H}}_F^\circ = \frac{\hbar^2}{2m} \left(\hat{k}_x \hat{\mathbb{1}} - k_L \hat{\sigma}_z \right)^2 + \frac{\hbar\Omega}{2} \hat{\sigma}_x - \frac{\hbar\delta}{2} \hat{\sigma}_z, \quad (3.15)$$

where the detuning,

$$\delta = \frac{\mu_B g_F}{\hbar} B_0 - \omega_\Delta, \quad (3.16)$$

is the difference between the Zeeman splitting due to the bias magnetic field, and the frequency difference between the lasers. Similarly, we define the Rabi frequency of the Raman coupling as,

$$\Omega = \frac{u_v g_F}{\hbar g_J} \tilde{E}_A \tilde{E}_B. \quad (3.17)$$

Due to the spin-momentum coupling, the eigenstates of this Hamiltonian will be labeled by both degrees of freedom. We may understand it further by computing the eigenvalues as a function of k_x (*i.e.* the dispersion curves), for different choices of Ω and δ . The recoil momentum k_L imposes a natural scale for momentum, and energy; it is convenient then to work in units of k_L for momentum, and for energy we define the corresponding recoil energy $E_L = \hbar^2 k_L^2 / 2m$.

First, in the absence of Raman fields, $\Omega = \delta = 0$, only the kinetic energy term survives, but the spins are still shifted in momentum by $2k_L$ with respect to each other, due to our choice of basis. Looking at the dispersion, we see two parabolas corresponding to the nominal kinetic energies of the atoms, with their respective shifts, as shown by the curves in Fig. 3.1 for $\Omega = 0$. Each spin has equal and

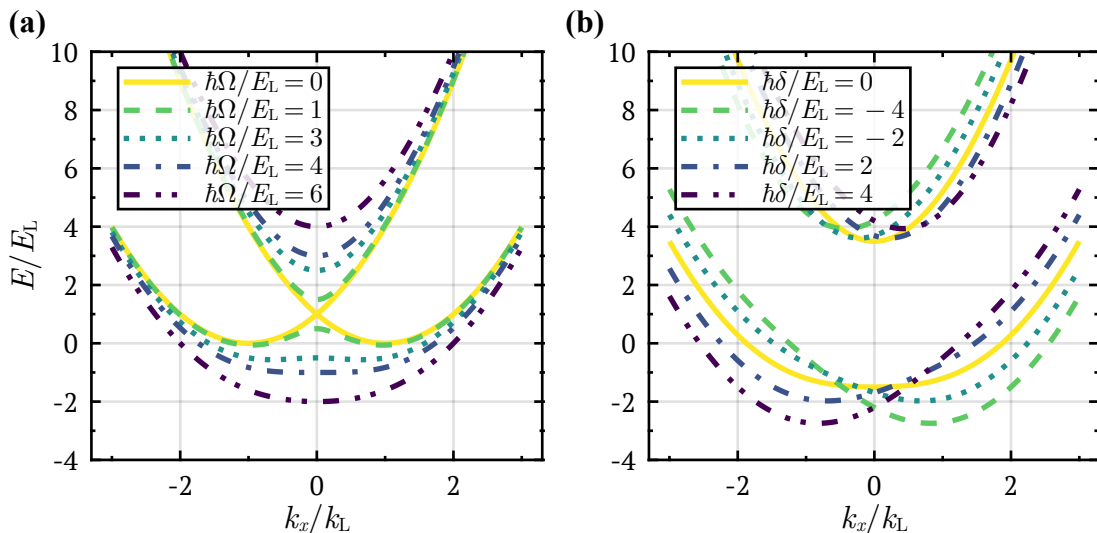


Figure 3.1: Energy eigenvalues of the Hamiltonian Eq. 3.15, for various values of Ω and δ , as a function of the momentum k_x ; results are reminiscent of band structure in crystals. (a) For $\Omega = 0$, the dispersion consists of the bare atomic kinetic energies, separated by $2k_L$ due to our choice of basis. For $\Omega > 0$, an avoided-crossing between bands opens, separating into an upper and lower curve, where the spacing between upper and lower bands at $k_x = 0$ is Ω . If $\Omega/E_L < 4$ the lower band has two minima, and if $\Omega/E_L \geq 4$ they merge into a single minimum. (b) The detuning acts to shift the minimum of the bands from $k_x = 0$, as shown here for $\Omega/E_L = 5$. In all cases, the regions near the band minima are approximately parabolic, as with the bare atomic dispersions.

opposite momentum.

Once the Raman coupling is enabled, $\Omega > 0$, the energies of the new eigenstates, *dressed states*, exhibit avoided-crossings due to the coupling, as seen in Fig. 3.1(a). The dressed states are superpositions of the bare atomic spins, where the relative population of each spin component depends on the values of Ω and δ . The dispersion thus consists of an upper and lower curve, the energies of the associated dressed states, with a gap between them. The spacing at $k_x = 0$ is Ω (with no detuning), representing the coupling strength between the dressed states. The dispersion is now reminiscent of band structure in solids. For Rabi frequencies $\Omega/E_L < 4$, the lower band has two minima, which move closer together with increasing Ω ; once $\Omega/E_L \geq 4$ they merge. We therefore define Rabi frequencies $\Omega/E_L > 4$ as the *strong coupling* regime, as the physics here are quite different from the case with two minima [30, 34]. We will mainly focus on the strong coupling regime for the remainder of this chapter.

The effect of the detuning δ is to shift the minima of the upper and lower bands away from $k_x = 0$, according to its sign. This may be understood as one spin state becoming more energetically favorable than the other. Since the spins have opposing momenta, there is a corresponding shift in the minima away from $k_x = 0$. We




Figure 3.2: Dressed state dispersions, with the state composition in terms of the spin components indicated by line coloring; $|\uparrow\rangle$ and $|\downarrow\rangle$ are the ± 1 eigenstates of $\hat{\sigma}_z$, respectively. Due to each spin state having equal and opposite momenta, an imbalanced superposition of spins leads to a corresponding momentum. In both (a) and (b), $\Omega/E_L = 5$. (a) $\delta/E_L = 0$, no spin component is energetically favorable over the other, so each dressed state has minimum energy when comprised of an equal superposition of both, and therefore have no net momentum. (b) $\delta/E_L = 2$, detuning imposes a spin imbalance in the composition of the minimum energy dressed states, with a corresponding shift to finite momenta; each band is shifted in opposite directions, as the spin imbalance in each dressed state is opposite (though not equal).

may see this by looking at the relative population between spin components along the bands, as shown in Fig. 3.2. For the uncoupled case, each band is comprised entirely of one spin or the other. With coupling, the new eigenstates, the dressed states, are mixtures of the spin components. Since the spin components carry opposing momenta, a dressed state with some finite momentum must be comprised predominantly of the corresponding spin. Conversely, a dressed state with some spin superposition must have a corresponding momentum, determined by the relative amplitudes of each spin. When $\delta = 0$ specifically, the dressed states are equal mixtures of the spins, for $k_x = 0$. These states have no net momentum, since they are comprised equally of the two spin components, with equal and opposite momenta. Each dressed state therefore has minimum energy when they are an equal superposition of the spins, in the absence of detuning. When detuning is added, the minima of each band are shifted to opposite sides in k -space, indicating that the dressed states now energetically prefer one spin to the other, and have corresponding momenta due to this imbalance.

We now focus our attention to the lower band specifically, which is approximately parabolic in shape. The energy therefore takes the form,

$$E \approx \frac{\hbar^2}{2m^*} \left(\hat{\mathbf{k}} - \frac{q\mathbf{A}}{\hbar c} \right)^2 + q\varphi. \quad (3.18)$$

This expression for a parabola in k -space is purposefully evocative, since it represents the energy of a particle with charge q , and mass m^* , subjected to vector potential \mathbf{A} and scalar potential φ ; c here is the speed of light. This system of Raman dressed neutral atoms is therefore akin to that of a charged particle moving in an electromagnetic field. Atoms obtain a new effective mass $m^* > m$, given by the curvature of the band. Since the momentum of atoms is only shifted along the x -axis, the vector potential here has a single component, $\mathbf{A} = A\mathbf{e}_x$. The magnitude A is determined by the location of the band minimum, $k_{\min.}$; it is more convenient to work in units of momentum here, so we will neglect c and the charge q , writing $A = \hbar k_{\min.}$. To be explicitly clear, these are Abelian gauge fields acting on the lower dressed state, which is a mixture of spins. The upper band follows a similar prescription, so it could also be described by a different set of Abelian gauge fields, but we will ignore the upper band here, assuming states may be prepared purely in the lower band. The important point is that, in either dressed state, spins are subjected the *same* gauge field; if they were to differ, then this would be akin to a non-Abelian gauge field, though there are additional criteria for this (see Ch. 4). In the weak-coupling regime, with $\Omega/E_L < 4$, spins are split into two separate minima, creating a form of synthetic spin-orbit coupling as from solid-state systems [29, 34]; in this case, each spin is subject to a different set of fields, but we will not discuss such a situation here any further.

From this representation, we should expect that from the vector and scalar potentials, effective electric and magnetic fields may emerge according to the usual relations from classical electromagnetism,

$$\mathbf{E} = -\nabla\varphi - \frac{\partial\mathbf{A}}{\partial t}, \quad (3.19)$$

$$\mathbf{B} = \nabla \times \mathbf{A}. \quad (3.20)$$

If all the parameters of the Hamiltonian (Eq. 3.15) remain fixed, the vector and scalar potentials are uniform in space-time, resulting in no effective electric or magnetic fields [28]. Alternatively, if by tuning the parameters of the Hamiltonian we are able to engineer either a time-dependent vector potential, or gradient of the scalar potential, then an electric field may be produced. Similarly, a vector potential with a non-zero curl will produce an effective magnetic field. Both of these situations have been realized in experiment [32, 33].

To see how such fields may be engineered, consider the effect of Ω and δ discussed above, and shown in Fig. 3.1: in the absence of detuning, Ω shifts the spacing between bands, which is akin to tuning the scalar potential. The detuning shifts

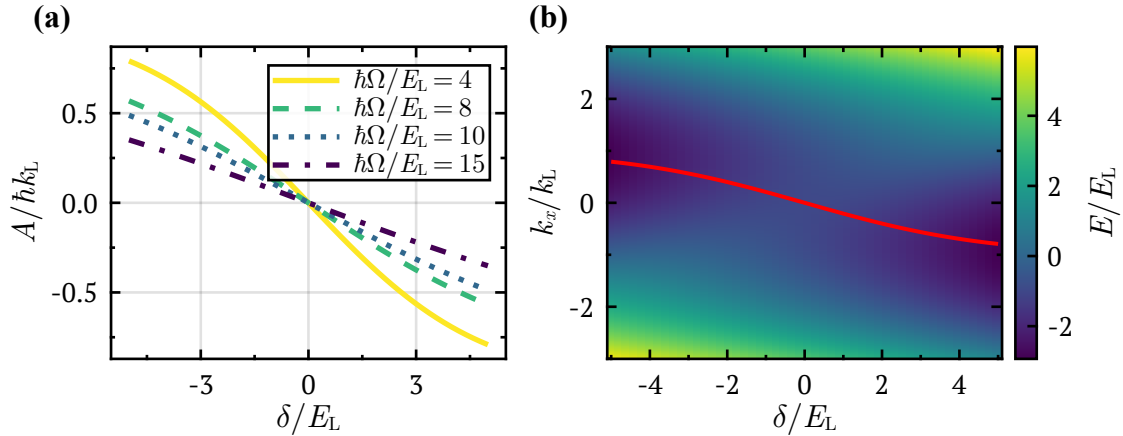


Figure 3.3: Properties of the lowest energy dressed state (band) of the Hamiltonian Eq. 3.15. (a) For various choices of Rabi frequency Ω in the strong-coupling regime, we plot the momentum corresponding to the band minimum for different detunings, which is the artificial vector potential $A = \hbar k_{\min.}$; results are given in units of the recoil momentum k_L . The resulting vector potentials are approximately linear near $\delta = 0$. (b) The energy surface of the lower band as it varies with detuning, with $\Omega/E_L = 4$, and the minimum indicated by the solid red line.

the minima of the bands along k_x , which resembles changes to the vector potential³. Therefore, producing an Ω or δ with a spatial dependence, or detuning with a time-dependence, effective electric and magnetic fields may be engineered to act upon neutral atoms.

As an example of this, in Fig. 3.3(a) we calculate $k_{\min.} = A/\hbar$ with varied detuning, and $\Omega/E_L \geq 4$; we find that in the vicinity of $\delta = 0$, the vector potential is approximately linear. Suppose then that a spatial gradient of the detuning is applied along the y -axis; we will assume its linear, with slope β , and centered at $y = 0$, *i.e.* $\delta(y) = \beta y$. The vector potential thus has the form, $\mathbf{A} = A\mathbf{e}_x$, with $A \propto \beta y$, and hence a magnetic field, $\mathbf{B} = -\partial_y A\mathbf{e}_z$ will arise. This will result in a Lorentz force acting on atoms in the ensemble, producing circulating currents.

The pragmatist might point out that there is no need to identify effective gauge potentials in the Hamiltonian, Eq. 3.15, in order to explain the physics at hand, and they would be correct; however, since this effective description is possible (at least in a particular gauge, since we do not have independent control of scalar and vector potentials [36, 37]), either explanation is valid. This is the core argument of quantum simulation. There are many extensions of these fairly simple ideas of tuneable electromagnetic potentials in ultracold atoms, as the seminal work of *Lin et al.* 2009 [28] has inspired a generation of physicist (such as this author) to implement various interesting alterations to these protocols [38–42]. To this end, we developed a flexible package in Python to simulate the effects of various such fields

³In both cases, changing Ω or δ , the effective mass m^* also changes slightly.

in BECs [1], with the hope of some day implementing them in the lab. We instead found that these artificial gauge fields, which are well understood through classical electromagnetism, are in fact an interesting way to probe the many-body physics of large interacting systems. The remainder of this chapter will focus on these efforts.

3.2 Bose-Einstein Condensates

As we alluded to in the previous section, Sec. 3.1, the rather simple Abelian gauge fields from classical electromagnetism may equip us with a means to study the many body physics underlying a system of interacting particles, *i.e.* a BEC. Some of the earliest experiments in BECs looked at the effects of rotation, which permitted the nucleation of *quantized vortices* in the gas [43]. Vortices are excitations resulting from the BEC having angular momentum (hence the rotation experiments), permitting circulating currents of atoms. Since atoms in a BECs are described together through a single macroscopic wavefunction, the motion of atoms in the gas is the result of a spatial gradient in the phase of the wavefunction. The phase profile in a BEC must remain continuous at all times for the wavefunction to be coherent. A circulating current is therefore caused by a phase winding; for it to remain continuous, it must therefore wind an integer multiple of 2π times. In such a case, within such a winding, there must be a discontinuity in the phase, and hence the BEC density drops to precisely zero at this location. This is the nature of a quantized vortex: an integer multiple of 2π phase winding, wherein the BEC density drops sharply to zero, where this location is called the *vortex core*. Such effects are manifestations of interactions in the gas, and are of immense interest. As we discussed in Sec. 3.1.2, artificial gauge fields may provide a means to study such phenomena.

In this section we will cover the mean-field theory of interacting Bose gases, only to the extent required to understand the numerical techniques used to compute the density profiles of such gases in Sec. 3.3. There are many excellent reviews of BECs [9, 44, 45], so here we aim to introduce only the basic concepts, and provide some intuition along the way.

3.2.1 Gross-Pitaevskii Equation

In the work that follows, we will mainly be concerned with finding the ground state wavefunction of BECs subjected to various external conditions, such as Raman-dressed artificial gauge fields. We will model the BEC through mean-field theory, wherein we approximate the macroscopic wavefunction of the gas as a complex scalar field [9]; this is essentially a classical field approximation of the quantized

bosonic field operators. Excitations of the gas outside of the ground state, called the condensate *depletion*, will be ignored.

To start, we will derive the energy functional for a system of identical bosons, which are interacting. Understanding the energy of the gas in this way will allow us to derive some useful expressions. We will assume that the interactions are low energy (*s*-wave), so the scattering is constant throughout momentum space, resulting in a Dirac-delta potential in coordinate space. For dilute gases, with interactions mostly over large scales, this approximation holds well [9]. The many-body Hamiltonian for N atoms is,

$$\hat{H} = \sum_{i=1}^N \left(-\frac{\hbar^2 \nabla_i^2}{2m} + V(\mathbf{r}_i) \right) + g \sum_{i<j} \delta(\mathbf{r}_i - \mathbf{r}_j) \quad (3.21)$$

where V is the external trapping potential, and the interaction strength in terms of the atomic scattering length a is [9],

$$g = \frac{4\pi\hbar^2 a}{m}. \quad (3.22)$$

We write the many-body wavefunction Ψ as a product state of the single-particle wavefunctions ϕ ,

$$\Psi(\mathbf{r}_1, \dots, \mathbf{r}_N) = \prod_{i=1}^N \phi(\mathbf{r}_i), \quad (3.23)$$

with $\int d\mathbf{r} |\phi(\mathbf{r})|^2 = 1$.

Computing the energy expectation value $E = \langle \Psi | \hat{H} | \Psi \rangle$, we have,

$$\begin{aligned} E &= \int d\mathbf{r}_1 \dots d\mathbf{r}_N \phi^*(\mathbf{r}_1) \dots \phi^*(\mathbf{r}_N) \left\{ \sum_{i=1}^N \left[-\frac{\hbar^2 \nabla_i^2}{2m} + V(\mathbf{r}_i) \right] \right. \\ &\quad \left. + g \sum_{i<j} \delta(\mathbf{r}_i - \mathbf{r}_j) \right\} \phi(\mathbf{r}_1) \dots \phi(\mathbf{r}_N), \\ &= \sum_{i=1}^N \int d\mathbf{r}_i \left[-\frac{\hbar^2}{2m} \phi^*(\mathbf{r}_i) \nabla_i^2 \phi(\mathbf{r}_i) + V(\mathbf{r}_i) |\phi(\mathbf{r}_i)|^2 \right] + g \sum_{i<j} \int d\mathbf{r}_i |\phi(\mathbf{r}_i)|^4, \end{aligned} \quad (3.24)$$

where to obtain the second line, we were able to carry out the integration of all wavefunctions not acted upon by each of the energy operators, all yielding unity due to normalization. For the first term, involving the kinetic and potential energies, we notice that each term is identical, and there are N such terms. The same is true of the interaction term, except that there are $N(N-1)/2$ such terms. The energy

becomes,

$$E = N \int d\mathbf{r} \left[-\frac{\hbar^2}{2m} \phi^*(\mathbf{r}) \nabla^2 \phi(\mathbf{r}) + V(\mathbf{r}) |\phi(\mathbf{r})|^2 + \frac{N-1}{2} g |\phi(\mathbf{r})|^4 \right]. \quad (3.25)$$

We write the mean-field wavefunction as $\psi(\mathbf{r}) = \sqrt{N} \phi(\mathbf{r})$, such that the density distribution of all particles is $n(\mathbf{r}) = |\psi(\mathbf{r})|^2$. The field $\psi(\mathbf{r})$ is often called the condensate wavefunction, or **BEC** order parameter. It is normalized by the total particle number as,

$$\int d\mathbf{r} |\psi(\mathbf{r})|^2 = \int d\mathbf{r} n(\mathbf{r}) = N. \quad (3.26)$$

It is this wavefunction ψ that we eventually aim to solve for, by minimizing the energy functional.

In the thermodynamic limit, where $N \gg 1$, the energy functional becomes,

$$E(\psi) \approx \int d\mathbf{r} \left[-\frac{\hbar^2}{2m} \psi^*(\mathbf{r}) \nabla^2 \psi(\mathbf{r}) + V(\mathbf{r}) |\psi(\mathbf{r})|^2 + \frac{g}{2} |\psi(\mathbf{r})|^4 \right], \quad (3.27)$$

representing the energy for a condensate with order parameter $\psi(\mathbf{r})$. To determine the lowest energy state, we may minimize it under the constraint that there are N particles. Specifically, we compute,

$$\left(\frac{\partial E}{\partial \psi^*} - \mu \frac{\partial N}{\partial \psi^*} \right) + \left(\frac{\partial E}{\partial \psi} - \mu \frac{\partial N}{\partial \psi} \right) = 0, \quad (3.28)$$

with the Lagrange multiplier μ associated with particle number conservation being the chemical potential. Both terms yield the same result, so here we will look only at the derivatives with respect to ψ^* . We have,

$$\frac{\partial E}{\partial \psi^*} - \mu \frac{\partial N}{\partial \psi^*} = 0 = \int d\mathbf{r} \left[-\frac{\hbar^2}{2m} \nabla^2 \psi(\mathbf{r}) + V(\mathbf{r}) \psi(\mathbf{r}) + g |\psi(\mathbf{r})|^2 \psi(\mathbf{r}) - \mu \psi(\mathbf{r}) \right]. \quad (3.29)$$

Setting the integrand to zero and rearranging, we obtain,

$$\mu \psi(\mathbf{r}) = \left[-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + g |\psi(\mathbf{r})|^2 \right] \psi(\mathbf{r}). \quad (3.30)$$

This is the time-independent **GPE**, which bears striking similarity with the time-independent Schrödinger equation, aside from the nonlinear interaction term (and the energy eigenvalue is replaced by the chemical potential). There is a time-dependent version of the **GPE** as well, which is,

$$i\hbar \frac{\partial}{\partial t} \psi(\mathbf{r}, t) = \left[-\frac{\hbar^2}{2m} \nabla^2 + V(\mathbf{r}) + g |\psi(\mathbf{r})|^2 \right] \psi(\mathbf{r}). \quad (3.31)$$

The former, Eq. 3.30, may be derived from the latter, Eq. 3.31, by assuming the condensate wavefunctions acquire phase according to their energy, $\psi(\mathbf{r}, t) = \psi(\mathbf{r}) \exp(-i\mu t/\hbar)$ [9]. As such, we may utilize either version just as we do the time-independent Schrödinger equation and [time-dependent Schrödinger equation \(TDSE\)](#): the first gives the eigenstates of the system in question, while the second is used to propagate those eigenstates in time, where each state acquires phase according to its energy. Here, we may use Eq. 3.30 to solve for the ground state wavefunctions of the [BEC](#), and then use Eq. 3.31 to propagate the resulting complex wave packet in time. Together, both versions of the [GPE](#) give us immense computational power, though this requires us to solve a nonlinear differential equation, potentially in both space and time; a method for accomplishing this is discussed in Sec. 3.3.

3.2.2 Thomas-Fermi Approximation

In the absence of a nonlinear differential equation solver, or known analytical solution to Eq. 3.30, there are still ways in which we may predict various properties of [BECs](#); in particular, through the [TF](#) approximation. This consists of assuming that the kinetic energy term in Eq. 3.30 is negligible. We might argue that the atoms are cold, and therefore have low kinetic energy, but to be more precise, the kinetic energy is only large where the density gradients are high, which is near the edges of the gas (for trapped atoms). In the regions where the density $n(\mathbf{r})$ is high, we therefore expect the interactions to have a more substantial role in the energetics. Under this approximation, the [GPE](#) becomes,

$$\mu = V(\mathbf{r}) + g |\psi(\mathbf{r})|^2, \quad (3.32)$$

and therefore,

$$n(\mathbf{r}) = |\psi(\mathbf{r})|^2 = \frac{\mu - V(\mathbf{r})}{g}. \quad (3.33)$$

Typically, [BECs](#) are confined in [ODTs](#) with an approximately harmonic potential,

$$V(\mathbf{r}) = \frac{1}{2}m (\omega_x^2 x^2 + \omega_y^2 y^2 + \omega_z^2 z^2), \quad (3.34)$$

where ω_q for $q \in \{x, y, z\}$ are the *trap frequencies*. These are the rates at which atoms oscillate in the [ODT](#), along the indicated directions⁴. The density profile of the gas under the [TF](#) approximation is therefore an inverted parabola, owing its shape to the trapping potential. Where the density reaches zero along any given

⁴In fact, observing such oscillations in the lab is precisely how the trap frequencies are measured.

direction q , we define the **TF** radii,

$$R_q = \frac{1}{\omega_q} \sqrt{\frac{2\mu}{m}}. \quad (3.35)$$

Beyond R_q the density is assumed to be zero.

The density profile thus has the form,

$$n(\mathbf{r}) = \frac{\mu}{g} \left(1 - \frac{x^2}{R_x^2} - \frac{y^2}{R_y^2} - \frac{z^2}{R_z^2} \right). \quad (3.36)$$

From its normalization condition, $N = \int d\mathbf{r} n(\mathbf{r})$, we may also derive an expression for the chemical potential. The integral may be performed in generalized spherical coordinates, with $d\mathbf{r} = R_x R_y R_z \rho^2 \sin\theta d\rho d\theta d\phi$, radius $\rho \in [0, 1]$, polar angle $\theta \in [0, \pi)$, and azimuthal angle $\phi \in [0, 2\pi)$. The Cartesian coordinates assume the typical form,

$$x = R_x \rho \sin\theta \cos\phi, \quad (3.37)$$

$$y = R_y \rho \sin\theta \sin\phi, \quad (3.38)$$

$$z = R_z \rho \cos\theta, \quad (3.39)$$

such that the integrand in Eq. 3.36 is just $1 - \rho^2$. Integrating, and substituting in our expressions for the **TF** radii (Eq. 3.35) we obtain,

$$N = \frac{8\pi}{15} \frac{\mu}{g} \left(\frac{2\mu}{m\bar{\omega}^2} \right)^{3/2}, \quad (3.40)$$

where we have defined the geometric average of the trap frequencies as $\bar{\omega} = (\omega_x \omega_y \omega_z)^{1/3}$. Rearranging for the chemical potential,

$$\mu = \frac{\hbar\bar{\omega}}{2} \left(\frac{15Na}{\bar{a}} \right)^{2/5}, \quad (3.41)$$

where we have substituted in for g (Eq. 3.22), and $\bar{a} = \sqrt{\hbar/m\bar{\omega}}$ is called the harmonic oscillator length. This is a convenient form for μ , in terms of experimentally available parameters, namely the trap frequencies along each axis ω_q , the atom number N , and the scattering length a that has been measured in various cold atomic species; for ^{87}Rb , the scattering interactions are repulsive, $a > 0$. The chemical potential sets the relevant energy scale in the **BEC**. In our experiments, μ is typically on the order of 1 kHz, mainly depending on the atom number, which varies between about 10^4 to 10^5 .

To demonstrate the effectiveness of the **TF** approximation, Fig. 3.4 shows a com-

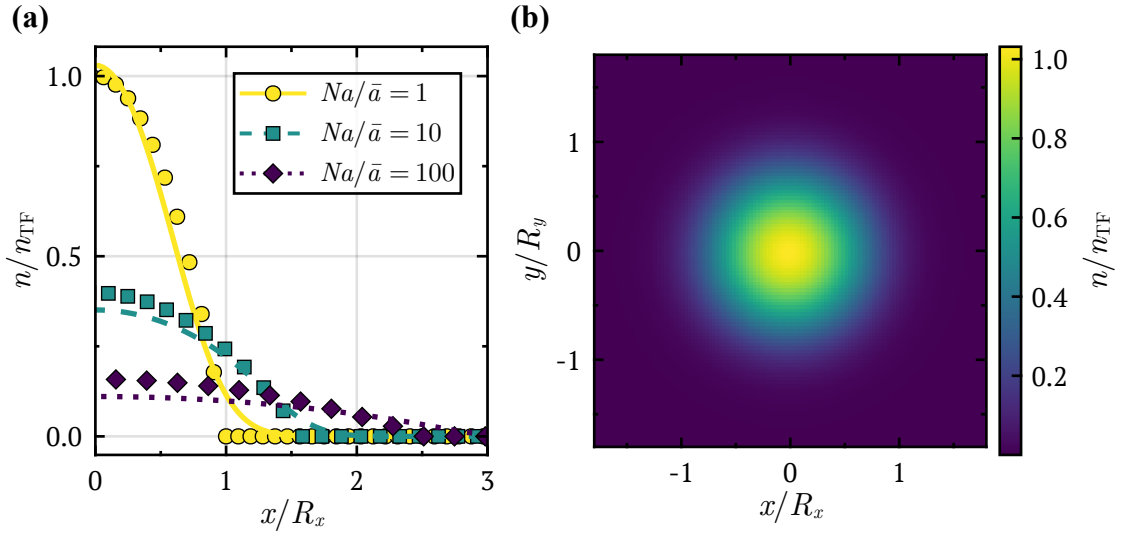


Figure 3.4: (a) Slices of BEC densities along x -axis, from the center of the BEC outwards, comparing results from the TF approximation (points) to the numerical simulations (lines) described in Sec. 3.3. Results show good agreement between density profiles, deviating most at the condensate edge where the kinetic energy is most significant. Curves are shown for BECs with three different interaction volumes; for larger values, interactions play a bigger role in the energetics, as seen by the repulsive interactions forcing atoms outwards to reduce the peak density. The x -axis is normalized by the TF radius of the $Na/\bar{a} = 1$ curve, and the densities by its peak density. A similar plot may be found in Ref. [9]. (b) Full density distribution in 2D, for condensate with $Na/\bar{a} = 1$, solved for numerically.

parison between the TF density (Eq. 3.36) and results from numerical integration of the GPE (Eq. 3.30). The density profile in the x -direction is shown in this case, from the center of the BEC outwards, but the results generalize to the other spatial dimensions. We can see that the TF density follows the numerical solutions well, deviating the most at the edge of the condensates where the kinetic energy (spatial gradient) is the largest. Effectively, adding the kinetic energy into the picture acts to “round off” the edges. An example density distribution found numerically (see Sec. 3.3) in 2D is shown in Fig. 3.4(b), displaying the full density profile for a uniform trap.

From Fig. 3.4(a), the density profiles are shown for three different *interaction volumes*, Na/\bar{a} , which determines the extent to which interactions play a role in the energetics, as compared to the kinetic energy and confinement [9]. For N atoms interacting over lengths a , the relative size of these interactions depends on how many are packed within the scale of their confinement, given by \bar{a} . The larger the interaction volume, the more significant interactions are; this may be seen in Fig. 3.4(a) where for higher values, the repulsive interactions in the gas minimize the energy by reducing the peak density, resulting in a larger condensate.

3.2.3 Spinors

The BECs discussed so far have no spin degree of freedom, or rather the spin has been ignored. To implement the artificial gauge fields described in Sec. 3.1.1, this must be remedied. Thankfully, generalizing the results to *spinor* gases (that is, those with a spin degree of freedom) is relatively easy [46].

Consider an ensemble with two spin components; we may describe the system through two wavefunctions, ψ_\uparrow and ψ_\downarrow , one for each spin. Each wavefunction independently follows the GPE shown before (Eq. 3.30): their kinetic and potential energies have the same form, in addition to the interactions between atoms in each component. Atoms from different components, however, are also permitted to interact. This results in two coupled GPEs [47, 48],

$$\mu\psi_\uparrow(\mathbf{r}) = \left[-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) + g_\uparrow|\psi_\uparrow(\mathbf{r})|^2 + g_\downarrow|\psi_\downarrow(\mathbf{r})|^2 \right] \psi_\uparrow(\mathbf{r}), \quad (3.42)$$

$$\mu\psi_\downarrow(\mathbf{r}) = \left[-\frac{\hbar^2}{2m}\nabla^2 + V(\mathbf{r}) + g_\downarrow|\psi_\downarrow(\mathbf{r})|^2 + g_\uparrow|\psi_\uparrow(\mathbf{r})|^2 \right] \psi_\downarrow(\mathbf{r}), \quad (3.43)$$

where g_\uparrow and g_\downarrow are the intra-atomic interaction strengths for each spin, and g_\uparrow is the inter-atomic interaction strength. Interactions between atoms in the same spin scale with their density, while the interactions between spins scale with the density of the other. For larger spin manifolds we may continue adding interactions in this manner, resulting in many coupled equations that will become expensive to solve, especially in 3+1 dimensions. In general though, it is possible. With the aim of simulating artificial electromagnetic potentials in BECs, we will be looking at solving the spinor GPEs, Eq. 3.42 and Eq. 3.43, numerically.

3.3 Numerical Methods

The role of interactions in ultracold gases are substantial, and yet even in the context of mean-field theory where the quantization of matter fields is ignored, solving the resulting GPEs to determine properties of the system remains difficult. The GPE contains a cubic nonlinearity in the wavefunction ψ , due to the density dependence of the mean-field interactions. Solutions to such a differential equation are not easily accessible, especially in cases where the trapping potentials or other external fields lack symmetries to aid in determining solutions. Hence, numerical techniques provide significant value to researchers, though highly performant code is not ubiquitous in the community, especially among experimentalists.

In this section, we will cover one of the common numerical techniques for solving both time-dependent and independent GPEs (Eq. 3.31 and Eq. 3.30). The technique

will then be adapted to solving spinor **GPEs** (Eqs. 3.42- 3.43), with functionality to add Raman-dressing as in Sec. 3.1.1, with the purpose of exploring many-body ground states of trapped atoms subjected to such artificial gauge-fields. Last, the data-parallel nature of the algorithm permits it to be accelerated on **GPUs**, resulting in substantially faster run-times than on **central processing units (CPUs)**; this allows researchers to spend more time in the lab performing experiments, rather than running and debugging **BEC** simulations [1].

3.3.1 Time Splitting Spectral Method

A common numerical technique for solving **GPEs**, and closely related nonlinear Schrödinger equations, is the **time-splitting spectral method (TSSM)** [49]. It is widely known that finite-difference derivatives are computationally expensive; the kinetic energy, which is a Laplacian operator in coordinate space, therefore imposes a numerical challenge. The **TSSM** takes advantage of the fact that, in momentum space, the kinetic energy is a diagonal operator, with the Laplacian $\nabla^2 \rightarrow \mathbf{k}^2$. Wavefunctions in coordinate-space ψ are related to those in momentum space $\tilde{\psi}$ by a Fourier transform. Computationally, performing **fast Fourier transforms (FFTs)** is significantly less expensive than numerical gradients. This is the key observation that the **TSSM** takes advantage of.

We will describe the **TSSM** as applied to the time-dependent **GPE**, Eq. 3.31; in this case, we will discuss it as applied to simulations in quasi-2D, where in the third dimension an analytic density profile is considered (such as a Gaussian). This allows the densities predicted in simulation to correspond well with truly 3D **BECs**, at low computational cost, but works best when the confinement along one axis is stronger than the others⁵. The simulations presented in this chapter were performed in such a manner [1], though the algorithm extends to 3D, at significant memory cost.

To start, the wavefunctions must be discretized in both coordinate space and momentum space, hereafter called r -space and k -space, respectively. We designate the two spatial dimensions as x and y , with z the out of plane dimension, over which the density may be assumed as Gaussian, or some other density profile such as that of the **TF** approximation. Grids of dimension $N_x \times N_y$ are initialized; in r -space, they extend over the range $(\pm x^{\max}, \pm y^{\max})$, with the corresponding grid spacing of $\Delta q = 2q^{\max}/N_q$, $q \in \{x, y\}$. From this, the grids in k -space are initialized with the same dimensions. The limits in k -space are determined by the spatial resolution, $k_q^{\max} = \pm\pi/\Delta q$, and the spacings by the grid size $\Delta k_q = \pi/q^{\max}$. The wavefunctions, either in r - or k -space, are therefore $N_x \times N_y$ arrays of complex numbers (double

⁵This is particularly applicable to crossed **ODTs**, such as that in our system, since the trap frequencies in the plane of the beams are lower than the out of plane trap frequency.

precision, here), with each element corresponding to the amplitude at a particular location in that space. We will discuss practical considerations when choosing these spacings, and other parameters, further below.

To propagate a wavefunction forward in time, we discretize time into small steps of duration Δt . We may write the [GPE](#) in the following form,

$$-i\hbar\frac{\partial}{\partial t}\psi = \hat{H}_{\text{eff.}}\psi, \quad (3.44)$$

in terms of an effective Hamiltonian,

$$\hat{H}_{\text{eff.}} = \frac{\hbar^2}{2m}\nabla_{xy}^2 + V(x, y) + g|\psi(x, y)|^2, \quad (3.45)$$

which is itself a function of ψ due to the interactions⁶. The wavefunction at each step is therefore computed by the application of the evolution operator,

$$\hat{U}(\Delta t) = \exp\left(-\frac{i}{\hbar}\hat{H}_{\text{eff.}}\Delta t\right) \quad (3.46)$$

with $\hat{H}_{\text{eff.}}$ computed from the previous step's wavefunction. The evolution operator at each time step is therefore a spatial grid, and computing $\hat{U}\psi$ is a *Hadamard product* (element-wise multiplication of arrays), where the wavefunction amplitude at each location on the grid obtains phase according to the corresponding element in \hat{U} .

In principle, $\hat{H}_{\text{eff.}}$ could be computed in full at each time step, and then applied as above, but as stated before this would require the kinetic energy term to be calculated from gradients (finite-difference or otherwise). The [TSSM](#) avoids this by further decomposing each time step. Specifically, the evolution operator for the kinetic energy is computed in k -space, and the potential energy and interaction terms in r -space. The corresponding evolution operators are therefore calculated and applied to the wavefunction separately. The evolution operator for the potential and interactions, $\hat{U}_{V,g}$ is applied to the r -space wavefunction ψ . Prior to applying the kinetic energy, \hat{U}_k , we transform the wavefunction to k -space by a 2D [FFT](#), $\psi \rightarrow \tilde{\psi}$; the kinetic energy term is applied to $\tilde{\psi}$, and then the state is transformed back to r -space by an inverse-[FFT](#) for the next time-step. This cycle is iterated over for the desired duration. Importantly, the kinetic energy does not commute with either the potential or interaction energies, so to reduce the commutation error generated by assuming $\hat{U}(\Delta t) = \hat{U}_k(\Delta t)\hat{U}_{V,g}(\Delta t)$ we adopt the common *Strang* Trotterization

⁶Note that we have expressed the Laplacian operator ∇^2 in 2D here.

(also called Strang splitting) [50]. This takes the form,

$$\hat{\mathcal{U}}(\Delta t) = \hat{\mathcal{U}}_k \left(\frac{\Delta t}{2} \right) \hat{\mathcal{U}}_{V,g}(\Delta t) \hat{\mathcal{U}}_k \left(\frac{\Delta t}{2} \right), \quad (3.47)$$

where the kinetic energy is applied for a half time step $\Delta t/2$, followed by potential and interaction term for a full time step Δt , and finished by another half from kinetic energy. Higher order splittings also exist [51]. Every time step therefore contains four 2D FFTs, and three Hadamard products.

Here we have described the process for taking a known wavefunction ψ and propagating it forward in time according to the time-dependent GPE, Eq. 3.31. Of course, as we mentioned in Sec. 3.2.1, solutions to the time-independent GPE are not generally accessible. Prior to propagating wavefunctions in time as described, we must first determine a suitable ground state to start with. Thankfully, the algorithm described above may also be applied to finding BEC ground states, which are solutions to Eq. 3.30. This is done by starting with some initial wavefunction that is approximately the BEC ground state, such as the TF density, and propagating it forward in *imaginary* time.

To understand this, consider the evolution of states according to the TDSE; decomposing the wavefunction into the eigenbasis of the Hamiltonian, we may write the time evolution as,

$$|\psi(t)\rangle = \hat{\mathcal{U}}(t) |\psi_0\rangle = \sum_j \tilde{c}_j e^{-i\epsilon_j t/\hbar} |\psi_j\rangle \quad (3.48)$$

where the eigenstate $|\psi_j\rangle$ obtains phase according to its energy eigenvalue ϵ_j , and $\tilde{c}_j = \langle \psi_0 | \psi_j \rangle$. If we make time imaginary, by substituting $t \rightarrow -i\tau$, then we find that rather than obtaining phase as τ increases, wavefunctions are exponentially suppressed according to their energy. States with higher energy are therefore more suppressed, so the dominant term in the expansion after some “time” τ will be that with the lowest energy, *i.e.* the ground state. Imaginary time propagation is akin to taking the zero-temperature limit [52].

Imaginary time propagation through the TSSM is done this very way, by replacing the discrete time step with a complex one, $\Delta t \rightarrow -i\Delta\tau$, so that the evolution operators are no longer complex exponentials. We are free to renormalize the wavefunctions after each time step, so that after many such steps forward in τ , the wavefunction will converge to the ground state. In the simplest case, starting with the TF density, one would witness the transition from the sharp inverted parabola, to a smooth distribution such as those shown in Fig. 3.4. In this way, we may obtain reasonable ground states to then propagate in real-time through the same algorithm.

With the algorithm described above, we are able to discretize wavefunctions in space, solve for the ground states, and propagate them forward in similarly discretized time. The grid parameters should be chosen so that the boundaries of the grid extend far beyond the size of the BEC, for all times in the simulation, otherwise edge effects will occur. A natural choice of length scale is therefore the TF radius, where the grid size should exceed this. Increasing the grid boundary far beyond the region occupied by atoms, however, must also be accompanied by an increase in the grid dimensions, $N_x \times N_y$, in order to maintain the required resolution (which sets the k -space dimension). The operations performed by the CPU during each step mostly consist of Hadamard products between arrays and 2D FFTs, with arrays comprised of double precision complex floating point numbers⁷. The time steps, Δt , should also be sufficiently small to reduce integration errors; the chemical potential μ serves as a good energy scale for this, where time steps $\Delta t \ll \hbar/\mu$. Reducing the time step too much may reduce the numerical error, but at the cost of significantly longer run-times. Therefore, with better spatial and temporal resolution, simulations quickly become computationally expensive.

Raman-Coupled Spinors

The TSSM described above may also be generalized to spinor gases, which are described by coupled GPEs, such as Eqs. 3.42-3.43 for the case of spin-1/2 systems. The algorithm is similar to before, except that we simultaneously propagate two components of the full wavefunction, $\Psi = (\psi_\uparrow, \psi_\downarrow)^\top$. Here we are interested in going beyond the simple case of just kinetic, potential, and interaction energies; we wish to simulate the effects of Raman-dressed artificial gauge fields, as described in Sec. 3.1.1. Specifically, we aim to realize the Hamiltonian Eq. 3.15, with a spatially dependent detuning $\delta = \delta(\mathbf{r})$ so as to produce interesting artificial gauge fields (see Sec. 3.1.2). Such a system necessitates the addition of a spin degree of freedom; in the case of spin-1/2, we may write the coupled GPEs as [1, 31],

$$\mu\psi_\uparrow = \left[-\frac{\hbar^2}{2m} (\mathbf{k}^2 + k_L k_x) + V(\mathbf{r}) + \frac{\delta(\mathbf{r})}{2} + g_\uparrow |\psi_\uparrow|^2 + g_\downarrow |\psi_\downarrow|^2 \right] \psi_\uparrow + \frac{\Omega}{2} \psi_\downarrow, \quad (3.49)$$

$$\mu\psi_\downarrow = \left[-\frac{\hbar^2}{2m} (\mathbf{k}^2 - k_L k_x) + V(\mathbf{r}) - \frac{\delta(\mathbf{r})}{2} + g_\downarrow |\psi_\downarrow|^2 + g_\uparrow |\psi_\uparrow|^2 \right] \psi_\downarrow + \frac{\Omega}{2} \psi_\uparrow, \quad (3.50)$$

where several of the terms differ in sign, according to the Pauli matrices in Eq. 3.15. Additionally, we have expanded the kinetic energy term, ignoring the $k_L^2 \hat{\sigma}_z^2$ term since it is uniform across all degrees of freedom ($\hat{\sigma}_z^2 = \hat{\mathbb{1}}$). While the inter-spin interaction terms $\propto g_\uparrow$ assign an energetic cost to atoms in different components

⁷FFT's are more efficient on arrays whose dimensions are powers of two, so the dimensions should be set accordingly.

being near one another, this term does not couple spin components (*i.e.* these are not spin-exchange collisions). The additional terms $\propto \Omega$, on the other hand do produce coupling.

Looking at the various terms, we can determine in which order and coordinate representation they should be applied. The term $\hbar^2 (\mathbf{k}^2 \pm k_L k_x) / 2m$ is the only one which is diagonal in k -space, and hence it is applied there while all the others are applied to the r -space wavefunctions. We determined empirically that, following the kinetic terms, the interaction and coupling terms should be applied with the same Strang splitting as the kinetic energy; this is because, though they commute spatially, they do not commute over the spin degree of freedom [50]. They are therefore applied as a half step each. Given the spatial dependence of detuning here, it acts as a spin-dependent potential, and may therefore be applied along with the trap potential. The remaining half time step of the interaction and coupling terms is then applied, followed by the remaining kinetic terms.

Overall, adding the spin degree of freedom, along with the new Raman dressing terms, increases the complexity of each simulation substantially. Each individual step now consists of eight 2D FFTs, and $\gtrsim 20$ Hadamard products. The memory usage is also substantial, since we must now track two wavefunctions each in both r -space and k -space, in addition to the grids pertaining to each term in the GPE. Depending on the complexity of the coupling fields desired, and the desired spatial resolutions, finding many-body ground states takes on the order of several hours with commercially accessible hardware. Hence, the desire for more performant code is great, so as to drastically speed up the time it takes to devise new experiments, or explain existing ones.

3.3.2 GPU Acceleration

Throughout the development of a Python package to implement the algorithm described in Sec. 3.3.1, in both single-component BECs and spinors, we noticed that both of the key operations performed throughout the iteration loops, Hadamard products and FFTs, may be run on GPUs with high efficiency. This is because these operations are both *data* parallel. For an operation to be data parallel, each element in a collection may be operated on simultaneously by the same function; the Hadamard product is an example of this, since the elements in each of two arrays may be multiplied independently of the others. The 2D FFT is also data parallel, since it may be decomposed into many 1D FFTs. While the iteration loop itself is inherently serial, since the result of each iteration is the input to the next, the operations throughout each step are data parallel. The performance may therefore be readily improved by modern hardware acceleration techniques.

Over the years, GPUs have become increasingly efficient computational tools, capable of much more than the graphics processing they were designed to do [53–55]; the various stages of the graphics computation pipeline have been reprogrammed to also do linear algebra [56–58], FFTs [59], and more [60, 61]. This is in large part because of the rise of machine learning, which has driven innovation in GPU production. In contrast to CPUs, which perform tasks sequentially but with low latency, GPUs perform many tasks simultaneously, but with relatively high latency. There is at some point a trade-off in the time saved by doing many things at once, slowly, versus one thing at a time, quickly; this is the contest in which CPUs and GPU engage⁸. There is of course no definite victor, as some tasks are inherently serial in nature, where CPUs still excel. Identifying the strengths and weaknesses of various devices, to test the situations in which they thrive over the others, is therefore important.

It is now possible to program essentially all of the operations performed in the graphics pipeline, and GPUs are capable of performing operations on increasingly complex types. As such, even through commercially available GPU hardware, significant performance gains are readily available. Furthermore, due also to the rise of machine-learning, there are numerous open-source packages that make hardware acceleration relatively straightforward. With these considerations in mind, we adapted the TSSM to be run on either CPUs or GPUs, with minimal changes⁹.

Initially, our Python package worked through the NumPy library [62], so in order to accelerate simulations, the machine-learning library PyTorch [63] was chosen due to its syntactic similarities. It provides a similar interface as NumPy, offering wrappers for various CUDA libraries. Much like NumPy, functions in PyTorch are array based, with the significant difference that in the latter, arrays can be initialized on a specified device; for instance, if two arrays are initialized on the same device (such as a GPU), then operations between said arrays are performed by that device. In this way, execution of code on different devices is easily assigned by passing a single keyword argument to various array allocations, specifying the device on which they are stored.

The only drawback with PyTorch at the time the code was originally developed was that it did not support complex numbers. As such, the most major modification to the original code was to define our own data structure to encode the complex phases of the wavefunctions. We settled on stacking two arrays of double precision floating point number on top of each other, each array representing either the real or

⁸Keep in mind that other devices, such as field-programmable gate arrays, or application-specific integrated circuits, are also used in high-performance computing. Some day, perhaps quantum computers will be to.

⁹Specifically NVIDIA GPUs, through their CUDA API.

imaginary component of the wavefunction. We therefore had to define custom functions for various array operations, to handle the complex amplitudes correctly. Since then, PyTorch now supports complex data types, making hardware acceleration of similar code all the more simple.

3.3.3 Hardware Considerations

Unlike what is usually the case for CPUs, before blindly executing code on GPUs there are some important things to consider, since they are quite different in design (this is becoming less and less true with future innovations, it seems). Simply put, not all GPUs are capable of performing operations on all data types efficiently; there are differences in device *architectures*, which dictate the operations available to them. For the NVIDIA GPUs employed here, this is captured by the *compute capability* (CC), which is a number assigned to each device by NVIDIA to describe computational resources. For instance, devices with $CC > 6$ have the hardware necessary to perform 64-bit addition operations, while other devices cannot [1]. The CC is therefore an important metric to look at when choosing devices, as regardless of any other specifics, it determines the ability of the device to perform the various operations required. The device architecture, or CC, will become important when comparing performance between devices.

In addition to architecture, the other specifications that drive device performance are memory, and the number of virtual cores. Older GPUs are not equipped with as much memory as CPUs, even of the same era, though with each new generation this limitation is reduced. The same is true of the number of virtual cores, which roughly determines how many concurrent operations the GPU is able to perform. On a more speculative note, it will be interesting to see how the various “ARM64” chips emerging in recent years will be employed, as they provide unified memory between the CPU and GPU. Regardless, we are now roughly equipped with the means to understand the capabilities of GPUs, in order to compare their performance with each other, and with CPUs.

3.3.4 Results

To test the performance of our GPU-accelerated Python simulations of spinor BECs, we ran benchmarks on various available devices, summarized in Tab. 3.1. In each case, we ran simulations repeatedly, for various different grid sizes; as the grid size is increased, computations become more expensive. Results for each device in Tab. 3.1 are shown in Fig. 3.5(a). The time required to complete a pre-determined simulation is measured. We find that, as the array size is increased, the CPU run-times also increase; this is to be expected, since they perform all operations in a

GPUs	GeForce MX150	GeForce 980 Ti	TITAN V
CC	6.1	5.2	7.0
CUDA Cores	384	2816	5120
VRAM [GB]	2	6	12
CPUs	Intel-i5-7200U	AMD FX-6300	Intel i9-9900K
Clock (Boost) [GHz]	2.5 (3.1)	3.5 (4.1)	3.7 (5.0)
Available RAM [GB]	8	16	32

Table 3.1: List of relevant specifications for the devices used to benchmark spinor GPE simulations. The devices that share a column were installed on the same workstation. For the GPUs, the CC, number of virtual cores (CUDA cores), and available memory (VRAM) are provided. Similarly, the clock speeds of the CPUs are given with and without boost, along with the memory (RAM) available on the workstation in which they are installed. Note that the number of CPU cores is not relevant here, since the algorithm may only run on a single core at a time. A similar table may be found in Ref. [1]

serial manner. With increasing grid size, there are more operations to do so the process takes longer¹⁰.

Conversely, as shown in Fig. 3.5(a), the GPUs are roughly constant with increasing grid size, up to a certain point that differs for each. Each GPU is capable of performing parallel operations on arrays up to some size, limited by the number of virtual cores they have. Up to this array size, there is no additional cost to increasing the grid size, so the computation times are constant. Once arrays pass beyond this limit, then the device must start batching operations, increasing the computation time. Note that the GeForce MX150 is a newer architecture than the older GeForce 980 Ti, as shown by its CC in Tab. 3.1, so for small grid sizes it outperforms the 980 Ti; however, due to its lower number of virtual core, the 980 Ti eventually surpasses its capability. This demonstrates the importance of architecture in such benchmarks. Last, we are not capable of running simulations of all grid sizes on these devices, as they are limited by memory; the CPUs generally have more memory access, so this was less of an issue for them.

In Fig. 3.5(b), we show the *speedup* for each CPU-GPU pair, which is the ratio of the GPU and CPU run-times, $\tau_{\text{CPU}}/\tau_{\text{GPU}}$. We find that as the array size increases, the speedup also increases, until the GPUs begin batching operations, producing diminishing returns. The highest speedup recorded here is 85, though practically speaking the number is far higher; this is because we determined it would only be fair to compare devices running on the same workstation, due to other factors that could be constraining performance. As such, the best GPU, the TITAN V, is compared against the best CPU, the Intel i9-9900K; the latter boasts an impressive

¹⁰Note that the relative performance between the CPUs is not necessarily determined by their clock speeds, as the device which performed to poorest has the second highest clock speed. This demonstrates the nuance involved when performing such benchmarks, which is beyond the scope of this qualitative discussion.

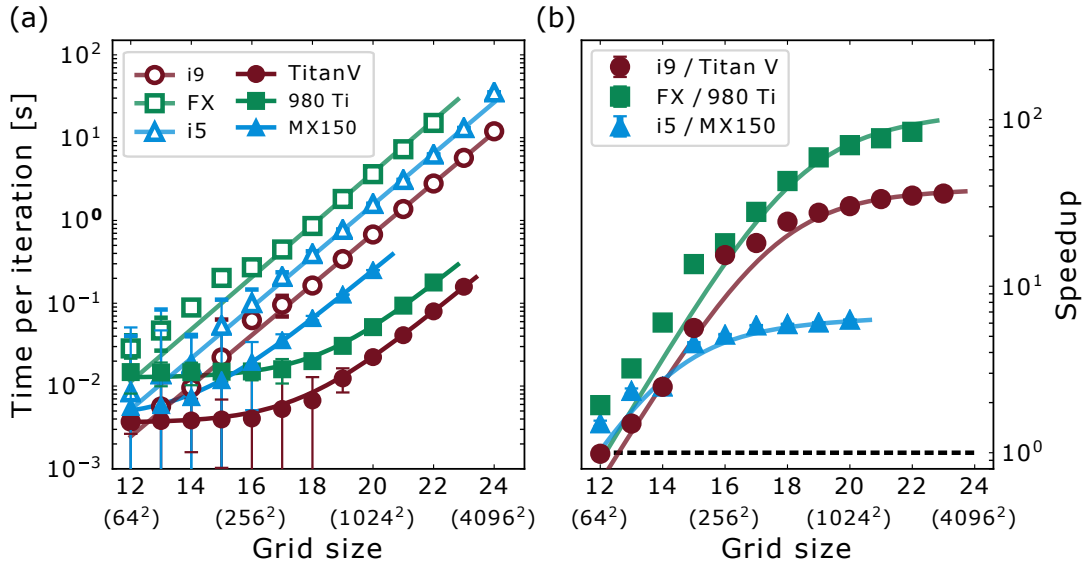


Figure 3.5: (a) Execution times of the various devices in Tab. 3.1 on spinor GPE simulations, with increasing grid sizes. Due to the FFT algorithms being more efficient on grid sizes of base two, the resolution was doubled in each step; the plot therefore has log-10/log-2 scaling. (b) Speedup of the various device pairs that are installed on the same workstation, as computed from the curves in (a) as the ratio of CPU to GPU execution times, $\tau_{\text{CPU}}/\tau_{\text{GPU}}$. We find that for all physically relevant grid sizes (> 64 in each axis) the GPUs have significant advantage. This figure is from Ref. [1].

5 GHz clock speed when running a single core, as it does with these simulations. Despite this, the highest measured speedup between these devices is 36, indicating just how powerful the TITAN V GPU is when applied in a suitable manner. If we compare this device to more commonly used CPUs, the results would be far more substantial. For a more detailed discussion of these benchmarks, and for more details on the hardware, see *Smith et al. 2022* [1].

Based on these results, we may conclude that the TSSM as applied to the problem of integrating the GPE, for both single-component and spinor gases, benefits highly from hardware acceleration on GPUs. It is worth noting that, due to the constant execution times of the GPU up to a certain grid size, there is almost no additional cost in running higher resolution simulations on such devices. The trade-off between sufficient resolution and large enough boundaries (Sec. 3.3.1) is therefore nullified by GPUs up to a threshold grid size, which depends on the particular device. The same can not be said of CPUs, as their execution times always increase accordingly with the grid size. Therefore, not only are GPUs faster, but they also require less of a cost-benefit analysis when planning large sets of simulations.

The path towards the results discussed here was obtained quite easily, all things considered, by adopting a package that was syntactically similar to the existing code, with back-end access to CUDA libraries. High performance computing is therefore

readily accessible, and open-source, allowing researchers more time in the lab, and shortening the development cycle for powerful numerics.

3.4 Artificial Magnetic Fields in Spinor BECs

We are now equipped with everything needed to solve for the ground states of the Hamiltonian Eq. 3.15, wherein the gas of neutral atoms behave as charged particles subject to electromagnetic gauge fields. Here, we show a series of simulations that were used to verify the spinor code discussed in Sec. 3.3.1 during its development; results were compared to similar simulations detailed in *Radić et al. 2011* [31].

The situation we consider, briefly discussed in Sec. 3.1.2, is of a gradient in the detuning along the y -axis, $\delta(y) = \beta y$, where the slope of the gradient is parameterized by β in units of recoil energy per TF radius, E_L/R_x . In the strong-coupling regime, with $\Omega/EL > 4$, the resulting parabolic dispersion of the lower dressed state band depends on the location along y , in such a way that the artificial vector potential produced, $\mathbf{A} \propto \beta y \mathbf{e}_x$, has a non-vanishing curl. The atoms therefore move according to an artificial magnetic field oriented along the z -axis. We will simulate these conditions for various combinations of β and the coupling strength Ω , finding the ground states through imaginary time propagation of the spinor GPEs (Eqs. 3.42-3.42).

For each simulation that follows, the parameters Ω and β were varied with all others constant (aside from details of the grids, as explained below). The atom number was fixed at $N = 10^4$, with trap frequencies along each direction $\omega_x = \omega_y = 2\pi \times 50$ Hz, and $\omega_z = 2\pi \times 200$ Hz. The trapping potential is therefore anisotropic, with stronger confinement along z . As discussed in Sec. 3.3.1, the quasi-2D GPE simulations employed here are well suited to such a situation. The interaction strengths were those of ^{87}Rb , with $g_{\uparrow} = g_{\downarrow} = 4\pi\hbar^2 a/m$; the inter-atomic interaction in ^{87}Rb is slightly larger, $g_{\uparrow} = 1.04g_{\uparrow}$, which causes atoms to energetically prefer interactions within their spin components, rather than between them. Simulations were performed with grids covering a region $2.4R_x$ to $3R_x$ wide along x , and $4.6R_x$ to $5R_x$ along y ; these sizes were changed slightly to accommodate the wider spatial extent of the ground states subjected to larger detuning gradients, which forces the spins apart. The grid dimensions also varied between simulations, in each case with either 256 or 512 pixels along each axis.

Additionally, in all cases, wavefunctions were initialized in the TF wavefunction with equal populations in each spin, and propagated in imaginary time without any Raman coupling or detuning. The result was then propagated again, just in the presence of the detuning gradient, since it acts as a potential to shift the spatial

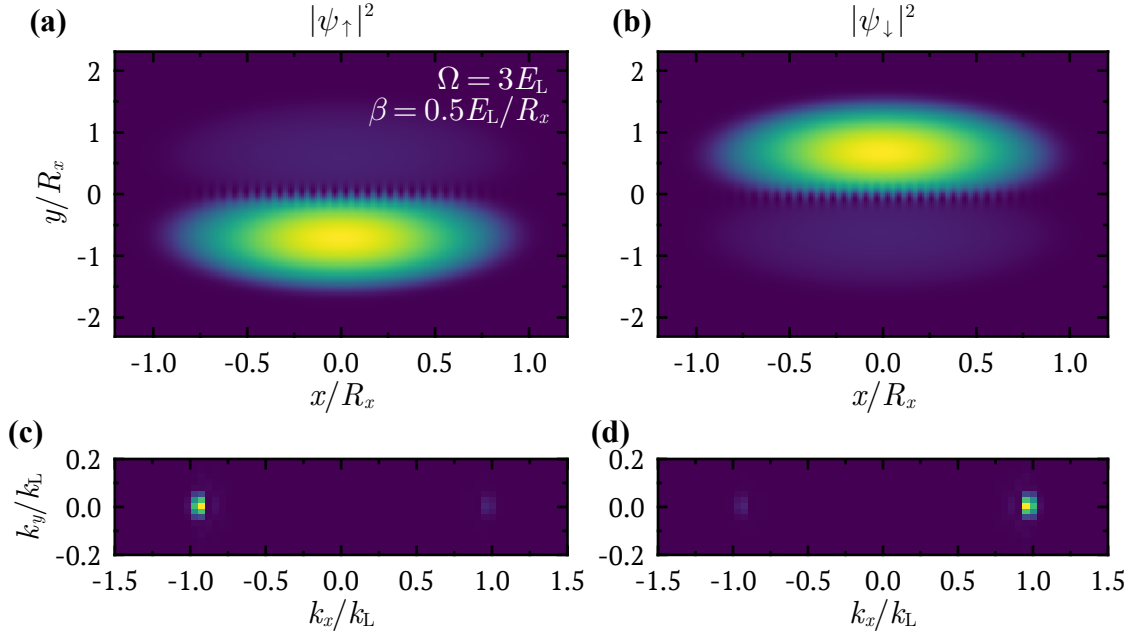


Figure 3.6: Numerically determined spinor BEC ground states, with Raman coupling $\Omega/E_L = 3$, and detuning gradient with slope $\beta = 0.5E_L/R_x$. The r -space densities are shown for ψ_{\uparrow} (a), and ψ_{\downarrow} (b). For the corresponding k -space plots, (c) and (d), a magnified region near zero is shown. Results show a phase separation due to the interplay between weak coupling, a potential energy shifted by detuning, and repulsive inter-spin interactions. In the middle region where the two components overlap, interference effects from the two oppositely propagating waves may be seen.

distributions of the spins apart. Finally, the result was propagated a third time, with all of the relevant terms applied, to determine the full ground state. We found that this process drastically reduced the time taken for results to converge. In this way, the various terms in the Hamiltonian were applied in sequence, allowing the result of each propagation loop to begin with a better estimate for the ground state. This is akin to the adiabatic preparation of states in experiment.

To begin, we will first look at the results for a relatively small detuning gradient, and in the weak-coupling regime. In Fig. 3.6 we show the ground state for each spin component in both r -space and k -space, with $\Omega/E_L = 3$, and $\beta = 0.5E_L/R_x$. In this case, since the coupling is weak, the dispersion is a double well, shifted on either side of $y = 0$ due to the small detuning gradient. The spins are therefore not substantially mixed; the relative populations remain balanced here, due to the symmetry of $\delta(y)$ across the origin. This is true for all other simulations shown here.

The detuning term $\delta(y)$ is a spatial gradient, and therefore applies a force on atoms. Effectively, the parabolic trap potential seen by each spin component has its minimum shifted away from $y = 0$; in this case, spins are pulled in different directions. In Fig. 3.6, with repulsive interactions and without sufficiently strong coupling, it is therefore energetically favorable for the spins to separate along the

y -axis. Another way to view this is through an artificial electric field, arising from the gradient of the scalar potential, since the height of bands changes along y (see Fig. 3.3). This field is zero at $y = 0$ and increases away from the origin, but due to the confining harmonic trap, atoms are still contained to a near central region. Due to the opposing momenta of each spin, in the region where they overlap slightly there is a notable interference pattern. This may be understood as two plane waves propagating in opposing directions, producing an interference pattern with a length scale given by the difference in their momenta, which here is $\approx 2k_L$. This simulation, Fig. 3.6, gives us all the ingredients necessary to understand the others.

In Fig. 3.7 we show the result of a simulation in the strong-coupling regime, with $\Omega/E_L = 5$, and $\beta = 0.8E_L/R_x$. This situation may therefore be understood through an Abelian magnetic field acting on the lower energy dressed state (see Sec. 3.1.2). The results look similar to before, though with the stronger coupling there is substantially more mixing between the spins. The dispersion being a single minimum also brings the momentum components closer to zero. The resulting interference pattern in the middle region is more pronounced as a result. In Fig. 3.8 we take a closer look at the interference pattern, plotting the total density (sum of the two spin components), alongside the phase of the spinor wavefunctions in the region near $y = 0$. The phase reveals that the interference pattern actually consists of a row of quantized vortices, each with phase winding 2π . The winding around each vortex, for both spin components, is in the same direction; this is because both spins in the lower band feel the *same* Abelian magnetic field. The Lorentz force has produced circulation in the spinor gas, resulting in the nucleation of vortices.

Looking now at the case where we increase the strength of the applied magnetic field, results from simulations with $\Omega/E_L = 4$ and $\beta = E_L/R_x$ are shown in Fig. 3.9. In this case, the coupling is slightly weaker, allowing the repulsive interactions between spins to overcome it; the resulting ground state is therefore more miscible in spin. Furthermore, several vortices appear in the high density regions of each spin, due to the strong artificial magnetic field imparting significant angular momentum on the gas. The middle row of vortices also persists. It is also worth noting that the momentum distributions have split into several peaks. The presence of vortices in the bulk necessitate such behavior, as more momentum components are required to reproduce the complicated density structure in r -space¹¹.

Finally, we look at the situation with both large coupling, $\Omega/E_L = 6$, and a large artificial magnetic field, with $\beta = E_L/R_x$. Results of such a simulation are shown in Fig. 3.10. Here we see a combination of all the previous results. The strong coupling produces significant mixing between spins, so they are less miscible,

¹¹This is just a statement regarding Fourier series, quite generally. For the structures seen in the density profiles, more Fourier components are needed to properly represent them.

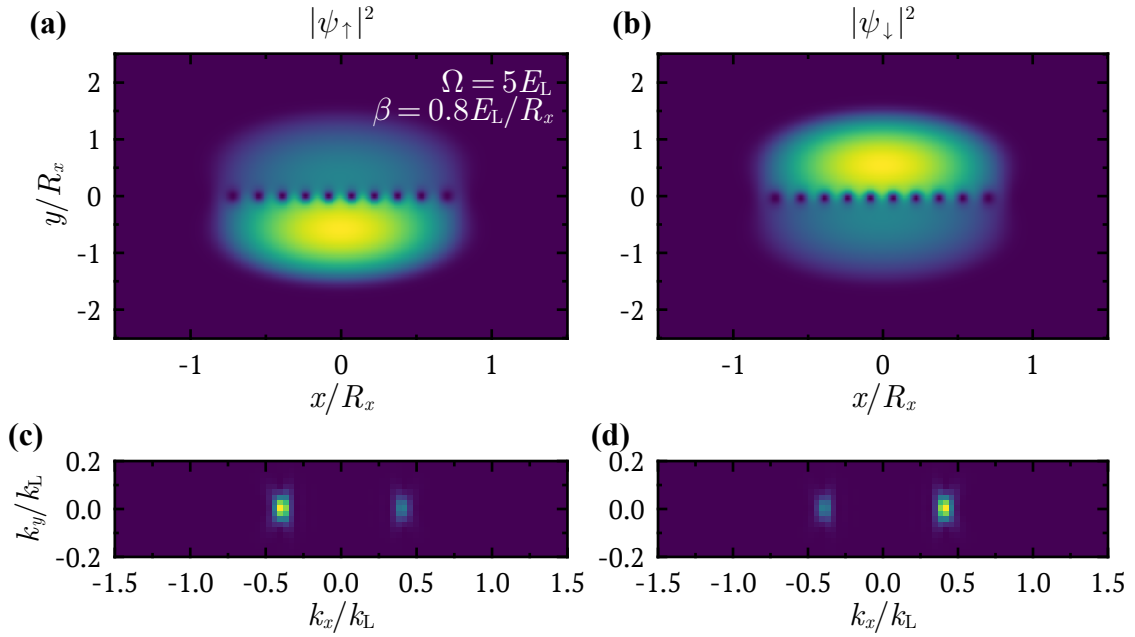


Figure 3.7: Ground state solutions in the strong-coupling regime, with $\Omega/E_L = 5$, and detuning gradient with slope $\beta = 0.8E_L/R_x$, with same layout as in Fig. 3.6. The more significant role of the coupling Ω has mixed the spins, bringing their momenta closer to zero. The resulting interference pattern from these oppositely traveling spin waves has a length scale of $\approx k_L$, producing a row of vortices where they overlap.

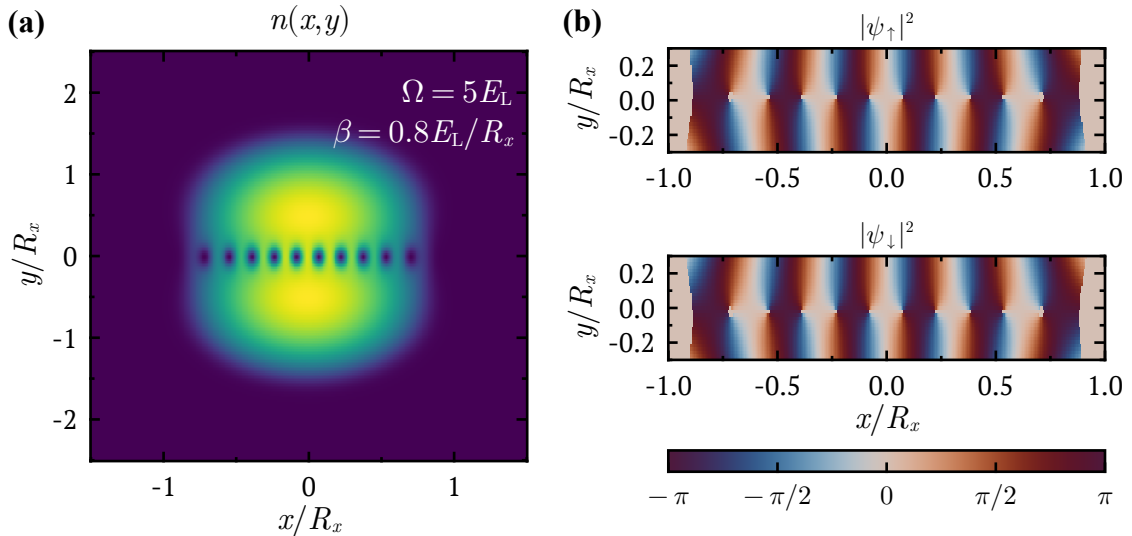


Figure 3.8: The same simulations as shown in Fig. 3.7, showing the total density (a), beside the phase distributions of each spin component (b). The full density displays an interference pattern between the oppositely traveling spin waves. Looking at the phases of each spin component in (b), with ψ_\uparrow (top) and ψ_\downarrow (bottom) each displaying a 2π phase winding around the regions where the density vanishes, indicating that these features are quantized vortices. The winding in each spin component is in the same direction, as they feel the effects of the same Abelian magnetic field.

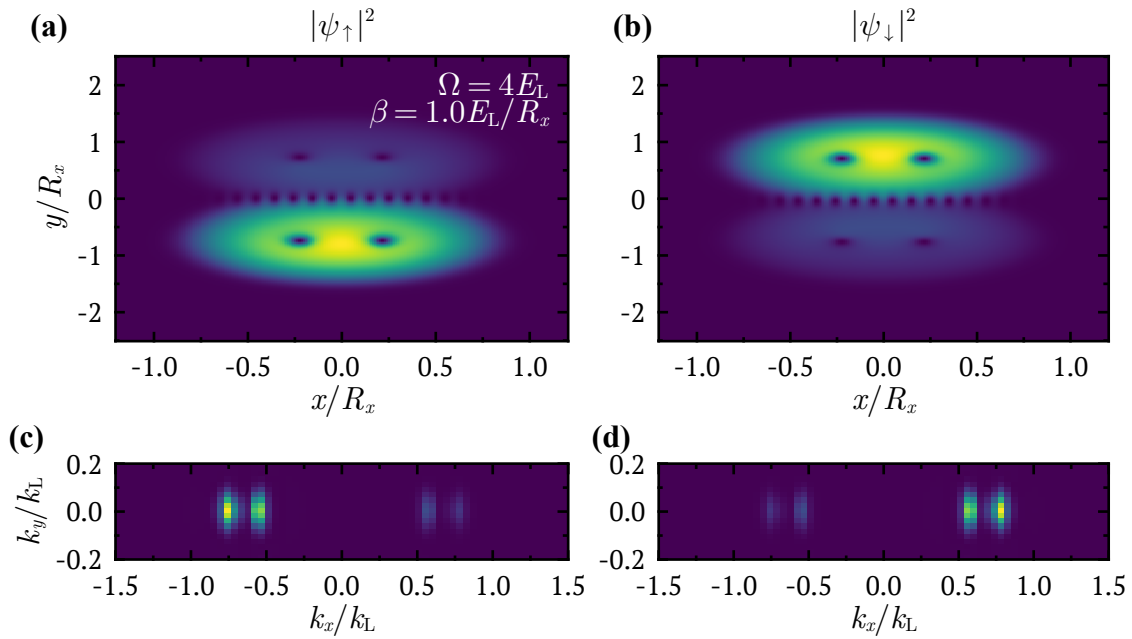


Figure 3.9: Spinor BEC densities in r -space and k -space, with the same layout as in Fig. 3.6, under similar conditions to Fig. 3.7 but with a stronger magnetic field. The coupling strength is $\Omega/E_L = 4$, and the detuning gradient has slope $\beta = E_L/R_x$. The lower coupling strength permits a higher degree of separation between spin components, but the larger gradient produces a stronger artificial magnetic field. In this case, the field is strong enough to permit the nucleation of vortices into the high-density regions. As before, all vortices have the same winding.




Figure 3.10: Spinor BEC ground states, shown with the same layout as in Fig. 3.6, with both strong coupling, $\Omega/E_L = 6$, and a large detuning gradient, with slope $\beta = E_L/R_x$. Results show spins that are less miscible than in other simulations, due to the strong coupling strength. There are several vortices, including some in the high density regions, due to the large magnitude of the artificial magnetic field.

and their momenta are brought closer together, increasing the wavelength of their interference pattern. Additionally, the large magnetic field permits vortex nucleation in the bulk.

These simulations exemplify the applicability of the GPE integration technique described in Sec. 3.3 to various complicated Hamiltonians. If many such simulations are to be performed, especially at high enough resolution to include the small vortex structures, or large enough boundaries to permit real-time propagation¹², the GPU accelerated code is expected to significantly reduce the time and resources required to undertake such investigations.

Altogether, the simulations discussed here demonstrate the complicated interplay between terms in the Hamiltonian, which lead to interesting many-body phenomena. In this case, we looked at the effect of a relatively simple Abelian electromagnetic potential, producing circulation in a spinor gas; the results show several different vortex structures [31], in only a limited region of the full parameter space. In addition to the simulations shown here, we were also able to simulate artificial electric fields by changing a spatially-uniform detuning in time, and other such

¹²If left in the trap, atoms are expected to “slosh” around to some extent, meaning the simulation boundaries must be extended. Otherwise, simulating the dynamics in TOF would be experimentally applicable, but involves the rapid expansion of the BEC requiring a significant increase in the grid size and resolution.

gauge fields. In each case, the artificial gauge fields discussed in Sec. 3.1.2 allow us to leverage our intuition from classical electromagnetism, and use it to explain the phases of many-body systems.

CHAPTER 4

Non-Abelian Geometric Phase

THE previous chapter, Ch. 3, concerned itself with quantum simulation. A technique for using Raman dressing (Sec. 2.2.2) to produce gauge-field like terms in the atomic Hamiltonian, allowed us to probe the many-body physics of an interacting Bose gas. In this case, the rather simple $U(1)$ gauge fields were more a tool, helping us to understand many-body physics, rather than the object of study itself. In this chapter, we will cover a different means of producing artificial gauge fields, wherein they will be the main topic of study and the many-body physics will be mostly ignored.

In this chapter we will focus our attention on geometric phase, which has a convenient description in the language of gauge theory. Geometric phase may arise in systems with a Hamiltonian that varies in time through a set of parameters (usually ones that can be controlled by the experimenter). One may find that the evolution of the system can be partitioned into two components: a dynamical component, which depends on the Hamiltonian and the specific dynamics of the evolution, and a geometric component, which does not depend on the Hamiltonian, but rather on the curvature of the Hilbert space as it varies with the parameters. Typically geometric phases are discussed in terms of adiabatic evolution [64, 65], where the dynamical component of the phase may be ignored, but there are non-adiabatic generalizations where they will become important [66].

In light of its connection to gauge theory, it is natural to extend our interest in generating various gauge fields to the study of geometric phase. Looking at the evolution of quantum systems, generated by geometric phase, may tell us something about the underlying gauge fields. For this reason, we will look at how geometric phases with gauge fields of interesting symmetries may be created in the lab. In particular, we will focus on non-Abelian gauge fields, which have traditionally been difficult to synthesize [30, 67]. It turns out phases of this nature also have interesting connections to several [quantum computing \(QC\)](#) approaches, which will be discussed

as well.

We begin in Sec. 4.1 by introducing the concept of geometric phase, through the adiabatic evolution of a parameterized Hamiltonian. This Hamiltonian is part of a family of Hamiltonians, all connected through the continuous variation of parameters. The instantaneous eigenstates of the Hamiltonian change as it evolves; these alterations correspond to curvature in parameter space, and so one naturally arrives at the condition for parallel transport of vectors in curved space. The effects of curvature are expressed through a vector potential, for which there is an associated magnetic field. Generally, both Abelian and non-Abelian geometric phases may arise, and the difference between the two may not be readily apparent [68]. As such, we will discuss a means to discern the difference between the two cases in a gauge-invariant manner, through *Wilson loops*.

Following this, we will more firmly draw the connection between artificial gauge fields and geometric phase, in Sec. 4.2. Some of the difficulties in studying non-Abelian gauge fields specifically will be discussed, as will their applications in various fields in physics. Later, in Sec. 4.3, we will motivate this work even further by discussing the implications of geometric phase in several [quantum information processing \(QIP\)](#) approaches. More specifically, there are several [QC](#) schemes that rely on geometric phase; they will be briefly introduced, in order to distinguish them from each other.

4.1 Geometric Phase

First we must introduce the concept of geometric phase as it arises from the Schrödinger equation. Here we will concern ourselves only with the adiabatic limit, but the result from a non-adiabatic generalization will become important in Sec. 5.3.

Suppose we have some Hamiltonian \hat{H} that depends on a set of parameters $\boldsymbol{\lambda}(t)$, which will evolve in time. We assume the parameters are under our control at all times, and take on values $\boldsymbol{\lambda} \in \mathcal{M}$. The [TDSE](#) equation is

$$i\hbar \frac{\partial}{\partial t} |\psi\rangle = \hat{H}[\boldsymbol{\lambda}(t)] |\psi\rangle. \quad (4.1)$$

We may write the Hamiltonian in a new basis, which changes with the parameters,

$$\hat{H}(\boldsymbol{\lambda}) \rightarrow \hat{U}^\dagger(\boldsymbol{\lambda}) \hat{H} \hat{U}(\boldsymbol{\lambda}) - i\hbar \hat{U}^\dagger(\boldsymbol{\lambda}) \partial_t \hat{U}(\boldsymbol{\lambda}), \quad (4.2)$$

for which the [TDSE](#) equation becomes,

$$i\hbar \frac{\partial}{\partial t} |\psi\rangle = \left(\hat{U}^\dagger \hat{H} \hat{U} - i\hbar \hat{U}^\dagger \frac{\partial}{\partial t} \hat{U} \right) |\psi\rangle. \quad (4.3)$$

All of the time dependence is contained within the parameters $\boldsymbol{\lambda}(t)$, so we may change the derivatives to those in parameter space as,

$$\begin{aligned}\frac{\partial}{\partial t} &= \frac{\partial \lambda^\mu}{\partial t} \frac{\partial}{\partial \lambda^\mu}, \\ &= \dot{\lambda}^\mu \partial_\lambda^\mu,\end{aligned}\tag{4.4}$$

where $\dot{\lambda}^\mu = \partial_t \lambda^\mu$, and $\partial_\lambda^\mu = \partial / \partial \lambda^\mu$. Summation over repeated indices is implied (Einstein summation). Changing the derivatives in Eq. 4.3 and rearranging yields,

$$i\hbar \left(\partial_\lambda^\mu + \hat{U}^\dagger \partial_\lambda^\mu \hat{U} \right) |\psi\rangle = \frac{1}{\dot{\lambda}^\mu} \hat{U}^\dagger \hat{H} \hat{U} |\psi\rangle.\tag{4.5}$$

We define $\hat{H}_\lambda = \hat{U}^\dagger \hat{H} \hat{U}$, so then we have,

$$i\hbar \nabla_\lambda^\mu |\psi\rangle = \frac{1}{\dot{\lambda}^\mu} \hat{H}_\lambda |\psi\rangle,\tag{4.6}$$

written in terms of the *covariant derivative* $\nabla_\lambda^\mu = \partial_\lambda^\mu - i\hat{A}^\mu$, where,

$$\hat{A}^\mu(\boldsymbol{\lambda}) = i\hat{U}^\dagger(\boldsymbol{\lambda}) \partial_\lambda^\mu \hat{U}(\boldsymbol{\lambda}),\tag{4.7}$$

is called the *connection*, which is an anti-Hermitian matrix. At the moment, Eq. 4.6 still contains the influence of \hat{H}_λ on the evolution, but under certain conditions discussed later, we may assume that $\hat{H}_\lambda |\psi\rangle = 0$. In this special case,

$$\nabla_\lambda^\mu |\psi\rangle = 0.\tag{4.8}$$

This is the condition for parallel transport of the vector $|\psi\rangle$ through parameter space \mathcal{M} , which has a curvature captured by the connection $\hat{\mathbf{A}}(\boldsymbol{\lambda})$. The evolution is fully described by the geometric properties of the Hilbert space, as it varies in \mathcal{M} . From the perspective of a physicist, the connection plays the role of a vector potential, from which we may then define a magnetic field that here is called the *curvature*. We will forgo providing further intuition for these relations until a later section, (Sec. 4.1.2).

4.1.1 Holonomies

To derive the solution to Eq. 4.6, which still contains contributions from the Hamiltonian \hat{H}_λ , it is convenient to first go back a few steps. Specifically, we begin with

$$i\hbar \frac{\partial}{\partial t} |\psi\rangle = \left(\hat{H}_\lambda - \hbar \dot{\lambda}^\mu \hat{A}^\mu \right) |\psi\rangle,\tag{4.9}$$

for which the solution is $\hat{\mathcal{U}}_{\boldsymbol{\lambda}}(t) |\psi\rangle$, where $\hat{\mathcal{U}}_{\boldsymbol{\lambda}}(t)$ is the time-ordered exponential¹,

$$\hat{\mathcal{U}}_{\boldsymbol{\lambda}}(t) = \mathcal{T} \exp \left\{ -\frac{i}{\hbar} \int_0^t dt' \left[\hat{H}_{\boldsymbol{\lambda}}(t') - \hbar \dot{\lambda}^\mu \hat{A}^\mu(\boldsymbol{\lambda}) \right] \right\}, \quad (4.10)$$

where \mathcal{T} is the time ordering operator. We will look at the case where the evolution is *cyclic*, so that after some time T , we have traversed a loop in parameter space such that $\boldsymbol{\lambda}(T) = \boldsymbol{\lambda}(0)$.

Next, we consider the case where the family of Hamiltonians, connected through $\boldsymbol{\lambda}$, are *iso-degenerate*, meaning that the degeneracy structure of \hat{H} is consistent throughout the entire loop in \mathcal{M} [69]. Hence, there is no crossing of eigenstates. Furthermore, we assume that the evolution is sufficiently adiabatic, so that there is no dynamical coupling between states due to changes in $\boldsymbol{\lambda}$. In this case, the integrand in the above evolution operator, Eq. 4.10, is *block-diagonal*; the n^{th} block acts on an N_n dimensional degenerate subspace, with quasi-energy $\epsilon_n(\boldsymbol{\lambda})$. Each block thus takes the form $\epsilon_n(\boldsymbol{\lambda}) \hat{\mathbb{1}}^{N_n} - \hbar \dot{\lambda}^\mu \hat{A}_n^\mu(\boldsymbol{\lambda})$, with $\hat{\mathbb{1}}^{N_n}$ the N_n dimensional identity operator. The evolution operator may then be expressed as a *direct sum* between subspaces,

$$\hat{\mathcal{U}}_{\boldsymbol{\lambda}}(T) = \bigoplus_{n=1}^D \mathcal{T} \exp \left[-\frac{i}{\hbar} \int_0^T dt' \epsilon_n(\boldsymbol{\lambda}) \hat{\mathbb{1}}^{N_n} + i \int_0^T dt' \dot{\lambda}^\mu \hat{A}_n^\mu(\boldsymbol{\lambda}) \right], \quad (4.11)$$

for D many subspaces [69]. Through the direct sum, each term acts on its own subspace, with no mixing between them. Since the first term in the exponential is proportional to identity, we can separate this into the product of two exponentials, by the Baker-Campbell-Hausdorff formula. Additionally, the second term may be changed to a closed line integral² over the loop ℓ in \mathcal{M} , and the time-ordering may therefore be changed to path-ordering, via operator \mathcal{P} . We have

$$\begin{aligned} \hat{\mathcal{U}}_{\boldsymbol{\lambda}}(T) &= \bigoplus_{n=1}^D \exp \left[-\frac{i}{\hbar} \int_0^T dt' \epsilon_n(\boldsymbol{\lambda}) \right] \mathcal{P} \exp \left[i \oint_{\ell} d\lambda^\mu \hat{A}_n^\mu(\boldsymbol{\lambda}) \right], \\ &= \bigoplus_{n=1}^D e^{-\frac{i}{\hbar} \phi_n(t)} \hat{\Gamma}_A^n(\ell). \end{aligned} \quad (4.12)$$

The dynamical contribution to the phase in subspace n is contained within $\phi_n(t) =$

¹Time-ordering is necessary because the integrand may not commute with itself at all times. Really, a time-ordered exponential is just a re-expression of the original differential equation, as the integration of the Hamiltonian has not yet been carried out. Path-ordering, which comes up later in this section, is similar.

²If we relaxed the constraint that the evolution is cyclic, the result would just be a line integral instead.

$\int_0^t dt' \epsilon_n [\boldsymbol{\lambda}(t)]$. The corresponding geometric contribution is,

$$\hat{\Gamma}_A^n(\ell) = \mathcal{P} \exp \left[i \oint_{\ell} d\boldsymbol{\lambda} \cdot \hat{\mathbf{A}}_n(\boldsymbol{\lambda}) \right], \quad (4.13)$$

which is called a *holonomy*³ [69]. Each subspace of the Hamiltonian therefore evolves with a dynamical phase according to the energy of the subspace, and a geometric contribution from the curvature.

4.1.2 Some Intuition

To develop further intuition on the geometric Schrödinger equation, Eq. 4.6, and the resulting form of the evolution operator (holonomy), we will remove the dynamical contribution to the phase. This is traditionally a choice that is made much earlier in the derivation of geometric phase, but the dynamical part is important to consider in many applications [71, 72]. Looking at a single subspace, we have that $\hat{H}_{\boldsymbol{\lambda}} |\psi\rangle = \epsilon |\psi\rangle$. This Hamiltonian, with Hilbert space of dimension N_n , therefore has an N_n -fold degeneracy (in the case where $N_n = 1$, we are concerning ourselves then with just a single state). We are free to shift the Hamiltonian in energy by ϵ such that $\hat{H}'_{\boldsymbol{\lambda}} |\psi\rangle = 0$, with $\hat{H}'_{\boldsymbol{\lambda}} = \hat{H}_{\boldsymbol{\lambda}} - \epsilon$. The Schrödinger equation therefore becomes the parallel transport condition, Eq. 4.8.

Rearranging Eq. 4.8 we obtain,

$$\partial_{\boldsymbol{\lambda}}^{\mu} |\psi\rangle = i \hat{A}^{\mu} |\psi\rangle. \quad (4.14)$$

The above equation is fully geometric; by eliminating the energetics of the states, what remains does not depend on the spectral properties of the Hamiltonian. For the typical dynamical Schrödinger equation (Eq. 4.1), we see that the action of changing the state vector in time is the same as a rotation by the Hamiltonian. Here we have a similar result: the action of moving the state vector in parameter space (the derivative on the left), is the same as a rotation by the connection. This is, therefore, the geometric equivalent of the typical Schrödinger equation that we are familiar with, where the connection has taken the place of the Hamiltonian as the generator of transformations. From the above, Eq. 4.14, we may also write the

³Also called *anholonomy* in the original work on geometric phase, as described in *Berry* 1988 [70]. This refers to a system in which quantities driving the evolution are returned to their original values, but several others are changed in this process. The system dynamics are, therefore, not fully captured by the driving quantities (nonintegrability). More recently, the term holonomy has been adopted, particularly in the quantum information community, so we will use this convention throughout the thesis.

connection in the conventional manner [65],

$$A_{nm}^\mu = -i \langle \psi_n | \partial_\lambda^\mu | \psi_m \rangle, \quad (4.15)$$

defined through the matrix elements for each vector component.

These points are made more clear by examining the form of the holonomy, Eq. 4.13. The holonomy, which is directly the solution to the above Schrödinger equation (Eq. 4.14), depends exclusively on the path ℓ taken in parameter space \mathcal{M} , as described by the connection $\hat{\mathbf{A}}$. It is important to understand this relationship; the holonomy is *not* a function of the connection, but of the path ℓ . The connection is pre-established by the structure of the eigenstates of $\hat{H}(\boldsymbol{\lambda})$ as they vary in parameter space. Through this lens, the structure of the Hilbert space defines the form of the connection, from which the various choices of path determine which transformations are generated.

Following this, if we consider the set of all possible loops \mathcal{L} , we obtain the *holonomy group* $\mathcal{G} = \left\{ \hat{\Gamma}_A(\ell) \ \forall \ell \in \mathcal{L} \right\}$ [69]. In general, the holonomy group \mathcal{G} is a subset of $U(N_n)$, where N_n is the subspace space dimension. This indicates that, based on the form of the connection, a subset of all possible unitary transformations in the subspace may be generated by the set of loops⁴. Similarly, when the holonomy group coincides with $U(N_n)$ then all possible unitaries may be generated, and the connection is said to be *irreducible*.

4.1.3 Abelian vs. Non-Abelian

Putting everything together, the evolution of the system is fully characterized by the connection; by studying this object, and the various forms it may take, we may obtain the rest of the results. There are two important cases to consider: nondegenerate and degenerate subspaces.

In the nondegenerate case, evolution of states is constrained to each subspace of dimension $N_n = 1$. The components of $\hat{\mathbf{A}}_n$ must therefore be real numbers⁵, which commute making the connection Abelian. Looking then at the form of the holonomy, Eq. 4.13, we can ignore the path ordering since the connection commutes with itself everywhere along the path. The result is that the holonomy group for each subspace is a subset of $U(1)$. Each eigenstate therefore obtains a global phase (global in the subspace), with a dynamical contribution from the energy of the state, and a geometric contribution from the curvature in the subspace. Altogether, an

⁴Note that the composition of several loops, such as $\ell_1 \circ \ell_2$, is itself a loop. The corresponding holonomy may be written $\hat{\Gamma}_A(\ell_1 \circ \ell_2) = \hat{\Gamma}_A(\ell_2) \hat{\Gamma}_A(\ell_1)$. Therefore, members of the holonomy group may be generated as products of holonomies, by the composition of loops.

⁵Looking at the holonomy in Eq. 4.13, for it to be unitary the result of the integration must be real.

arbitrary state across several subspaces would therefore obtain some relative phase through this process, but the population in each eigenstate would remain the same since there is no coupling between subspaces. In nondegenerate systems, in the adiabatic limit, the geometric phase is therefore *always* Abelian.

The result for the simple $N_n = 1$ dimension subspace is instructive, as it is reasonably easy to see that *the only* possibility is for the state to obtain a global phase. This can be seen from the time-independent Schrödinger equation, $\hat{H}(\boldsymbol{\lambda})|\psi\rangle = \epsilon|\psi\rangle$. Since the evolution is adiabatic and cyclic, at time T the state has to have returned to the initial eigenstate, as this is the only solution to the eigenvalue problem here, but a global phase is always permitted. Therefore, the only possible solution is $|\psi(T)\rangle = e^{i\phi}|\psi\rangle$. This is the case originally studied by Berry [64], and describes the geometric origin of the global phase seen in the seminal Aharonov-Bohm effect [73, 74].

The other case to consider is what form the phase takes on in degenerate manifolds, that is in subspaces with dimension $N_n > 1$. In this case the dynamical phases accumulated by all states in the manifold are the same, as they have identical quasi-energies throughout $\boldsymbol{\lambda} \in \mathcal{M}$ due to the iso-degenerate property of the Hamiltonian. The geometric contribution is matrix-valued, and in general need not be diagonal in the subspace eigenbasis. It may then be the case that the connection's components do not commute, making it non-Abelian, but this depends specifically on the form of the curvature; there may yet be matrix-valued Abelian phases, if the components of the connection always commute in all gauges. Degeneracy is therefore a necessary but insufficient condition for the existence of non-Abelian geometric phase. Later, in Sec. 4.1.4 we will introduce a means to distinguish these two cases.

Due to the non-commutative property of non-Abelian geometric phases, the coupling of states within a degenerate subspace is permitted; as with the $N_n = 1$ case, this is reasonably easy to see. Once again consider the time-independent Schrödinger equation $\hat{H}(\boldsymbol{\lambda})|\psi_k\rangle = \epsilon|\psi_k\rangle$, which now has a set of degenerate eigenstates $\{|\psi_k\rangle\}$. Suppose, *e.g.* at time $t = 0$ the system was prepared in one of the eigenstates, $|\psi_1\rangle$ (this choice is arbitrary). Following cyclic adiabatic evolution (and ignoring the global dynamical phase), there is no reason that the state must return to the initial, as any arbitrary superposition of the basis states satisfies the Schrödinger equation, $|\psi(T)\rangle = \frac{1}{\mathcal{N}} \sum_k \tilde{c}_k |\psi_k\rangle$. The holonomy is therefore, in general, a unitary operator which couples the degenerate manifold; which couplings are permitted therefore depend on the particular form of the connection, but as stated above, when it is irreducible, any arbitrary unitary transformation on the initial state is permitted. Unlike their Abelian counterparts, non-Abelian holonomies are therefore useful in quantum control, being the key ingredient in several approaches to QC (see Sec. 4.3).

4.1.4 Wilson Loops

The holonomy, Eq. 4.13, demonstrates that for the geometric phase to be non-Abelian, a degenerate subspace is required; however, as pointed out in Sec. 4.1.3, this is a necessary but insufficient condition. Depending on the structure of the degenerate eigenbasis in parameter space, *i.e.* the connection, the geometric phase could be matrix-valued, yet still Abelian [68]. This could result from a connection that is expressed in terms of only a single generator of the group, though, since the connection is gauge-dependent it is difficult to say based on this alone. Computing the field (curvature) from the connection may also be misleading; the field is only gauge invariant in Abelian cases, while in the more general non-Abelian case they are gauge covariant, due to the presence of a commutator between components of the connection, the *structure factor*. This means that the field does in general change with the choice of gauge (basis). In both of these situations, measurement is made difficult because observables are always gauge invariant⁶. As such, we require a rigorous way of distinguishing the two cases through available measurements. Furthermore, this method must be gauge-invariant, as the choice of basis can in several ways deceive us, as briefly described above.

The appropriate quantity of interest here is the Wilson loop [68, 75], which is relevant only in the case of cyclic evolution, where the system Hamiltonian, $\hat{H}(t=0) = \hat{H}(t=T)$, returns to its initial form after some time T . From the corresponding evolution operator \hat{U}_\circ , where we have used the subscript to denote cyclic evolution, the Wilson loop is defined as

$$\mathcal{W} = \text{tr} \left(\hat{U}_\circ \right). \quad (4.16)$$

This quantity is gauge invariant [76]. Under a global change of basis by unitary \hat{U} , this is easy to demonstrate: due to the cyclic invariance of the trace, we have

$$\begin{aligned} \mathcal{W} &= \text{tr} \left(\hat{U}^\dagger \hat{U}_\circ \hat{U} \right), \\ &= \text{tr} \left(\hat{U} \hat{U}^\dagger \hat{U}_\circ \right), \\ &= \text{tr} \left(\hat{U}_\circ \right), \end{aligned} \quad (4.17)$$

and hence the result doesn't depend on the basis chosen. In the case where non-cyclic evolution is considered, this may be referred to as the *Wilson line*, which is no longer gauge invariant.

⁶Consider in classical electrodynamics that the electric and magnetic fields may be measured; since this is an Abelian theory, the fields are gauge invariant. In the non-Abelian case this is no longer true. Observables then, must be comprised of several terms, each gauge covariant, such that together through the cancellation of terms, the observable itself is gauge invariant. Therefore, determining the degree to which the fields are only gauge covariant is a significant challenge.

Conceptually, the Wilson loop is a gauge-invariant sum of eigenvalues of the transformation $\hat{\mathcal{U}}_\circ$, whatever that transformation may be. The corresponding eigenvectors describe states that are unchanged by the evolution. We are mainly concerned about transformations $\hat{\mathcal{U}}_\circ$ that are holonomic (Eq. 4.13). In this context, the evolution operator $\hat{\Gamma}_A(\ell)$ arises from the adiabatic change of an eigenbasis through parameter space; the Wilson loop thus encodes the residual distortions of that eigenbasis over the path ℓ .

The importance of the Wilson loop here is that it may be used to distinguish Abelian phases from non-Abelian. The typical manifestation of a non-Abelian generator, be it dynamical, geometric, or some combination of the two, is *path-dependence* of the transformations⁷. For a particular Hamiltonian, different choices of parameters varied in time, effecting either dynamical or geometric transformations, result in different unitaries (this is essentially the dynamical extension of the holonomy group discussed in Sec. 4.1.2). If applied to some initial state, the result will depend on the order of such operations in the case of a non-Abelian phase. This is a direct consequence of the non-commutativity between the generators of such transformations. Unfortunately, this result is once again not necessarily gauge invariant [68]; however, this is where the Wilson loop comes in. Demonstrating path-dependence of the Wilson loop, rather than the evolution operators themselves, *is* a gauge invariant manifestation of non-Abelian generators.

To see how it may be used for this purpose, consider the evolution of a single degenerate subspace; as described above (Sec. 4.1.2), the dynamical contribution to the phase is global and may therefore be ignored. What remains is a purely geometric phase, described by the holonomy $\hat{\Gamma}_A(\ell)$. A path dependence of the Wilson loop implies that, for at least two different orderings of available loops, the corresponding Wilson loops are different. This therefore requires a minimum of three loops, due to the cyclic invariance of the trace⁸. Therefore, to demonstrate a path dependence of the Wilson loop we require a minimum of three loops, ℓ_1 , ℓ_2 , and ℓ_3 , with their corresponding holonomies. The Wilson loop for a chosen path order is therefore,

$$\mathcal{W}_{ijk} = \text{tr} \left(\hat{\Gamma}_A(\ell_k) \hat{\Gamma}_A(\ell_j) \hat{\Gamma}_A(\ell_i) \right). \quad (4.18)$$

If for some non-cyclic permutation of these loops,

$$\mathcal{W}_{ijk} - \mathcal{W}_{jik} \neq 0, \quad (4.19)$$

⁷We use the word ‘‘path’’ quite generally here: for dynamical phases, the path refers to time-ordering of operations, while for geometric phase it is path-ordering.

⁸With just two loops, ℓ_1 and ℓ_2 , the Wilson loop is $\text{tr} \left[\hat{\Gamma}_A(\ell_1) \hat{\Gamma}_A(\ell_2) \right] = \text{tr} \left(\hat{\Gamma}_A(\ell_2) \hat{\Gamma}_A(\ell_1) \right)$. There is no non-cyclic permutation of just two elements, so the Wilson loop will never depend on the ordering of these operations, even though the holonomies $\hat{\Gamma}_A(\ell_1 \circ \ell_2) \neq \hat{\Gamma}_A(\ell_2 \circ \ell_1)$.

then the phase is non-Abelian, otherwise it is Abelian. Expressed in terms of the holonomies themselves,

$$\mathcal{W}_{ijk} - \mathcal{W}_{jik} = \text{tr} \left\{ \hat{\Gamma}_A(\ell_k) \left[\hat{\Gamma}_A(\ell_j), \hat{\Gamma}_A(\ell_i) \right] \right\}, \quad (4.20)$$

where we refer to this quantity as the *trace-commutator* hereafter. While this is a special case of how one would test the path dependence of the Wilson loop, with three loops, this is just the minimum requirement. More generally, *any* non cyclic permutation of equal to or greater than three loops will suffice.

There is, however, an important caveat to the condition in Eq. 4.20: even for a truly non-Abelian connection, one can not use any arbitrary set of loops. The easiest way to see this is if each of the loops produce the same generator (even if they are different loops) then the holonomies in Eq. 4.20 trivially commute, and one would incorrectly infer that the connection is Abelian. The three chosen loops must in fact produce three distinct non-commutative generators for the trace-commutator to not vanish; this is akin to having a linearly independent set of holonomies. In some restricted control circumstances, it may not be possible to produce enough generators, in which case this method will not work. Otherwise, an appropriate set of loops should be identified. The numerical value of the trace-commutator will also depend on the magnitude of the generated phases and their commutators, so it is also possible that from the chosen loops, the non-Abelian signature is arbitrarily small, which will also pose experimental difficulties.

Applying the Wilson loop in this way equips us with the necessary test to distinguish Abelian from non-Abelian transformations. If the trace-commutator (Eq. 4.20) does not vanish, this is a gauge-invariant manifestation of the non-commutativity between the components of the connection. Importantly though, the Wilson loop does not indicate whether these transformations resulted from geometric phase or otherwise. As such, if one is looking to infer the nature of the connection, *all* non-trivial dynamical contributions to the phase must be eliminated or precisely accounted for; this point will become important later in Ch. 5 and Ch. 6.

With the necessary and sufficient condition identified, one is left with the task of performing such a measurement. Measuring the holonomy $\hat{\Gamma}_A(\ell_i \circ \ell_j \circ \ell_k)$, for different orders of the path composition is necessary. This requires some fairly sophisticated state-tomography (App. C), which will not scale well with the size of the subspace as many more measurements are required. Furthermore, it is not always the case that projections within the subspace in question are available, which will pose additional complications for the experimenter. For each subspace basis state, $|\psi_n\rangle$, the state following the transformation, $|\psi'_n\rangle = \hat{\Gamma}_A(\ell_i \circ \ell_j \circ \ell_k) |\psi_n\rangle$, must be fully characterized. Knowing how the basis transforms yields the full transformation; the

result from the n^{th} state is the n^{th} column of the holonomy. Once both holonomies are measured, the difference in their traces may be computed. Thankfully, in the case where one is attempting to demonstrate a non-Abelian phase, we are looking for the trace-commutator to *not* vanish; it is much easier to experimentally show that something is not zero than otherwise⁹.

4.2 Geometric Phases as Artificial Gauge Fields

The importance of gauge theory to quantum systems, and its deeply geometric origins, was made apparent through the Aharonov-Bohm effect [73, 74], where a charged particle may obtain a phase by encircling a solenoid containing a magnetic field, which remains confined inside the solenoid. Despite the particle never encountering the magnetic field itself, as it is confined to the solenoid, the vector potential extends beyond, resulting in a geometric phase that was first described in *Berry* 1984 [64]. This phase is the result of *real* electromagnetic fields, in this case, though the entire notion of geometric phase as described in Sec. 4.1 follows the same fundamental principles. The motion of a particle through some parameter space, be it coordinate space or otherwise, may obtain geometric phase arising from the curvature in that space. The quantum simulation aspect to this is therefore that parameter space is treated quite literally as coordinate space, so even the evolution of neutral atoms may be made equivalent to that of charged particles moving in a manifestly real gauge field [38, 77].

Thus far we have mainly concerned ourselves with discussing the connection, but just as in classical electromagnetism we may also look at the associated magnetic fields (curvature) [77]. Looking at the form of the holonomy, Eq. 4.13, the geometric phase is expressed through a line integral around the loop ℓ in parameter space. In the Abelian case, for instance, using Stokes' theorem we may re-write this integral as,

$$\oint_{\ell} d\lambda \cdot \mathbf{A} = \iint_{\mathcal{S}} d\mathbf{S} \cdot \boldsymbol{\Omega}, \quad (4.21)$$

$$\boldsymbol{\Omega} = \nabla_{\lambda} \times \mathbf{A}, \quad (4.22)$$

where $\boldsymbol{\Omega}$ is the curvature (Berry curvature, more specifically). In this representation, we may think of geometric phase as arising from the magnetic flux through the surface \mathcal{S} enclosed by the loop ℓ . This extends beyond Abelian geometric phases,

⁹Consider making a measurement with the corresponding uncertainty. When attempting to show that a quantity vanishes, how small must the error bar be to demonstrate this? Conversely, when demonstrating that a quantity does not vanish, the success is easily determined by the relative distance of the value from zero, weighted by the uncertainty.

for which this form is quite directly analogous with those from electromagnetism, as there are non-Abelian generalizations of Stokes' theorem¹⁰ [78].

The Abelian Berry connection [64] from the original Aharonov-Bohm effect, formulated in the context of particles without internal structure, was generalized to degenerate systems, resulting in the Wilczek-Zee connection [65], contained within the holonomy (Eq. 4.13 discussed here). Similarly, this was further generalized to non-adiabatic systems, resulting in the Anandan connection [66] (see Sec. 4.3.1 for further discussion). It should be no surprise then, that this was generalized even further to open quantum systems, with the Uhlmann connection [79], and corresponding holonomy that acts on density matrices instead of state vectors. As a result, gauge fields (non-Abelian ones in particular) have played a crucial role in understanding the evolution of quantum systems, across many disciplines.

As mentioned prior, geometric phase can serve as a tool with which to understand other phenomena, such as the many body physics described in Ch. 3. Furthermore, it is interesting in its own right as a means of quantum control, as we will discuss in the next section in the context of QC, Sec. 4.3. In addition to these applications, the gauge fields themselves are also of immense interest [30, 77, 80]. In particular, the ultimate goal for many is the realization of *dynamical* gauge theories in quantum simulators [81–83]. The fields described so far act on particles, imparting transformations on them; there is, however, a missing ingredient for this to be a full gauge theory: the particles must, in turn, act back on the gauge fields. For instance, in classical electromagnetism, magnetic fields arise from current density, and electric fields from charge density; these fields are inextricably connected to the matter upon which they act.

The scheme for producing artificial electromagnetic potentials acting on neutral atoms, discussed in Ch. 3, is limited in this respect [28, 29, 32–34]. Changing laser parameters allows us to tune the fields produced, and the neutral atoms behave as charged particles, responding to such fields in the expected way. This gauge theory is, however, incomplete, because these artificial charges only respond to the fields we produce; there is no notion of charge or current density producing their own fields here.

Simulating a complete gauge theory is difficult, as allowing particles to create gauge fields of their own requires interactions. One approach involves the generation of density-dependent gauge fields [84]; these are fields that depend on the density distribution of the particles they act on, therefore including the necessary feedback in dynamical gauge theories. In one such study, researchers demonstrated a topological gauge theory in an optically dressed BEC [85]. They produced a gauge-field that

¹⁰The non-Abelian curvature includes the commutators between the components of the connection (structure factors), and is therefore only gauge covariant.

depends on the density of atoms in the ensemble. Allowing the BEC to expand in 1D, for instance, produces an electric field due to the time-changing density. As the condensate expands, atoms feel an electric force along the expansion axis, resulting in asymmetric expansion. Relatedly, researchers produced artificial gauge fields in *shaken* optical lattices (periodically modulated), which appear in the system Hamiltonian in a similar way as those discussed in Ch. 3, but the magnitude of the fields are density-dependent [86]. These experiments provide impactful steps forward in simulating complete gauge theories, though the fields produced in both cases are Abelian; a non-Abelian extension of this would be a significant step forwards.

Quantum simulators capable of producing artificial dynamical gauge fields would permit the controlled study of many systems that are otherwise experimentally inaccessible. For instance, there is immense interest in the simulation of lattice gauge theories for understanding materials [81, 82, 87], synthetic spin-orbit interactions [34, 80, 88, 89], and aspects of high-energy physics such as quantum chromodynamics [83]. In all cases, the models being simulated may be both dynamical and non-Abelian, making such experiments technically demanding.

Many challenges remain in realizing such gauge-fields, though there have been several promising experimental demonstrations, for density-dependent gauge fields [85, 86], and lattice gauge theories [90]. There has, therefore, been significant interest in realizing increasingly sophisticated geometric phases, with the ultimate goal of finding ways to engineer the desired gauge fields. Despite the numerous proposals for realizing both Abelian and non-Abelian gauge fields [91–103], including through digital quantum simulators [104–106], there have been remarkably few experimental demonstrations [39, 67, 107–110] in the context of artificial gauge fields¹¹. As such, there is more work to be done in finding ways to generate even non-dynamical fields, to eventually realize quantum simulation of dynamical gauge theories. This is the primary motivation for the work discussed in Ch. 5 and Ch. 6.

4.3 Geometric Quantum Computing

In addition to their interpretation as artificial gauge fields (Sec. 4.2), geometric phases also have their applications in quantum control. We have seen how, even in the Abelian case, the adiabatic motion of a state through parameter space can effect transformations that are independent of system dynamics. The natural extension of this is to see how geometric phases might be used to impart specific transformations that could be useful in QIP, and determine what benefits there may be to this approach over more conventional dynamical ones, if any. It turns out that there

¹¹There has, however, been significant progress in generating various geometric phases for quantum computing, as discussed in Sec. 4.3, but the interests there lie more in robust quantum control.

are several approaches to QC that rely on geometric phase; we will refer to these approaches, together, as **geometric quantum computing (GQC)**. Before we discuss several of these geometric approaches, we will very briefly introduce some of the important principles of more conventional circuit-based QC; this is *not* intended to be an exhaustive description, there are many great reviews with such treatment [111–113].

We ought to be aware of the binary nature of classical computing, where information is encoded into bits, and processed through *gate* operations between bits. Due to the inherent binary nature of classical bits, there are a limited set of operations that may be done. When designing an algorithm, eventually it must be broken down into this set of primitive gate operations. A direct analogue of this to quantum systems would be to encode information into two orthogonal quantum states of a system, $\{|\psi_1\rangle, |\psi_2\rangle\}$. Through evolution under the TDSE, all of the same primitive gates may in general be applied to these quantum states. Various quantum control techniques may be used to generate unitary transformations that are equivalent to classical gates, *i.e.* classical computing can be done with the quantum states of a system (whatever the states may physically represent, such as spin). The broad interest in QC comes from the notion that, perhaps with resources only available to quantum systems, we may be able to perform computations much faster than on classical devices [5, 114–116]; this is referred to as *quantum advantage*.

The key insight with QC is that quantum systems can go beyond binary encoding; states are not limited to having full population in one constituent level or the other, they may be in a superposition. The quantum bit, or *qubit*, more generally can be written as,

$$|\psi\rangle = c_1 |\psi_1\rangle + c_2 e^{i\phi} |\psi_2\rangle, \quad (4.23)$$

where the coefficients c_1, c_2 are the probability amplitudes for each of the qubit levels, and ϕ is a relative phase. Since a qubit may be in a superposition of its levels, this already introduces more available gate operations that can be done. This on its own though is not enough to see computational advantages over classical computers.

The true power of a QC comes from having multiple qubits. Due to entanglement between states, the size of the full Hilbert space scales as 2^N , for N many qubits. The amount of information that can be encoded into just a handful of qubits may therefore far exceed the capacities of classical machines. In practice, this compression advantage is complicated though. Algorithms that would require this information to be read out from the QC after processing (such as performing full state tomography of the output register) would not have avoided the storage issue at all. This form of quantum advantage may only be considered in cases where only a small part of the output is necessary.

The other key idea behind [QC](#) is that entanglement may be used to generate non-trivial correlations in systems, which can be leveraged for computational speedup. This ultimately comes down to the fact that, due to the size of the Hilbert space, and the notion of superposition, there are more primitive gates available in quantum systems. Since quantum computers have these additional capabilities, more sophisticated algorithms may be designed, in principle, which may have more favorable scaling in problem size when compared to their classical counterparts. This is still an active area of research [[117–119](#)].

Despite all of the significant progress in this field, there remains a lot of work to be done before quantum advantage is conclusively demonstrated on large-scale for broadly-applicable problems. This is made more complicated by the fact that new classical algorithms are still being created, so much of the remaining work must be spent on the development of quantum algorithms. Aside from this, existing devices have limited capabilities. Much of this inefficiency can be attributed to the qubits themselves, which suffer from decoherence. Furthermore, the control schemes are prone to noise and miscalibration, reducing the fidelity of gate operations. Remarkably, there are *error correction* schemes which can drastically reduce the influence of errors [[120](#)], but they require multiple “physical” qubits to be combined together into a single “logical” qubit; sometimes, even 10-1000 physical qubits are necessary, which already far exceeds the number available on even the largest devices.

There are, therefore, many areas of [QC](#) that need attention before quantum advantage is demonstrated. Here we introduce three main [GQC](#) approaches that each address different issues in the more conventional schemes; they all differ in significant ways, but ultimately share a dependence on geometric phase. We place particular emphasis on the first, [holonomic quantum computing \(HQC\)](#), as it is the most relevant to the experiments composing the primary work in [Ch. 6](#). We then briefly discuss [topological quantum computing \(TQC\)](#), and [adiabatic quantum computing \(AQC\)](#), in order to distinguish the three from each other. Together, these techniques are of great interest, identifying yet another way in which geometric phase is worthy of further study.

4.3.1 Holonomic Quantum Computing

The first [QC](#) scheme to introduce, which is also the most relevant to this work, is [HQC](#) [[69](#), [72](#), [121](#), [122](#)]. This protocol aims to perform gate operations geometrically, rather than the more conventional dynamical gates. In [HQC](#), logical qubits are encoded onto degenerate manifolds. Through the adiabatic change of a parameterized Hamiltonian, the unitary evolution is described precisely by the holonomy, [Eq. 4.13](#), with a non-Abelian connection. As discussed in [Sec. 4.1.3](#), a non-Abelian

geometric phase is capable of coupling states within the degenerate manifold, which is a necessary requirement for performing even classical gates, such as a bit flip operation. This can not be accomplished by Abelian geometric phases. Gates between multiple qubits are generally more complicated, but there are holonomic approaches to this as well following the same structure: engineer a degeneracy in the entangled Hilbert space, and subsequently couple qubits through non-Abelian holonomies.

For both single-qubit and multiple qubits, different gates are performed by the choice of path ℓ in parameter space. The requirement for universal HQC is that *any* unitary transformation on the set of qubits is possible, *i.e.* the holonomy group should coincide with $U(2^N)$ for N qubits. As stated in Sec. 4.1.3, this means that the connection should be irreducible. Practically, we do not need to concern ourselves with every possible loop in order to produce arbitrary unitaries. As with dynamical gates, applying consecutive holonomies (loop composition) can lead to the desired results. This is advantageous, since computing the path-ordered exponential in Eq. 4.13 for arbitrary paths is challenging. It has been demonstrated that a universal set of gates is achievable on many different platforms [72].

Holonomic QC requires that qubits host degenerate levels; as stated before this is a necessary but insufficient condition for there to be non-Abelian geometric phases. Oftentimes, this is the most complicated requirement to meet. The degeneracies must be robust to perturbations, and the energetic structure must also be consistent throughout the parameter space of the control Hamiltonian (iso-degenerate). For this reason, intermediate levels are often required; multiple internal levels are mutually coupled, such that the Hamiltonian permits a degenerate subspace, composed of dressed states (Sec. 2.2). Dressed states are superpositions of the more physically accessible *bare* states of a system¹², that is, the eigenstates permitted by the Hamiltonian in the absence of control fields (such as the eigenstates of \hat{F}_z in atomic systems). This presents new challenges, as performing measurements a subspace composed of non-trivial superpositions of bare states can be difficult [106]. Furthermore, the necessity for intermediate levels allows for additional errors and decoherence mechanisms.

Supposing the required degeneracies have been acquired, there are some more general concerns in performing measurements in the subspace, since degenerate states can not be energetically distinguished¹³. As such, at the time of measurement, it is often necessary to coherently lift the degeneracy [69, 122]; having this capability inherently means that the degeneracy is *not* robust to certain controls that

¹²In order to couple the bare states of a system, the control (dressed) Hamiltonian *must* be off-diagonal. It therefore does not commute with the bare Hamiltonian, hosting a new eigenbasis composed of superpositions of the bare levels.

¹³In a nondegenerate system, control pulses resonant with one qubit level but not the other can be used to read out the relative populations.

are in place (noise in those controls can perturb the degeneracy during operations, contributing to computational errors).

Despite these complications, there are some potentially valuable advantages to performing gates holonomically. Specifically, HQC is thought to host intrinsic *fault tolerance* when compared to dynamical gates [72]. This can be attributed to the adiabatic and geometric nature of the holonomy. First, due to the adiabaticity, holonomies are impervious to changes in the rate at which loops are traversed, so long as the adiabatic condition holds. This is because the holonomy does not depend at all on the duration of the gate. Furthermore, geometric gates are generally tolerant to high frequency fluctuations of parameters; these fluctuations average out over the loop. Intuitively, consider the description in Sec. 4.2 of the holonomy in terms of the curvature, where the integration of the flux through the surface enclosed by the loop yields the resulting phase. Continuous perturbations of the loop that leaves the flux unchanged has no impact on the phase. In the non-Abelian case, this is a complicated condition to uphold, but generally this is the intuition ascribed to holonomic evolution.

There are, however, reasons to doubt the truth behind this intuition. In particular, a study by *Colmenar et al.* 2022 [123] investigated whether there are fundamental differences between dynamical and geometric gates, when it comes to errors arising from imperfect control. They developed an analytical means to transform fully geometric gates into fully dynamical ones; the transformation was continuous, so that the two extremes are connected by a family of gates that contain both geometric and dynamical contributions to the phase. This family of gates had, by construction, precisely the same duration, resulting unitary operation, and most importantly, identical noise spectra with respect to control parameters. Therefore, the fault tolerance ascribed to any given gate has nothing to do with the nature of the phase, be it geometric or dynamical, or some combination of the two, but rather with the details of the control scheme itself. The result strongly challenges the thought that geometric gates are intrinsically fault tolerant; whatever degree of fault tolerance there may be has nothing to do with the geometric nature of the phase, but everything to do with the particular control environment that enables such transformations. It should be noted that this result does not discredit the numerous studies that have quantified the degree of fault tolerance in holonomic gates against their dynamical counterparts [124–131], it merely puts the reasoning for such relative tolerance into question. Regardless, holonomic gates are generally more tolerant to faults than their dynamical counterparts, though perhaps the mechanism by which this is true may be debated.

Non-Adiabatic Holonomic Quantum Computing

Despite the proposed advantages of [HQC](#) over conventional dynamical techniques, one of its sources of strength is simultaneously perhaps its greatest weakness: holonomic gates are adiabatic, and are therefore slow, in comparison to dynamical gates. Unlike classical computers that can store bits almost indefinitely, quantum computers suffer from decoherence effects. The longer gates take to perform, the fewer that can be done within the coherence time of the device. In the language of [QC](#), this affects the *circuit depth*, which can be defined as, among all the qubits in the circuit, the highest number of gates applied to any given one of them. The circuit depth therefore sets the minimum duration of the computation¹⁴. The adiabatic condition for holonomic gates is problematic for practical implementations in the current noisy intermediate-scale quantum era.

In response to this limitation, a protocol that is closely related to [HQC](#) has been introduced, in which the adiabatic condition has been relaxed; this protocol is referred to as [non-adiabatic holonomic quantum computing \(NAHQC\)](#) [71]. Removing the adiabatic condition has several interesting consequences. First, and most obviously, loops may be performed faster. Although, the resulting dynamical effects, such as the coupling between subspaces, may no longer be ignored. Gates are constructed in specific ways so that these dynamical effects effectively cancel out; this imposes more strict conditions on the gates, but result in fully geometric gates as in traditional [HQC](#). Furthermore, it turns out that the need for a degenerate subspace is also no longer necessary, and yet there remains a non-Abelian connection. More specifically, this connection is not the same as that of Wilczek and Zee [65] (Eq. 4.7). The holonomy takes on a new form [66],

$$\hat{\Gamma}_A(t) = \mathcal{T} \exp \left\{ -\frac{i}{\hbar} \int_0^t dt' \left[\hat{H}(t') - \hbar \partial_{t'} \boldsymbol{\lambda} \cdot \hat{\mathbf{A}}(\boldsymbol{\lambda}) \right] \right\}, \quad (4.24)$$

where $\boldsymbol{\lambda} = \boldsymbol{\lambda}(t)$.

The geometric term in Eq. 4.24 looks generally as it did before, though it may now act across subspaces; this is the Anandan connection [66]. It is still fully geometric, in that it does not depend on the details of the Hamiltonian, and is non-Abelian since its components do not, in general, commute. Furthermore, the Hamiltonian can no longer be factored out, as it does not commute with the connection. As a result,

¹⁴In general there are far more complicated optimizations to be done. A computation typically has a pre-defined target error threshold. In addition to the error rates of the qubits themselves (which may depend on the specific qubit, due to fabrication imperfections), the various gates have different error rates. In order to achieve the target error threshold, not only does the total time of the computation matter, but also the gates that the circuit is composed of. There are, therefore, optimizations to be done here, such as choosing a longer circuit with better gates as opposed to the converse.

the form of this non-adiabatic holonomy is generally quite complicated; it may be expressed as an infinite series expansion of the products of these two terms¹⁵, at each point in time. While it may seem that we can write the two terms separately, as the integrand in Eq. 4.24, due to the time ordering and this complicated exponential expansion, the effects of the two parts are inextricably linked. In NAHQ, specific loops are chosen over which the dynamical term vanishes [71].

Since the original seminal work on traditional (adiabatic) HQ [69, 121, 122], there have been many developments. In addition to many general characterizations of holonomies [132], different qubit encodings [133], and more fault-tolerant approaches [134, 135], there have also been various more specific proposals for practical implementation [136–141], even including a quantum memory protocol in neutral atoms [142]. The approach has been experimentally realized in liquid nuclear magnetic resonance (NMR) systems [143], trapped ions [144], photonics devices [145], nitrogen-vacancy centers [146, 147], and neutral atoms [107, 148, 149].

More recently, NAHQ [71] has seen increased attention in the literature due not only to its speed advantage over traditional HQ, but also because it does not rely on degeneracy. Much work has been done in speeding up gates even further [150–152], making them more robust to error [153–155], adapting new qubit (or qudit) encodings [156–160], and making it more generally applicable [161–163]. Advances in control techniques include using shorter paths [164–168], inverse engineering [169], and optimal control [170]. Various fault-tolerant approaches have also been proposed, utilizing dynamical techniques [171–175], tailored paths [176, 177], measurement-based techniques [178], or addressing coherence issues [179, 180]. In addition to all of these improvements, there are proposals to realize NAHQ in optical cavities [152, 181], nitrogen-vacancy centers [182, 183], superconducting qubits [184–187], and Rydberg atoms [188–192]. The NAHQ approach has been experimentally demonstrated in some of the same platforms as HQ mentioned above, including liquid NMR systems [193–195], trapped ions [196], nitrogen-vacancy centers [197, 198], photonics systems [199], and neutral atoms [200], in addition to several other systems including, Rydberg atoms [201], and superconducting qubits [202–205].

4.3.2 Adiabatic and Topological Quantum Computing

There are two other main GQ methods worth mentioning, so as to distinguish them all from each other. First, AQ [206] is intuitively the most related to quantum simulation, such as the work described in Ch. 3. The technique relies on finding some Hamiltonian, \hat{H}_{sol} , for which the ground state is the solution to a target problem; if

¹⁵The matrix exponential $e^{\hat{X}} = \sum_{k=0}^{\infty} \frac{1}{k!} \hat{X}^k$, is an infinite series of the powers of the matrix. For $\hat{X} = \hat{A} + \hat{B}$ where $[\hat{A}, \hat{B}] \neq 0$ this quickly becomes intractable.

one can find the ground state then they have the solution. Adiabatic QC approaches this by first preparing the ground state of a different Hamiltonian, \hat{H}_G , that is easier to realize, and then adiabatically evolving the system through the family of Hamiltonians that connect \hat{H}_G to \hat{H}_{sol} . Through the adiabatic theorem, the system will follow the ground state of the system at all times, so once the target Hamiltonian is obtained, the answer may be found through measurement of the state. This is highly reminiscent of the evolution captured by Eq. 4.11. Methods such as *quantum annealing* are derived from this idea [207]. Unlike quantum annealing, AQC more generally has been shown to be equivalent to the circuit model of quantum computing, and is therefore universal [206].

The important distinctions between AQC and HQC are twofold: first, the evolution is non-cyclic. The starting Hamiltonian is evolved adiabatically to some target, without returning. Second, the computation is (more or less) independent of the chosen initial state, and the means of evolution. The entire computation relies on identifying the Hamiltonian for which the ground state is the solution; once this is accomplished, the rest are just experimental details. As such, this differs greatly from the logic-gate based HQC, having more semblance to analog quantum computation [5].

The other GQC technique worth mentioning here is TQC, which is often confused with HQC since it also relies on non-Abelian holonomic evolution as its working principle [208–211]. The fundamental ingredient in TQC is the *anyon*, a quasi-particle that has more complicated exchange statistics than bosons and fermions. Under exchange, fermions acquire a global Abelian phase of π , while bosons obtain a phase of 0; in three spatial dimensions (3D), these are the *only* possibilities [208]. This fact relies on a topological argument: if the particles are exchanged twice, this is equivalent to one making a loop adiabatically around the other. In 3D, this loop is equivalent to that which does not encircle the other particle, and hence the phase from two consecutive exchanges must cancel. The key insight that enables particles with more complicated exchange statistics is that, in 2D, a loop that encircles a particle is no longer equivalent to one that does not. This permits more exotic phases to be possible. The anyon, therefore, is a particle that only exists in 2D, which obtains a phase different from that of bosons and fermions under exchange. Importantly, the phase need not be Abelian.

From the description of exchange as the adiabatic evolution of particles in loops around one another, geometric phase naturally emerges. For TQC, the specific interest is in non-Abelian anyons. These particles evolve according to a holonomy with a non-Abelian connection (Eq. 4.13), when making loops in 2D space. In this case, unlike in HQC, the phase depends only on the presence (or lack thereof) of

other charges within the loops¹⁶.

The motion of particles in loops around each other is referred to as *braiding*; these operations serve as the logic gates in TQC. Specifically, qubits are encoded onto the internal states of several paired non-Abelian anyons, called *fusion results*. Particles are then braided around each other, effecting holonomic transformations that, when strung together, result in a series of logic gates. It has been shown that a universal gate set is possible through these processes [211].

Topological quantum computers rely on remaining in a many-body ground state comprised of anyons, which are immune to local perturbations. Furthermore, the gate operations are also mostly immune to error since they rely only on the topology of the loops, and not the detailed geometry or dynamics [208]. As such, TQC is thought to be intrinsically fault-tolerant, to a much larger extent than other approaches to QC. There are, however, noise sources that are unique to these systems, but the error rates are as yet unknown since these systems have yet to be verifiably realized in the lab. It is thought that some of the ground states of fractional quantum Hall systems may host non-Abelian anyons, though it is unclear if they will be controllable in the manner needed for TQC, nor are all of the states thought to exhibit universal gate sets. Other quantum simulation platforms, boasting advanced control capabilities, may hold the advantage here [213].

Altogether, the GQC methods discussed here demonstrate that geometric phases have important applications outside of their description through artificial gauge fields (Sec. 4.2). This idea will serve as a primary motivation behind the work in Ch. 5 and Ch. 6, where non-Abelian geometric phases are produced in a way that circumvents the need for explicit degeneracies, or intermediate level couplings. The technique presented there may be interpreted as a means to realize single-qubit holonomic gates, as in HQC, and may be potentially extended to multi-qubit gates. As such, QC has framed the way in which results are discussed, so the ideas presented here will be revisited in later sections.

¹⁶This is due to a localization of the gauge fields on the charges. In infinite dimensional Hilbert spaces, such as states in coordinate-space, the localization of these gauge fields is guaranteed. In finite sized systems, it approximately holds [212].

CHAPTER 5

Floquet-Engineered Non-Abelian Gauge Fields

NON-ABELIAN geometric phases have interesting applications in high-energy physics, and in QIP (particularly for its role in HQC), as discussed in Sec. 4.2 and Sec. 4.3. Producing such phases, however, has traditionally been difficult [67]. In the case of adiabatic evolution, the Wilczek-Zee connection 4.1, this is in part because of the necessity for robust degeneracies. As previously mentioned, the degeneracy structure must remain consistent in parameter space, and any perturbations of this degeneracy result in non-trivial dynamical contributions to the phase. Furthermore, measurements in degenerate subspaces may be difficult, often requiring the intentional lifting of degeneracy [106].

In this chapter, we discuss a promising workaround to this issue, using *Floquet engineering* [214–216]. For periodically driven Hamiltonians, there is a representation in the *Floquet* basis in which the Hamiltonian is fully degenerate; its responses to slow changes (slow with respect to the driving frequency) through a set of tuneable parameters may result in non-Abelian geometric phases [217, 218]. Importantly, the entire Hamiltonian must be periodic for this to occur. The Floquet Hamiltonian will be derived, and the resulting geometric transformations will be explored through some numerical simulations.

Next, the robustness of these transformations will be investigated; the effects of undriven terms in the Hamiltonian will be considered. Such terms ultimately result in a breaking of the adiabatic degeneracies in the Floquet basis, leading to large dynamical contributions to the phase. These results are important to any practical implementation of such a technique, such as that covered in Ch. 6. Furthermore, these implications generalize to various other non-Abelian geometric phases, such as those used in HQC or related techniques.

5.1 Floquet Engineered Holonomies

The preliminary work in Floquet engineering for producing non-Abelian geometric phases [217] investigated fairly generic periodic Hamiltonians, with slow modulations. This was later generalized, resulting in a more complete adiabatic condition for the modulations [218]; the example of a spin in an oscillating magnetic field was studied, and authors demonstrated that non-Abelian geometric phases can be produced in this way. Following this, there was a proposal to realize Hamiltonians of this nature in ultracold neutral atoms, using a Raman Λ -scheme (Sec. 2.2.2) [219], and another related proposal for universal holonomic gates (Sec. 4.3.1) using Floquet engineering in Rydberg atom arrays [220].

Here, we will begin with the driven Hamiltonian, deriving its Floquet basis representation and the subsequent holonomy; this derivation follows closely¹ to those presented in previous work [2, 218]. Consider the Hamiltonian for a spin in a magnetic field

$$\hat{H}_0(t) = \Omega_0 \boldsymbol{\lambda}(t) \cdot \hat{\boldsymbol{F}}, \quad (5.1)$$

where $\boldsymbol{\lambda}(t)$ is a unit vector describing the time-dependent orientation of the magnetic field with effective strength Ω_0 . The properties of this Hamiltonian are similar to those in Sec. 2.2, specifically Eqs. 2.19 and 2.37. The dynamics may therefore be understood as a set of nondegenerate spins undergoing Rabi oscillations, depending on the particular orientation of the field.

Suppose now that the Hamiltonian, Eq. 5.1, is driven by a periodic envelope,

$$\hat{H}_\lambda(t) = \hat{H}_0[\boldsymbol{\lambda}(t)] \cos \omega t, \quad (5.2)$$

where we have chosen a cosine function for the driving. The drive frequency ω is referred to as the *Floquet frequency*, which will remain much larger than any rate of change to $\boldsymbol{\lambda}(t)$. The basis in which this Hamiltonian is currently represented will hereafter be referred to as the rotating frame, due to its correspondence with that in Eq. 2.19. We will go through the steps, in detail, to arrive at the effective Hamiltonian in the Floquet basis, which is described by a non-Abelian connection in the parameter space of $\boldsymbol{\lambda} \in \mathcal{M}$.

5.1.1 Micromotion

The Hamiltonian (Eq. 5.2) may be analyzed through Floquet theory. We may express the system in a new basis defined by the periodicity of the driving, in a way that is analogous to Bloch's theorem in spatially-periodic systems. Floquet-

¹Credit for the version shown here goes to Mason Protter.

engineering is used in many different systems to realize physics beyond conventional capabilities [77, 87], but a key difference here is that the *entire* Hamiltonian is driven, rather than just a few terms. We will change basis according to the unitary,

$$\hat{U}(t) = \exp \left[-\frac{i}{\hbar} \frac{\sin \omega t}{\omega} \hat{H}_0(t), \right] \quad (5.3)$$

which is colloquially referred to as the *micromotion operator*². This transformation commutes with the modulated Hamiltonian \hat{H}_0 , and $\sin \omega t / \omega$ was chosen since its derivative is equal to the Floquet drive, $\cos \omega t$. Furthermore, for every $t = n\pi / \omega$ with $n \in \mathbb{Z}$, the micromotion operator $\hat{U} = \hat{\mathbb{1}}$; the Floquet basis is therefore *stroboscopically* equivalent to the initial basis.

Similar to how we derived geometric phase in Sec. 4.1, we may write the TDSE in this basis as,

$$i\hbar \frac{\partial}{\partial t} |\psi\rangle = \left(\hat{U}^\dagger \hat{H}_\lambda \hat{U} - i\hbar \hat{U}^\dagger \frac{\partial}{\partial t} \hat{U} \right) |\psi\rangle. \quad (5.4)$$

Since the micromotion operator \hat{U} is just \hat{H}_0 exponentiated, $[\hat{H}_\lambda, \hat{U}] = 0$. To compute the term $\hat{U}^\dagger \partial_t \hat{U}$ we can use the Baker-Campbell-Hausdorff lemma,

$$e^{-\hat{A}(x)} \partial_x e^{\hat{A}(x)} = \partial_x \hat{A}(x) - \frac{1}{2!} [\hat{A}(x), \partial_x \hat{A}(x)] - \frac{1}{3!} [\hat{A}(x), [\hat{A}(x), \partial_x \hat{A}(x)]] - \dots \quad (5.5)$$

Expanding the full Hamiltonian, $\hat{H}_U = \hat{H}_\lambda - i\hbar \hat{U}^\dagger \partial_t \hat{U}$, and writing $s = \sin \omega t / \omega$ for simplicity, yields,

$$\hat{H}_U = \hat{H}_0 \cos(\omega t) - i\hbar \left\{ -\frac{i}{\hbar} \partial_t s \hat{H}_0 - \frac{1}{2!} s \left(\frac{i}{\hbar} \right)^2 [\hat{H}_0, \partial_t s \hat{H}_0] - \frac{1}{3!} s^2 \left(\frac{i}{\hbar} \right)^3 [\hat{H}_0, [\hat{H}_0, \partial_t s \hat{H}_0]] - \dots \right\}. \quad (5.6)$$

The commutator

$$[\hat{H}_0, \partial_t s \hat{H}_0] = s [\hat{H}_0, \partial_t \hat{H}_0] + \partial_t s [\hat{H}_0, \hat{H}_0], \quad (5.7)$$

where the second term is zero. Inserting this into Eq. 5.6, and performing the first

²It is called micromotion, separating the fast oscillations of the basis states from the slow changes to the Hamiltonian, akin to the motion of ions in a Paul trap [215].

derivative in the curly brackets yields,

$$\begin{aligned}
\hat{H}_U &= \hat{H}_0 \cos(\omega t) - i\hbar \left\{ -\frac{i}{\hbar} \left[\hat{H}_0 \cos(\omega t) + s \partial_t \hat{H}_0 \right] \right. \\
&\quad \left. - \frac{1}{2!} \left(\frac{is}{\hbar} \right)^2 \left[\hat{H}_0, \partial_t \hat{H}_0 \right] - \frac{1}{3!} \left(\frac{is}{\hbar} \right)^3 \left[\hat{H}_0, \left[\hat{H}_0, \partial_t \hat{H}_0 \right] \right] - \dots \right\}, \\
&= i\hbar \left\{ \frac{i}{\hbar} s \partial_t \hat{H}_0 + \frac{1}{2!} \left(\frac{is}{\hbar} \right)^2 \left[\hat{H}_0, \partial_t \hat{H}_0 \right] \right. \\
&\quad \left. + \frac{1}{3!} \left(\frac{is}{\hbar} \right)^3 \left[\hat{H}_0, \left[\hat{H}_0, \partial_t \hat{H}_0 \right] \right] - \dots \right\}. \quad (5.8)
\end{aligned}$$

Looking now at the form of the nested commutators, we have for the first one in the series,

$$\left[\hat{H}_0, \partial_t \hat{H}_0 \right] = \Omega_0^2 \left[\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}, \partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{F}} \right]. \quad (5.9)$$

We may use the identity,

$$[\mathbf{a} \cdot \hat{\mathbf{x}}, \mathbf{b} \cdot \hat{\mathbf{x}}] = a^\mu b^\nu [\hat{x}^\mu, \hat{x}^\nu]. \quad (5.10)$$

Therefore,

$$\begin{aligned}
\left[\hat{H}_0, \partial_t \hat{H}_0 \right] &= \Omega_0^2 \lambda^\mu \partial_t \lambda^\nu \left[\hat{F}^\mu, \hat{F}^\nu \right], \\
&= \Omega_0^2 \lambda^\mu \partial_t \lambda^\nu \left(i\hbar \varepsilon^{\mu\nu\alpha} \hat{F}^\alpha \right), \\
&= i\hbar \Omega_0^2 \left(\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda} \right) \cdot \hat{\mathbf{F}}, \quad (5.11)
\end{aligned}$$

where we have used the definition of the cross product, in terms of the Levi-Civita symbol $\varepsilon^{\mu\nu\alpha}$.

Similarly, we have for the second commutator in the series,

$$\begin{aligned}
\left[\hat{H}_0, \left[\hat{H}_0, \partial_t \hat{H}_0 \right] \right] &= i\hbar \Omega_0^3 \left[\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}, \left(\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda} \right) \cdot \hat{\mathbf{F}} \right], \\
&= -\hbar^2 \Omega_0^3 \left[\boldsymbol{\lambda} \times \left(\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda} \right) \right] \cdot \hat{\mathbf{F}}, \\
&= -\hbar^2 \Omega_0^3 \left[\left(\boldsymbol{\lambda} \cdot \partial_t \boldsymbol{\lambda} \right) \boldsymbol{\lambda} - \left(\boldsymbol{\lambda} \cdot \boldsymbol{\lambda} \right) \partial_t \boldsymbol{\lambda} \right] \cdot \hat{\mathbf{F}}. \quad (5.12)
\end{aligned}$$

We assume that the field described by the Hamiltonian Eq. 5.1 may only change its orientation, not its magnitude; therefore, $\boldsymbol{\lambda} \cdot \partial_t \boldsymbol{\lambda} = 0$. This, together with the fact that $\boldsymbol{\lambda}$ is of unit length, yields,

$$\left[\hat{H}_0, \left[\hat{H}_0, \partial_t \hat{H}_0 \right] \right] = \hbar^2 \Omega_0^3 \partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{F}}. \quad (5.13)$$

For the third commutator in the series,

$$\begin{aligned} \left[\hat{H}_0, \left[\hat{H}_0, \left[\hat{H}_0, \partial_t \hat{H}_0 \right] \right] \right] &= \hbar^3 \Omega_0^4 \left[\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}, \partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{F}} \right], \\ &= \hbar^3 \Omega_0^2 \left[\hat{H}_0, \partial_t \hat{H}_0 \right], \end{aligned} \quad (5.14)$$

and so the commutators form a recurrence relation. Putting these results together, we have,

$$\begin{aligned} \hat{H}_U &= i \left\{ \partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{F}} \sum_{n=0}^{\infty} \frac{(is\Omega_0)^{2n+1}}{(2n+1)!} + i(\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda}) \cdot \hat{\mathbf{F}} \sum_{n=1}^{\infty} \frac{(is\Omega_0)^{2n}}{(2n)!} \right\}, \\ &= i \left\{ i\partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{F}} \sum_{n=0}^{\infty} \frac{(-1)^n (s\Omega_0)^{2n+1}}{(2n+1)!} + i(\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda}) \cdot \hat{\mathbf{F}} \sum_{n=1}^{\infty} \frac{(-1)^n (s\Omega_0)^{2n}}{(2n)!} \right\}, \end{aligned} \quad (5.15)$$

and so, in the transformed basis the Hamiltonian is,

$$\hat{H}_U = -\sin(s\Omega_0) \partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{F}} - [\cos(s\Omega_0) - 1] (\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda}) \cdot \hat{\mathbf{F}}. \quad (5.16)$$

5.1.2 Bloch's Theorem

Since the Hamiltonian is temporally periodic, with changes to $\boldsymbol{\lambda}(t)$ slow in comparison to ω , we may assume that the states follow this periodicity [2, 217, 218]. The wavefunctions may thus be expanded in a Fourier series; this is in direct analogy to Bloch's theorem in condensed matter systems with spatial periodicity. We may decompose the wavefunctions as

$$|\psi(t)\rangle = \sum_{n=-\infty}^{\infty} e^{in\omega t} |\phi^{(n)}(t)\rangle, \quad (5.17)$$

where $n \in \mathbb{Z}$ denotes the harmonics of ω , and label states reminiscent to bands in condensed matter systems³. This constitutes an expansion of the Hilbert space into an infinite set of states, which will eventually allow us to describe the time-evolution of the system through a set of time-independent Hamiltonians acting on the new basis states. The Hamiltonian has a similar decomposition,

$$\hat{H}_U(t) = \sum_{n=-\infty}^{\infty} e^{in\omega t} \hat{\mathcal{H}}^{(n)}(t). \quad (5.18)$$

³Time is playing the role of momentum, in this analogy, where for each time t there are a family of states (bands) denoted by n .

The Fourier components are computed over a single period of the Floquet frequency,

$$\hat{\mathcal{H}}^{(n)}(t) = \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} dt' \hat{H}_U(t, t') e^{-in\omega t'}. \quad (5.19)$$

Since $\lambda(t)$ is approximately static over a single Floquet period, we may separate the time-dependence of the Hamiltonian \hat{H}_U into the fast and slow parts; the slow parts may be factored out of the integral, while the fast parts are integrated over. For this reason, we call this a *quasi-static* Fourier series [2]. We will hold off from computing these components for now, to first examine their action on the states $|\phi^{(n)}(t)\rangle$. The TDSE becomes,

$$i\hbar \sum_{n=-\infty}^{\infty} \frac{\partial}{\partial t} e^{in\omega t} |\phi^{(n)}(t)\rangle = \sum_{m,n=-\infty}^{\infty} e^{im\omega t} \hat{\mathcal{H}}^{(m)}(t) e^{in\omega t} |\phi^{(n)}(t)\rangle. \quad (5.20)$$

To reveal some more useful information about the components $\hat{\mathcal{H}}^{(n)}(t)$ we will look at the matrix elements, multiplying both sides by $\sum_{l=-\infty}^{\infty} e^{-il\omega t} \langle \phi^{(l)}(t) |$, and integrating over a single Floquet period. Starting with the left-hand side, we have,

$$\begin{aligned} i\hbar \frac{\omega}{2\pi} \sum_{l,n=-\infty}^{\infty} \int_0^{\frac{2\pi}{\omega}} dt' \langle \phi^{(l)}(t) | e^{-il\omega t'} \partial_t e^{in\omega t'} |\phi^{(n)}(t)\rangle \\ &= i\hbar \frac{\omega}{2\pi} \sum_{l,n=-\infty}^{\infty} \int_0^{\frac{2\pi}{\omega}} dt' \langle \phi^{(l)}(t) | e^{i(n-l)\omega t'} (in\omega + \partial_t) |\phi^{(n)}(t)\rangle, \\ &= i\hbar \frac{\omega}{2\pi} \sum_{l,n=-\infty}^{\infty} \langle \phi^{(l)}(t) | (in\omega + \partial_t) |\phi^{(n)}(t)\rangle \int_0^{\frac{2\pi}{\omega}} dt' e^{i(n-l)\omega t'}, \\ &= \sum_{l,n=-\infty}^{\infty} \langle \phi^{(l)}(t) | (i\hbar \partial_t - n\hbar\omega) |\phi^{(n)}(t)\rangle \delta_{l,n}, \\ &= \sum_{n=-\infty}^{\infty} \langle \phi^{(n)}(t) | (i\hbar \partial_t - n\hbar\omega) |\phi^{(n)}(t)\rangle, \end{aligned} \quad (5.21)$$

where δ_{ln} is the Kronecker delta, and the second to last step in the above uses that,

$$\int_0^{2\pi} d\theta e^{i(j-k)\theta} = 2\pi \delta_{jk}, \quad (5.22)$$

for $j, k \in \mathbb{Z}$.

Looking now at the right-hand side we have,

$$\begin{aligned}
\frac{\omega}{2\pi} \sum_{l,m,n=-\infty}^{\infty} \int_0^{\frac{2\pi}{\omega}} dt' \langle \phi^{(l)}(t) | e^{-il\omega t'} \hat{\mathcal{H}}^{(m)}(t) e^{im\omega t'} e^{in\omega t'} | \phi^{(n)}(t) \rangle, \\
= \frac{\omega}{2\pi} \sum_{l,m,n=-\infty}^{\infty} \langle \phi^{(l)}(t) | \hat{\mathcal{H}}^{(m)}(t) | \phi^{(n)}(t) \rangle \int_0^{\frac{2\pi}{\omega}} dt' e^{i(m+n-l)\omega t'}, \\
= \sum_{l,n=-\infty}^{\infty} \langle \phi^{(l)}(t) | \hat{\mathcal{H}}^{(l-n)}(t) | \phi^{(n)}(t) \rangle. \tag{5.23}
\end{aligned}$$

where in this case, the only non-vanishing terms after integration are for $m = l - n$. Putting the left- and right-hand sides together, we see that the equality must also hold element-wise, yielding,

$$i\hbar \frac{\partial}{\partial t} | \phi^{(n)}(t) \rangle = \sum_{l=-\infty}^{\infty} \left(n\hbar\omega\delta_{nl} + \hat{\mathcal{H}}^{(n-l)}(t) \right) | \phi^{(l)}(t) \rangle. \tag{5.24}$$

This reveals the structure of the Hamiltonian, Eq. 5.16, in terms of the individual Fourier components.

The operator $\hat{\mathcal{H}}^{(0)}(t)$ acts within a single band, while the element $\hat{\mathcal{H}}^{(m)}(t)$ couples bands that are m units apart [218]. The “band-structure” in this case is thus an infinite ladder of identical bands, with structure determined by $\hat{\mathcal{H}}^{(0)}(t)$, each separated in quasi-energy by $\hbar\omega$.

5.1.3 Fourier Components

Now, we may compute the explicit form of the Fourier components of the Hamiltonian, Eq. 5.16, through Eq. 5.19. We have,

$$\begin{aligned}
\hat{\mathcal{H}}^{(n)}(t) &= \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} dt' e^{-in\omega t'} \left\{ -\sin \left[\frac{\Omega_0 \sin(\omega t')}{\omega} \right] \partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{F}} \right. \\
&\quad \left. - \left(\cos \left[\frac{\Omega_0 \sin(\omega t')}{\omega} \right] - 1 \right) (\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda}) \cdot \hat{\mathbf{F}} \right\}, \\
&= - \left(\partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{F}} \right) \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} dt' \sin \left[\frac{\Omega_0 \sin(\omega t')}{\omega} \right] e^{-in\omega t'} \\
&\quad - \left[(\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda}) \cdot \hat{\mathbf{F}} \right] \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} dt' \cos \left[\frac{\Omega_0 \sin(\omega t')}{\omega} \right] e^{-in\omega t'} \\
&\quad + \left[(\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda}) \cdot \hat{\mathbf{F}} \right] \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} dt' e^{-in\omega t'}, \tag{5.25}
\end{aligned}$$

where we have factored the “slow” time-dependence out of the integral, leaving only the parts that change notably over a Floquet period. The final term in the above

is only non-vanishing for $n = 0$. For the remaining terms, we make use of the identities,

$$\frac{1}{2\pi} \int_0^{2\pi} d\theta e^{-in\theta} \sin[a \sin(\theta)] = i \left[\frac{1 - (-1)^n}{2} \right] J_n(|a|), \quad (5.26)$$

$$\frac{1}{2\pi} \int_0^{2\pi} d\theta e^{-in\theta} \cos[a \sin(\theta)] = \left[\frac{1 - (-1)^{n+1}}{2} \right] J_n(|a|), \quad (5.27)$$

where $J_n(a)$ is the n^{th} order Bessel function of the first kind. The Fourier components are therefore

$$\begin{aligned} \hat{\mathcal{H}}^{(n)}(t) = & \left\{ \delta_{n0} - \left[\frac{1 - (-1)^{n+1}}{2} \right] J_n \left(\frac{\Omega_0}{\omega} \right) \right\} (\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda}) \cdot \hat{\mathbf{F}} \\ & - i \left[\frac{1 - (-1)^n}{2} \right] J_n \left(\frac{\Omega_0}{\omega} \right) \partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{F}}. \end{aligned} \quad (5.28)$$

The odd terms are

$$\hat{\mathcal{H}}_{\text{odd}}^{(n)}(t) = -i J_n(\Omega_0/\omega) \partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{F}}, \quad (5.29)$$

while the even terms are

$$\hat{\mathcal{H}}_{\text{even}}^{(n)}(t) = [\delta_{n0} - J_n(\Omega_0/\omega)] (\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda}) \cdot \hat{\mathbf{F}}; \quad (5.30)$$

in both cases, the magnitude of the matrix elements are bounded,

$$\langle \phi^{(n)}(t) | \hat{\mathcal{H}}^{(n-m)}(t) | \phi^{(m)} \rangle \leq F |\partial_t \boldsymbol{\lambda}|, \quad (5.31)$$

since $J_n(a) \leq 1$, and $\boldsymbol{\lambda}$ has unit length. In the limit that changes to the parameters are slower than the Floquet period, the energetics within the bands are therefore more significant than the couplings between them; this is the condition for adiabaticity, written explicitly as,

$$\left| \frac{\partial \boldsymbol{\lambda}}{\partial t} \right| \ll \omega. \quad (5.32)$$

If this condition is met, then we may restrict our attention to the the 0^{th} order term,

$$\begin{aligned} \hat{\mathcal{H}}^{(0)}(t) &= \left[1 - J_0 \left(\frac{\Omega_0}{\omega} \right) \right] (\boldsymbol{\lambda} \times \partial_t \boldsymbol{\lambda}) \cdot \hat{\mathbf{F}} \\ &= \left[1 - J_0 \left(\frac{\Omega_0}{\omega} \right) \right] \partial_t \boldsymbol{\lambda} \cdot (\hat{\mathbf{F}} \times \boldsymbol{\lambda}). \end{aligned} \quad (5.33)$$

In the adiabatic limit, the Floquet Hamiltonian is therefore,

$$\hat{\mathcal{H}}_{\text{Floq.}}(t) = \partial_t \boldsymbol{\lambda} \cdot \hat{\mathbf{A}}(\boldsymbol{\lambda}), \quad (5.34)$$

with the non-Abelian connection,

$$\hat{\mathbf{A}}(\boldsymbol{\lambda}) = g \left(\hat{\mathbf{F}} \times \boldsymbol{\lambda} \right). \quad (5.35)$$

The magnitude of $\hat{\mathbf{A}}$ is $g = 1 - J_0 (\Omega_0/\omega)$. The Hamiltonian, Eq. 5.34, is *adiabatically degenerate*, since as $\partial_t \boldsymbol{\lambda} \rightarrow 0 \Rightarrow \hat{\mathcal{H}}_{\text{Floq.}} \rightarrow 0$. If slow changes to $\boldsymbol{\lambda}(t)$ are permitted, then the quasi-energy structure of the Hamiltonian is described by the connection, given in terms of the spin matrices (Eq. 5.35).

The connection, Eq. 5.35, is *not* diagonal, due to the appearance of both \hat{F}_x and \hat{F}_y in its vector components, resulting in coupling between Floquet states (depending on the loop). The Floquet basis is therefore not the eigenbasis of $\hat{\mathbf{A}}$. Furthermore, if we look at the instantaneous eigenvalues of the Hamiltonian (Eq. 5.34) we find that there is, in general, an energy splitting between the eigenstates, which may seem confusing given that we have just claimed that the Hamiltonian is degenerate. For there to be coupling in quantum systems, there *must* be some splitting in energy between the states being coupled, proportional to the coupling rate. The splitting between eigenstates is therefore a necessary condition for non-Abelian phases to couple states in the degenerate manifold, though it seems to contradict the notion of degeneracy here. We must consider that for time-dependent Hamiltonians, we may only define *quasi-energies*, hence the specification before that the Hamiltonian is *adiabatically* degenerate. In this context, the magnitude of the splitting between states is proportional to $|\partial_t \boldsymbol{\lambda}|$, which vanishes in the adiabatic limit.

To see more clearly that this is the connection, responsible for a fully geometric phase, we may write the evolution operator for this Hamiltonian (Eq. 5.34),

$$\begin{aligned} \hat{\mathcal{U}}(t) &= \mathcal{T} \exp \left\{ -\frac{i}{\hbar} \int_0^t dt' \partial_{t'} \boldsymbol{\lambda} \cdot \hat{\mathbf{A}}[\boldsymbol{\lambda}(t')] \right\}, \\ &= \mathcal{P} \exp \left[-\frac{i}{\hbar} \int_{\ell} d\boldsymbol{\lambda} \cdot \hat{\mathbf{A}}(\boldsymbol{\lambda}) \right]. \end{aligned} \quad (5.36)$$

For cyclic evolution, over a loop ℓ in the parameter space $\boldsymbol{\lambda} \in \mathcal{M}$, this becomes the holonomy (Sec. 6.2.2),

$$\hat{\Gamma}_A(\ell) = \mathcal{P} \exp \left[-\frac{i}{\hbar} \oint_{\ell} d\boldsymbol{\lambda} \cdot \hat{\mathbf{A}}(\boldsymbol{\lambda}) \right]. \quad (5.37)$$

Since the connection is given in terms of the spin matrices, it covers the $\mathfrak{su}(2)$ algebra, therefore generating transformations in $SU(2)$. This is a result of the specific form of the Hamiltonian in the rotating frame. In general, such a technique could be applied to *any* Hamiltonian. The form of the connection in other instances will similarly depend on the original Hamiltonian, and while they are not guaranteed

to host interesting or useful curvature, the application of Floquet driving to other systems is worthy of study.

5.2 Simulated Evolution

To compare the evolution of the system in the rotating basis, Eq. 5.2, with the Floquet basis, Eq. 5.34, we have performed a variety of numerical simulations in Julia [221], using a custom package for integrating the TDSE⁴. Specifically, for the Hamiltonian of interest $\hat{H}(t)$, we integrate,

$$\frac{\partial}{\partial t}\hat{U}(t) = -i\hat{H}(t)\hat{U}(t), \quad (5.38)$$

solving for the evolution operator $\hat{U}(t)$, with the initial condition that $\hat{U}(t=0) = \hat{1}$. The package provides several functions with which to use a generic ordinary differential equation solver, `OrdinaryDiffEq.jl` [222], to solve for the evolution operator in time, given a time-dependent Hamiltonian. Additionally, it includes functionality to work in any spin manifold F . Using this, we can simulate the evolution in both bases, to better understand the differences in these representations.

The Hamiltonian Eq. 6.8 is parameterized by $\boldsymbol{\lambda}(t)$, a vector that describes the orientation of the magnetic field on the unit sphere. We may therefore write $\boldsymbol{\lambda}$ in polar coordinates as,

$$\boldsymbol{\lambda}(t) = [\sin \Theta(t) \cos \Phi(t), \sin \Theta(t) \sin \Phi(t), \cos \Theta(t)]^T, \quad (5.39)$$

with $\Theta(t)$ the polar angle, and $\Phi(t)$ the azimuthal angle. The parameter space \mathcal{M} is therefore the collection of points on the unit sphere, defined through $\{\Theta, \Phi\}$, with $\Theta \in [0, \pi]$ and $\Phi \in [0, 2\pi)$. If we were to allow the amplitude of the fictitious field to change, Ω_0 , then the parameter space would include points within the sphere; an alteration to the Hamiltonian in the Floquet basis would have to accompany this addition [218].

There is a large degree of freedom in choosing the various control parameters in the Hamiltonian, Eq. 6.8, aside from choosing the loops themselves. There are three important parameters to consider setting: the amplitude of the fictitious field, Ω_0 , the Floquet frequency, ω , and the duration of the loops, T . For the duration of the loops, it is often more convenient to specify the corresponding frequency of the loop, $\Omega = 2\pi/T$, where this can be thought of as the rate of loop traversal, but also carries additional significance. The frequency Ω corresponds to the coupling rate between states in the Floquet basis, as seen by the $\partial_t \boldsymbol{\lambda}$ coefficient in the Floquet

⁴The base functions of the package were developed by Mason Protter.

Hamiltonian (Eq. 5.34) described before.

From a practical viewpoint, the amplitude of the fictitious field, Ω_0 , determines the scale for every other parameter. Based on the choice of Ω_0 , the other parameters may be chosen so that they maintain the adiabatic condition, Eq. 5.32. When this condition holds, the inter-band coupling in the Floquet basis may be neglected [2, 217, 218]. After choosing either ω or Ω , the other may be appropriately set so as to satisfy this condition.

Conceptually, Ω_0 is the Rabi frequency of the undriven Hamiltonian, controlling the rate at which spins are coupled. Together with the Floquet frequency, this controls the strength of the connection, g (Eq. 5.35), which ultimately defines the magnitude of the acquired geometric phase. Therefore, based on the choice of Ω_0 , the Floquet frequency should be set so as to obtain the desired amount of phase over the loops. For the demonstrations considered here (informed by the experiments in Ch. 6), we set $\Omega_0/\omega = 1$, yielding $g \approx 0.23$. Given this choice, Ω may be adjusted to reduce the significance of non-adiabatic corrections. The other consideration is that the spin-basis and Floquet basis are stroboscopically equivalent, that is, at times $t = 2\pi m/\omega$, $m \in \mathbb{Z}$, the bases coincide [219]. Therefore, setting Ω as the n^{th} subharmonic of ω is of practical relevance; measurements of states following these loops will thus be made simultaneously in both bases, which may reduce some experimental complexities (see Ch. 6, Sec. 6.1.2), though this is not explicitly necessary.

With these considerations in mind, there are several choices for primitive loops; here we focus on those of relevance to the experiments presented later in Ch. 6, summarized in Tab. 5.1 along with their corresponding holonomies. These loops are also shown in parameter space⁵ in Fig. 5.1. Importantly, a closed-form solution of the holonomy (Eq. 5.37) is not always evident; if the connection does not commute with itself at each point along the loop, then the path-ordering makes it difficult to obtain an analytical solution (see App. E). For this reason, the numerical methods are particularly useful. Paths that follow a geodesic, on the other hand, have readily available analytic solutions [78].

The three loops shown in Fig. 5.1(a), ℓ_1 , ℓ_2 , and ℓ_3 , all geodesics, are significant in that they produce a phase each with one of the spin matrices, as shown by their corresponding holonomies in Tab. 5.1; this demonstrates that the connection generates transformations in $SU(2)$ (although this is not proof that it is irreducible). The loop ℓ_4 , also a geodesic, demonstrates that combinations of such generators may

⁵Note that in each case, both Θ and Φ are permitted to vary over the interval $[0, 2\pi]$. Allowing Θ to vary over this extended interval permits simpler parametrizations of the loops in which the angles vary smoothly, avoiding the use of step functions that would accomplish the same task. This is especially relevant to the experimental demonstration in Ch. 6 where these parameters must be controlled by physical devices.

Loops	$\Theta(t)$	$\Phi(t)$	$\hat{\Gamma}_A(\ell)$
ℓ_1	Ωt	0	$\exp\left(-i2\pi g \hat{F}_y / \hbar\right)$
ℓ_2	Ωt	$\pi/2$	$\exp\left(i2\pi g \hat{F}_x / \hbar\right)$
ℓ_3	$\pi/2$	Ωt	$\exp\left(-i2\pi g \hat{F}_z / \hbar\right)$
ℓ_4	Ωt	$\pi/4$	$\exp\left[i\sqrt{2}\pi g \left(\hat{F}_x - \hat{F}_y\right) / \hbar\right]$
ℓ_5	$\pi/4$	Ωt	-
ℓ_6	Ωt	Ωt	-

Table 5.1: Summary of the loops ℓ implemented in experiment (Ch. 6), parameterized by $\{\Theta(t), \Phi(t)\}$, and the corresponding holonomies $\hat{\Gamma}_A(\ell)$. Loops are performed over a single period of $\Omega = \omega/10$, acquiring a phase proportional to $g = 1 - J_0(\Omega_0/\omega)$; for the ratio $\Omega_0/\omega = 1$ used here, $g \approx 0.23$. The loops 1-4 are geodesics, so the holonomies may be computed analytically (App. E). For loops 5 and 6, closed form solutions of the corresponding holonomies are not known, due to the connection not commuting with itself at the various points along the path. The geometric phases are given in terms of the spin matrices, making these holonomies members of $SU(2)$. These loops are also depicted in Fig. 5.1.

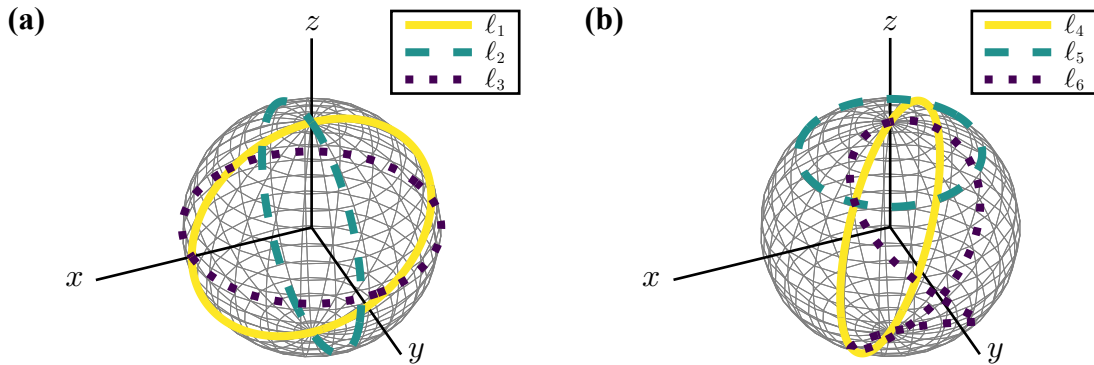


Figure 5.1: Loops in parameter space, denoting the orientation of the magnetic field in Eq. 6.8 over time, as parameterized by the polar and azimuthal angles. Example parametrizations of these loops are shown in Tab. 5.1, along with the corresponding holonomies. (a) Geodesic loops, each a great circle oriented along the various Cartesian planes. The holonomies for these loops each may be expressed by one of the three spin matrices, demonstrating how the choice of path may generate different transformations. (b) Some more complicated loops, which generate more non-trivial combinations of spin matrices.

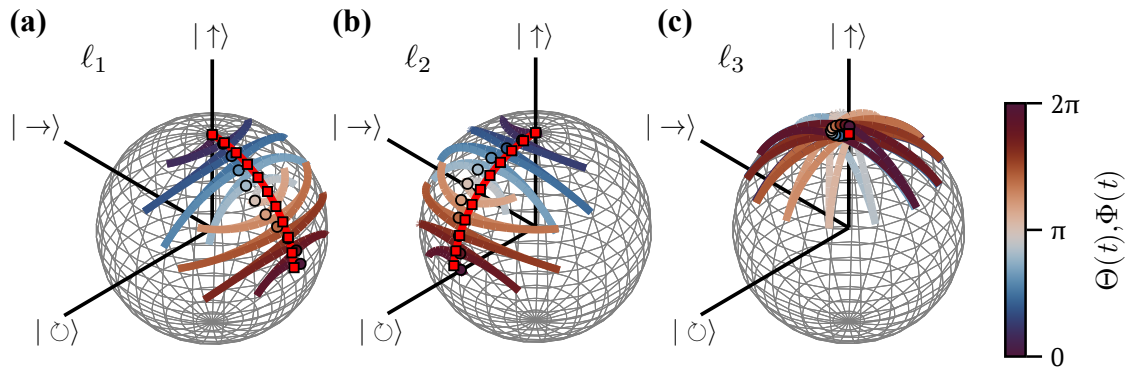


Figure 5.2: Numerical simulations comparing the evolution of spin-1/2 basis states in the Floquet basis (Eq. 5.34) and rotating basis (Eq. 5.2), shown on the Bloch sphere, for loops (a) ℓ_1 , (b) ℓ_2 , and (c) ℓ_3 . Parameters are $\Omega_0/\omega = 1$, and $\Omega = \omega/10$. The axes are indicated by the +1 eigenstates of the spin matrices: $|\rightarrow\rangle$, $|\odot\rangle$, and $|\uparrow\rangle$ for \hat{F}_x , \hat{F}_y , and \hat{F}_z respectively. State trajectories in the rotating frame are depicted by the oscillatory lines, with points (circles) sampled stroboscopically at times $t = 2\pi n/\omega$, $n \in \mathbb{Z}$; color gradient depicts time, or progress along the loop through $\Theta(t)$ for (a), (b), and $\Phi(t)$ for (c). Evolution in the Floquet basis is indicated by the solid (red) line, and markers (square) are stroboscopic samples. The markers for each basis do not perfectly coincide, due to a violation of the adiabatic condition from finite loop durations, which may be mitigated by reducing the rate of loop traversal, Ω .

also be produced. The loops ℓ_5 and ℓ_6 do not have known closed-form solutions, due to path-ordering, but ℓ_6 in particular was motivated by the experiments, for the sake of measuring the Wilson loop (see Sec. 4.1.4) trace-commutator, which will be discussed later in Sec. 5.4.

From these choices of loops, we may simulate the evolution. The simplest case, for an $F = 1/2$ manifold, the state trajectories may be displayed on the Bloch-sphere. The results for loops 1-3 in Tab. 5.1 are shown in Fig. 5.2. The axes on the Bloch sphere are indicated by the +1 eigenstates of \hat{F}_x , \hat{F}_y , and \hat{F}_z , where the initial state in each case is $|\uparrow\rangle$. The calculations in the Floquet basis result in a straight path (or in the case of ℓ_3 , no movement), while in the rotating basis the states exhibit highly oscillatory behavior. For loops ℓ_1 and ℓ_2 , the holonomy couples the eigenstates of \hat{F}_z , resulting in a superposition at the end of the loops. These trajectories take on the same basic form, but the rotations occur over orthogonal axes as indicated by the respective generators (Tab. 5.1); this ultimately results in states that differ only in relative phase. For loop ℓ_3 , the initial state is an eigenstate of the generator for this transformation, and so it remains in the initial state throughout the loop, in the Floquet basis; in the rotating basis, this is only true stroboscopically. From Fig. 5.2 we may also see the non-adiabatic error, due to the finite duration of the loops. The evolution in the rotating basis is sampled stroboscopically, along with that in the Floquet basis; these point do not coincide due to violations of the adiabatic condition (Eq. 5.32). A decreased loop rate Ω will mitigate these errors.

In order to experimentally reproduce the state trajectories depicted in Fig. 5.2, the state would need to be fully reconstructed through state tomography at each point in time (or a large sample of points). Perhaps a more convenient way of visualizing the evolution throughout the loop is to look at the populations in each eigenstate of \hat{F}_z in time; this corresponds more closely with the typical measurement capabilities, as it constitutes a projective measurement in the spin basis for times along the loop. This projection, however, will erase information of the relative phase. An example for loop ℓ_1 in $F = 1/2$ is shown in Fig. 5.3(a). Furthermore, this representation allows us to demonstrate the evolution in larger spin manifolds, which may not be displayed on the Bloch sphere. Examples for two experimentally relevant (hyperfine ground states of ^{87}Rb) manifolds are shown in Fig. 5.3(b) and Fig. 5.3(c), with $F = 1$ and $F = 2$ respectively. The former is for loop ℓ_1 , while the latter shows ℓ_2 , to demonstrate that the phase information is lost in these projections.

The results in Fig. 5.3 aid in demonstrating the $SU(2)$ symmetry of the holonomy; as the spin manifold is altered, the transformations still follow a predictable pattern. The loops used in Fig. 5.3, starting from a maximally polarized spin state $|\pm F\rangle$, result in states that are approximately eigenstates of \hat{F}_x or \hat{F}_y ; the populations at the end correspond to symmetric superpositions of the spins, each essentially equivalent to one another. This demonstrates that the spin manifold F really only changes the details, but the transformations themselves are fully understood by considering the minimal case of $F = 1/2$, as it is the lowest dimensional representation of $SU(2)$. Therefore, regardless of whichever spin manifold this technique is implemented in, the results may be easily generalized to any other spin manifold. Together, Fig. 5.2 and Fig. 5.3 (inconclusively) demonstrate the non-Abelian nature of the phase; a fully geometric phase results in the coupling of states in the manifold, which may only be accomplished through the existence of a non-Abelian connection (Sec. 4.1.3). This must be validated by a gauge-invariant measurement, however, such as the Wilson loop (Sec. 4.1.4, and Sec. 5.4).

To further demonstrate some more non-trivial loops, and our ability to numerically simulate relatively complicated time-dependent Hamiltonians, the evolution of spins through several more loops not given in Tab. 5.1 are shown in Fig. 5.4. The loops themselves are depicted in Fig. 5.4(a) along with their parametrizations; the closed form solution of the holonomy is not known for any of them. Due to these loop parametrizations containing frequency components that are larger than in all the previous simulations, non-adiabatic effects are more significant. These effects are mitigated in the simulations depicted in Fig. 5.4(b) by setting $\Omega = \omega/20$, demonstrating how the non-adiabatic error may be reduced by choice of parameters.

Together, the simulations in this section provide some insight into the connection, and the holonomic transformations it generates. If we consider the representation of

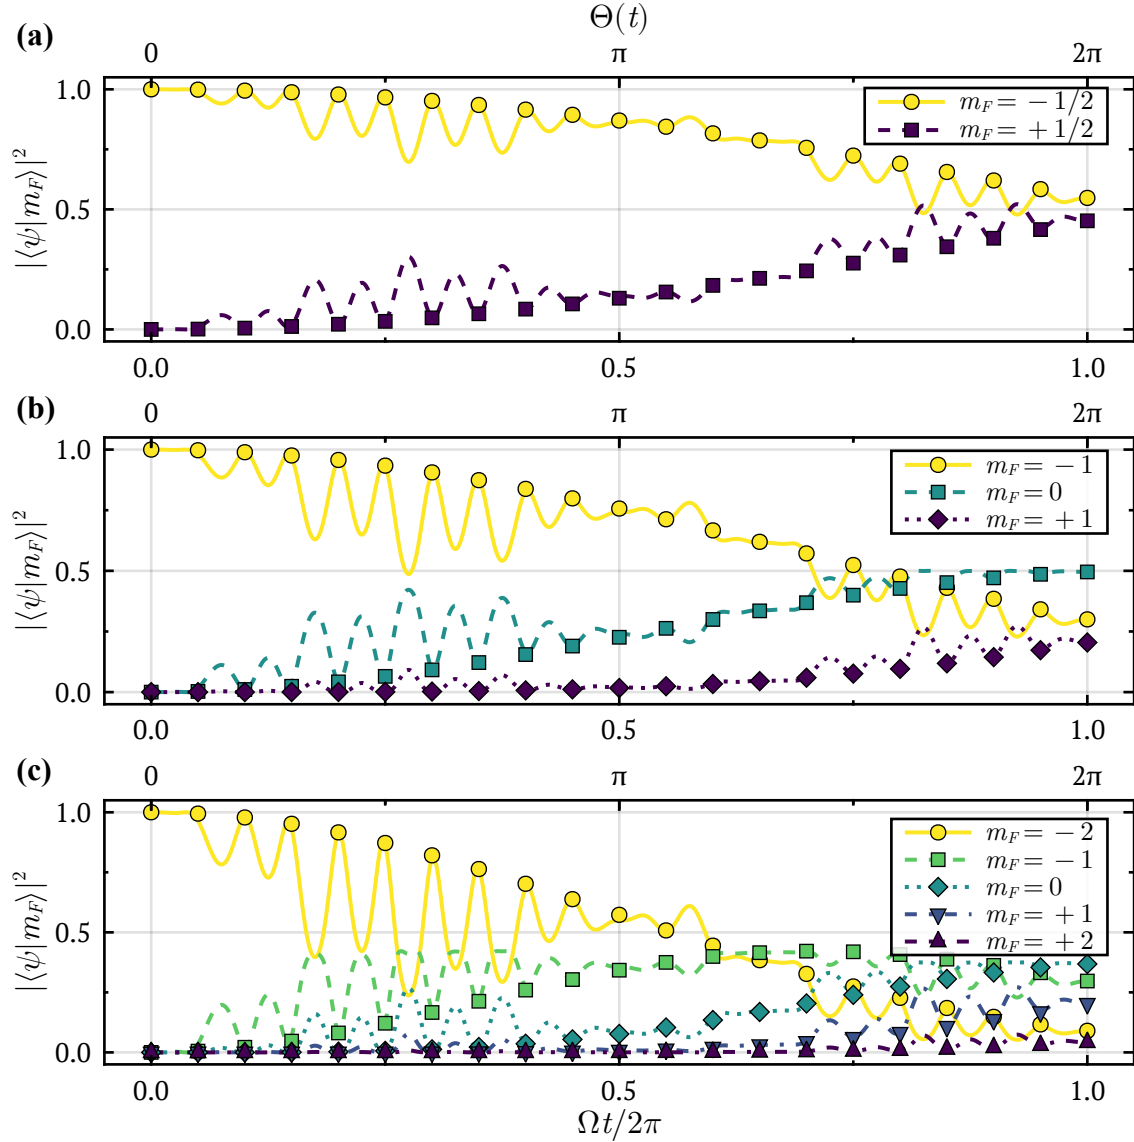


Figure 5.3: Simulated state evolution in the rotating frame and Floquet basis, shown through the relative populations in the spin eigenbasis, \hat{F}_z , for different spin manifolds: (a) $F = 1/2$, (b) $F = 1$, and (c) $F = 2$. Lines display results in the rotating frame, while points are stroboscopically sampled populations in the Floquet basis. (a) and (b) show results for loop ℓ_1 . (c) depicts ℓ_2 , demonstrating that the difference in relative phase is erased by the projections onto spin states, since these trajectories are the $F = 2$ equivalent to those in (a, b). Loop parameters are $\Omega_0/\omega = 1$ and $\Omega = \omega/10$.

the geometric phase through the curvature, as discussed in Sec. 4.2, the generated phase is given by the integrated flux through the surface enclosed by the loop. Looking at the orientation of these loops in Fig. 5.1 and Fig. 5.4, we can see that the relative orientation of this surface determines whether there is coupling between states or not. If the surface has some projection in either the xz - or yz -planes, the loop generally results in coupling between states, as the integrated flux contains components of \hat{F}_x and \hat{F}_y . Alternatively, if the surface only has a projection in the xy -plane, then it generally results in no net transfer of population. Importantly, this is only approximately true, since loops that do not follow geodesics produce non-commutative effects along the loop that deviate from this intuition; the loop ℓ_5 for instance produces a small amount of phase $\propto \hat{F}_y$ for this reason (verified numerically).

5.3 Detuning

The Floquet driving scheme presented in Sec. 5.1 is a way of taking a decidedly nondegenerate system, and through periodic modulation of the Hamiltonian, producing an equivalent representation in the Floquet basis that is adiabatically degenerate; subsequently, non-Abelian geometric phases may be permitted, depending on the form of the driven Hamiltonian. For the one considered there, the connection was given in terms of spin matrices, thus generating transformations in $SU(2)$. A key assumption behind this protocol is that the *entire* Hamiltonian is driven by a periodic function with zero time-average, over a Floquet period. Such a requirement will challenge experimenters, as imperfections in the control of quantum systems will impact the results. Motivated by the experiments in Ch. 6, here we will introduce, quite generally, an undriven term in the Hamiltonian. The form of this term in the Floquet basis, and its subsequent Fourier decomposition will reveal the impact of such control errors, which will be useful for any practical implementation of this protocol [2]. To demonstrate the impact more clearly, some additional simulations of these transformations with this source of error will be shown, for comparison with those in Sec. 5.2 without.

To start, we write the total Hamiltonian as the sum of the modulated Hamiltonian in the rotating basis, Eq. 5.2, and an additional unmodulated term,

$$\hat{H}_{\text{tot.}}(t) = \hat{H}_\lambda(t) + \hat{H}^\Delta, \quad (5.40)$$

where,

$$\hat{H}^\Delta = \mathbf{\Delta} \cdot \hat{\mathbf{F}}. \quad (5.41)$$

We assume $\mathbf{\Delta} = (\Delta_x, \Delta_y, \Delta_z)^\top$ is independent of time, or rather, that any changes to its components are much slower than the time scales we are concerned with. This term contains each spin matrix, and is therefore the most general unmodulated term (time-independent, that is) for a spin-1/2 system. For higher spin systems, unmodulated terms could be comprised of different matrices, such as the quadratic Zeeman shift which generates $SU(2F + 1)$ contributions to the phase (App. A); to compute the impact of such terms in the Floquet basis will follow the same procedure as that which follows, and generally will produce similar results.

This term is meant to resemble a systematic source of error in the control of the system, or a miscalibration. In the current basis, this term corresponds to some residual field $\mathbf{\Delta}$, shifting the time-average of the driven part from zero.

5.3.1 Micromotion

To determine the impact of an unmodulated term (detuning) on the evolution in the Floquet basis, we will go through a similar procedure as that in Sec. 5.1, starting with applying the micromotion operator (Eq. 5.3). Once again, we employ the Baker-Campbell-Hausdorff lemma (Eq. 5.5), yielding,

$$\hat{U}^\dagger \hat{H}^\Delta \hat{U} = \hat{H}^\Delta + \frac{i}{\hbar} s \left[\hat{H}_0, \hat{H}^\Delta \right] + \frac{1}{2!} \left(\frac{i}{\hbar} \right)^2 s^2 \left[\hat{H}_0, \left[\hat{H}_0, \hat{H}^\Delta \right] \right] + \dots, \quad (5.42)$$

with $s = \sin \omega t / \omega$. Looking at the commutators, the first is,

$$\begin{aligned} \left[\hat{H}_0, \hat{H}^\Delta \right] &= \Omega_0 \left[\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}, \boldsymbol{\Delta} \cdot \hat{\mathbf{F}} \right], \\ &= i\hbar\Omega_0 (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}}, \end{aligned} \quad (5.43)$$

where we have made use of the identity Eq. 5.10 once again. The second commutator is,

$$\begin{aligned} \left[\hat{H}_0, \left[\hat{H}_0, \hat{H}^\Delta \right] \right] &= i\hbar\Omega_0^2 \left[\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}, (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}} \right], \\ &= -\hbar^2\Omega_0^2 \left[(\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) \boldsymbol{\lambda} - (\boldsymbol{\lambda} \cdot \boldsymbol{\lambda}) \boldsymbol{\Delta} \right] \cdot \hat{\mathbf{F}}, \\ &= -\hbar^2\Omega_0^2 \left[(\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) \boldsymbol{\lambda} - \boldsymbol{\Delta} \right] \cdot \hat{\mathbf{F}}. \end{aligned} \quad (5.44)$$

Finally, for the third commutator,

$$\begin{aligned} \left[\hat{H}_0, \left[\hat{H}_0, \left[\hat{H}_0, \hat{H}^\Delta \right] \right] \right] &= -\hbar^2\Omega_0^3 \left[\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}, (\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) \boldsymbol{\lambda} \cdot \hat{\mathbf{F}} \right] + \hbar^2\Omega_0^3 \left[\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}, \boldsymbol{\Delta} \cdot \hat{\mathbf{F}} \right], \\ &= -\hbar^2\Omega_0^3 (\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) (\boldsymbol{\lambda} \times \boldsymbol{\lambda}) \cdot \hat{\mathbf{F}} + \hbar^2\Omega_0^2 \left[\hat{H}_0, \hat{H}^\Delta \right], \\ &= \hbar^2\Omega_0^2 \left[\hat{H}_0, \hat{H}^\Delta \right], \end{aligned} \quad (5.45)$$

where once again we have found a recurrence relation.

Putting all of the commutator results together yields,

$$\begin{aligned}
\hat{U}^\dagger \hat{H}^\Delta \hat{U} &= \hat{H}^\Delta + \frac{i}{\hbar} s \left[i\hbar\Omega_0 (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}} \right] + \frac{1}{2!} \left(\frac{is}{\hbar} \right)^2 \left\{ -\hbar^2\Omega_0^2 [(\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) \boldsymbol{\lambda} - \boldsymbol{\Delta}] \cdot \hat{\mathbf{F}} \right\} \\
&\quad + \frac{1}{3!} \left(\frac{is}{\hbar} \right)^3 \left[i\hbar^3\Omega_0^3 (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}} \right] + \dots, \\
&= \boldsymbol{\Delta} \cdot \hat{\mathbf{F}} - s\Omega_0 (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}} + \frac{1}{2!} s^2\Omega_0^2 [(\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) \boldsymbol{\lambda} - \boldsymbol{\Delta}] \cdot \hat{\mathbf{F}} \\
&\quad + \frac{1}{3!} s^3\Omega_0^3 (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}} + \dots, \\
&= \boldsymbol{\Delta} \cdot \hat{\mathbf{F}} - (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}} \sum_{n=0}^{\infty} \frac{(-1)^n}{(2n+1)!} (s\Omega_0)^{2n+1} \\
&\quad - [(\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) \boldsymbol{\lambda} - \boldsymbol{\Delta}] \cdot \hat{\mathbf{F}} \sum_{n=1}^{\infty} \frac{(-1)^n}{(2n)!} (s\Omega_0)^{2n}. \tag{5.46}
\end{aligned}$$

Therefore, in the Floquet basis the detuning term transforms to,

$$\hat{U}^\dagger \hat{H}^\Delta \hat{U} = \boldsymbol{\Delta} \cdot \hat{\mathbf{F}} - \sin(s\Omega_0) (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}} - [\cos(s\Omega_0) - 1] [(\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) \boldsymbol{\lambda} - \boldsymbol{\Delta}] \cdot \hat{\mathbf{F}}. \tag{5.47}$$

5.3.2 Fourier Components

Now we may decompose the detuning term into its Fourier components, Eq. 5.19, as we did in Sec. 5.1.3 to the modulated part of the Hamiltonian. Up until this point, it has not mattered that $\boldsymbol{\Delta}$ is time-independent. When we computed the Fourier components of the modulated Hamiltonian (in the Floquet basis) we assumed that changes to $\boldsymbol{\lambda}(t)$ were approximately static over a single Floquet period, $2\pi/\omega$. Therefore, if we wish to consider a detuning that depends on time, either a similar assumption must be made, or the explicit time-dependence must be known in order to proceed; here we are considering constant detuning, but the results would also apply to slowly varying detuning due to the former of these options.

The Fourier components are,

$$\begin{aligned}
\hat{\mathcal{H}}_{\Delta}^{(n)}(t) &= \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} dt' e^{-in\omega t'} \hat{U}^\dagger \hat{H}^\Delta \hat{U}, \\
&= \delta_{n0} (\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) (\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}) - (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}} \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} dt' e^{-in\omega t'} \sin \left[\frac{\Omega_0 \sin(\omega t')}{\omega} \right] \\
&\quad - [(\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) - \boldsymbol{\Delta}] \cdot \hat{\mathbf{F}} \frac{\omega}{2\pi} \int_0^{\frac{2\pi}{\omega}} dt' e^{-in\omega t'} \cos \left[\frac{\Omega_0 \sin(\omega t')}{\omega} \right], \\
&= \delta_{n0} (\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) (\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}) - i \left[\frac{1 - (-1)^n}{2} \right] J_n \left(\frac{\Omega_0}{\omega} \right) (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}} \\
&\quad - \left[\frac{1 - (-1)^{n+1}}{2} \right] J_n \left(\frac{\Omega_0}{\omega} \right) [(\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) - \boldsymbol{\Delta}] \cdot \hat{\mathbf{F}}, \tag{5.48}
\end{aligned}$$

where in the last step we made use of the integral identities, Eq. 5.26 and Eq. 5.27. The even components are,

$$\hat{\mathcal{H}}_{\Delta, \text{even}}^{(n)}(t) = \delta_{n0} (\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) (\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}) - J_n \left(\frac{\Omega_0}{\omega} \right) [(\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) - \boldsymbol{\Delta}] \cdot \hat{\mathbf{F}}, \tag{5.49}$$

for which the matrix elements are once again bounded,

$$\left| \langle \phi^{(n)}(t) | \hat{\mathcal{H}}_{\Delta, \text{even}}^{(n)}(t) | \phi^{(m)}(t) \rangle \right| \lesssim F |\boldsymbol{\Delta}|. \tag{5.50}$$

Similarly, the odd terms are,

$$\hat{\mathcal{H}}_{\Delta, \text{odd}}^{(n)}(t) = -i J_n \left(\frac{\Omega_0}{\omega} \right) (\boldsymbol{\lambda} \times \boldsymbol{\Delta}) \cdot \hat{\mathbf{F}}, \tag{5.51}$$

whose matrix elements are also bounded in the same way. Therefore, in the limit that the detuning is small with respect to the splitting between Floquet bands, $|\boldsymbol{\Delta}| \ll \omega$, we may once again narrow our attention to the zeroth order term. We have,

$$\hat{\mathcal{H}}_{\Delta}^{(0)}(t) = (\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) (\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}) - J_0 \left(\frac{\Omega_0}{\omega} \right) [(\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) - \boldsymbol{\Delta}] \cdot \hat{\mathbf{F}}. \tag{5.52}$$

Rearranging, and expressing in terms of the magnitude of the connection (Eq. 5.35), $g = 1 - J_0(\Omega_0/\omega)$, the detuning in the Floquet basis is,

$$\hat{\mathcal{H}}_{\text{Floq.}}^\Delta(t) = (1 - g) \boldsymbol{\Delta} \cdot \hat{\mathbf{F}} + g (\boldsymbol{\lambda} \cdot \boldsymbol{\Delta}) (\boldsymbol{\lambda} \cdot \hat{\mathbf{F}}). \tag{5.53}$$

The first term in the detuned Hamiltonian, Eq. 5.53, is fully dynamical; when integrated over time, it does not depend at all on the location on parameter space \mathcal{M} . Depending on the components of $\boldsymbol{\Delta}$, this term both breaks the degeneracy in

the Floquet band, and couples states within⁶. The second term in the detuning has a similar impact, though it does depend on the specific loop in parameter space. Importantly, however, the second term is also fully dynamical, in that the evolution operator may not be expressed through an integral⁷ over the loop ℓ . Therefore, terms in the Hamiltonian that are unmodulated by the Floquet drive correspond to dynamical terms in the Floquet basis, which break the adiabatic degeneracy, and dynamically couple states in the band. In more traditionally degenerate systems, small perturbations of the energetics break the degeneracy; here, unmodulated terms play the same role, but due to the time-dependence of the Floquet basis representation (the micromotion operator, Eq. 5.3), even static terms become time-dependent, and may result in dynamical coupling between states as well.

For the full Floquet Hamiltonian, the combination of the geometric Hamiltonian, Eq. 5.34, and the detuning, Eq. 5.53, the evolution operator takes the form of the non-adiabatic generalization of the holonomy (Eq. 4.24),

$$\hat{\Gamma}_A^\Delta(t) = \mathcal{T} \exp \left\{ -\frac{i}{\hbar} \int_0^t dt' \left[\frac{\partial \boldsymbol{\lambda}}{\partial t'} \cdot \hat{\mathbf{A}}(\boldsymbol{\lambda}) - \hat{H}_{\text{Floq.}}^\Delta(t') \right] \right\}, \quad (5.54)$$

written in the time-ordered form. The connection $\hat{\mathbf{A}}$ is therefore that of Anandan [66], which takes on the same form as the Wilczek-Zee connection [65] from before, but now since the detuning has broken the adiabatic degeneracy, it acts across subspaces.

As discussed in Sec. 4.3.1, the form of this holonomy may be deceiving; the geometric phase term is unchanged from before (Eq. 5.37), and we have simply added a dynamical phase described by $\hat{H}_{\text{Floq.}}^\Delta$. Although it may seem that the two contributions to the phase may be separated in this way, in general they do not commute. The expansion of this exponential must be expressed as a time-ordered sum of the nested commutators between each. The impact from each term on the generated transformations are therefore inextricably linked, depending on the specific path in parameter space, and the timing of the evolution. The detuning has fundamentally changed the nature of these transformations.

5.3.3 Detuned Simulations

Here we perform more numerical simulations to illustrate the impact of detuning on the holonomic evolution. To start, we may look at each of the components of $\boldsymbol{\Delta}$ independently, as applied to a particular loop. In Fig. 5.5 we show the results for

⁶The \hat{F}_z matrix is responsible for degeneracy breaking, and \hat{F}_x and \hat{F}_y for coupling.

⁷The time dependence is contained within $\boldsymbol{\lambda}(t)$; though this may therefore appear to be geometric, the absence of a $\partial_t \boldsymbol{\lambda}$ coefficient reveals its dynamical nature, since we may not trivially change the time-integral to a path-integral as we could before (Eq. 5.34).

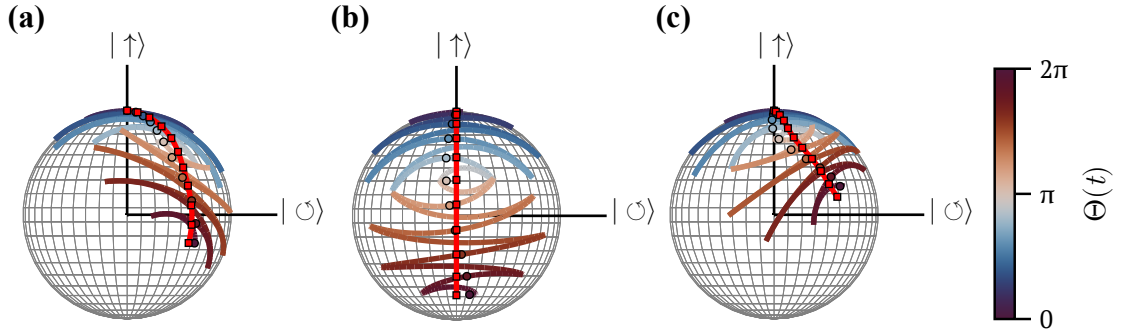


Figure 5.5: Detuned holonomy $\hat{\Gamma}_A^\Delta(\ell_1)$ shown on the Bloch sphere, with various detuning components. The viewing angle is directly along $|\odot\rangle$. State trajectories in the rotating frame are depicted by lines, with points (circles) sampled stroboscopically at times $t = 2\pi n/\omega$, $n \in \mathbb{Z}$, and color gradient that depicts time, or progress along the loop through $\Theta(t)$. Evolution in the Floquet basis is indicated by the solid (red) line, and markers (square) are stroboscopic samples. The detuning Δ is (a) $(\delta, 0, 0)^\top$, (b) $(0, \delta, 0)^\top$, and (c) $(0, 0, \delta)^\top$, where $\delta = \Omega_0/50$. Other loop parameters are $\Omega_0/\omega = 1$, and $\Omega = \omega/10$.

loop ℓ_1 on the Bloch sphere, as it contrasts with the undetuned result in Fig. 5.2(a).

Without detuning, we expect $\hat{\Gamma}_A(\ell_1)$ to rotate the state from $|\uparrow\rangle$ to approximately $|\leftarrow\rangle$; the phase is $\propto -\hat{F}_y$ thus rotating the state around the $|\odot\rangle$ axis, with a nearly $\pi/2$ pulse area. The results in Fig. 5.5(b) show $\hat{\Gamma}_A^\Delta(\ell_1)$ with a Δ_y component; from Eq. 5.53 we can see that only the first detuning term is non-zero in this case, and commutes with the geometric contribution to the phase, Eq. 5.34 for the entire loop. As such, the state still rotates around $|\odot\rangle$, but the pulse area is effectively larger due to the additional \hat{F}_y phase accumulated from detuning, resulting in a state that has “over-rotated” (gone past the target). If the sign of the detuning was inverted, the state will have “under-rotated”, as the detuning reduces the effective coupling strength instead. This example generalizes to situations in which the detuning commutes with the geometric contribution to the phase; the result is lesser or greater effective pulse area, depending on the relative sign between the dynamical and geometric contributions.

In the case that the detuning does not commute with the geometric phase, the results are more difficult to discern, but follow similar intuition. For Fig. 5.5(a), the results for $\hat{\Gamma}_A^\Delta(\ell_1)$ with a Δ_x component are shown. Similar to the case of a Δ_y component, there is additional coupling strength in this case due to the detuning, resulting in a larger effective pulse area (over-rotation); however, since the detuning no longer commutes with the geometric phase, the axis of rotation is shifted slightly towards $|\rightarrow\rangle$, deflecting the state-trajectory. Similarly, Fig. 5.5(c) has shifted the rotation axis towards $|\uparrow\rangle$, resulting in the accumulation of more relative phase. In both cases, since the dynamical phase does not commute with the geometric contribution, the path-ordering of the holonomy makes an analytic solution to $\hat{\Gamma}_A^\Delta(\ell_1)$

intractable.

The examples in Fig. 5.5 generalize to other loops, illustrating how the various detuning components impact the evolution, depending on whether they commute or not with the geometric contribution. Practically speaking, we may expect a mixture of all three detuning components simultaneously. In Fig. 5.6 we display the impact on spin populations over time, for loops ℓ_1 , ℓ_3 , and ℓ_6 . The results in both the Floquet and rotating bases are shown. In each case, 5000 detunings Δ are simulated, with each component sampled from a Gaussian distribution with mean $\mu = 0$ and standard deviation $\sigma = \Omega_0/80$. At each point in time, the interquartile range is computed from the distribution of results⁸, and is plotted as the colored bands over the results for no detuning, $\hat{\Gamma}_A(\ell)$. In addition to major deviations throughout the loop, the final populations are broadly distributed, as expected from the conclusions made from Fig. 5.5. With all three detuning components, the dynamical phase is responsible for breaking the degeneracy, causing relative phase errors between states, and coupling, resulting in both relative phase and pulse-area errors.

These simulation demonstrate the non-trivial nature of $\hat{\Gamma}_A^\Delta(\ell)$, where the interplay between the dynamical and geometric phase contributions follow some intuition, but the path-ordering makes quantitative conclusions available only through numerics. For larger spin systems, this may quickly become intractable. Furthermore, it is prudent to point out the scale of detuning considered in the simulations above, which are on the order of $\Omega_0/50$ (or equivalently $\omega/50$)⁹. A residual magnetic field significantly smaller than the driven field still results in major deviations from the desired transformations. Any practical demonstration of the Floquet engineering technique presented here must therefore very carefully isolate any residual detuning, else it will have a dramatic impact on the results.

5.4 Wilson Loops

To substantiate the claims that the connection produced by Floquet driving, Eq. 5.35, is non-Abelian, we may use the Wilson loop as outlined in Sec. 4.1.4. Specifically, we consider the trace-commutator (Eq. 4.20); if, for three different loops, this quantity is non-zero (and is thus path-dependent), this constitutes a gauge-

⁸This is the appropriate statistical metric for data that is bounded, as is the case for spin populations defined in the range $[0, 1]$. From a set of observations, the interquartile range is computed by taking the median, and separating the data into values larger than, and smaller than this value. The median of each sub-set is then computed, yielding the spread of values on either side of the combined median.

⁹In the case of Fig. 5.6, since we sampled each detuning component from a Gaussian with mean $\mu = 0$ and standard deviation $\sigma = \Omega_0/80$, the magnitude $|\Delta|$ follows a Maxwell-Boltzmann distribution with a mean of $2\sqrt{\frac{2}{\pi}}\sigma \approx \Omega_0/50$.




Figure 5.6: Numerical simulations of spin populations for the detuned holonomy $\hat{\Gamma}_A^\Delta(\ell)$, demonstrating the impact of a generalized detuning on the transformations, for loops (a) ℓ_1 , (b) ℓ_3 , and (c) ℓ_6 . The results for no detuning, $\hat{\Gamma}_A(\ell)$ are shown as the solid lines, in the Floquet basis (left) and rotating basis (right). In each case, a collection of 5000 detunings Δ were simulated, by sampling each component from a Gaussian with mean $\mu = 0$ and standard deviation $\sigma = \Omega_0/80$. For each point in time, the interquartile range of the population distributions was computed, and plotted as the colored bands. These show the spread of populations for detunings in this range. Other loop parameters are $\Omega_0/\omega = 1$, and $\Omega = \omega/10$.

invariant manifestation of the non-Abelian connection. The trace-commutator ought to be measured in experiment, but here we will provide some theoretical predictions for such measurements.

As stated in Sec. 4.1.4, the Wilson loop is not unconditionally path-dependent; the paths chosen must be different, and consist of three distinct non-commutative generators. We may demonstrate this in the particular case of $F = 1/2$, where the holonomies take the form (after integration over the loop),

$$\begin{aligned}\hat{\Gamma}_A(\ell) &= \exp\left[2i\gamma\left(\mathbf{q}\cdot\hat{\mathbf{F}}\right)/\hbar\right], \\ &= \exp\left[i\gamma\left(\mathbf{q}\cdot\hat{\boldsymbol{\sigma}}\right)\right], \\ &= \cos\gamma\hat{\mathbf{1}} + i\sin\gamma\left(\mathbf{q}\cdot\hat{\boldsymbol{\sigma}}\right),\end{aligned}\tag{5.55}$$

where \mathbf{q} is a unit vector, and γ is the amplitude of the acquired phase. Furthermore, we have expressed the operator in terms of Pauli matrices via $\hat{F}_i = \hbar\hat{\sigma}_i/2$ for $i = x, y, z$. The last line of Eq. 5.55 is true only of spin-1/2; for higher spins $F > 1/2$, there are additional terms in the expansion with various powers of the spin matrices, with $2F$ being the largest power (for the generators of $SU(2)$, higher powers than this recursively become lower ones, making this infinite series converge to a finite set of terms) [223]. As such, the result shown here is for $F = 1/2$ only, but a more general treatment may be found in *Cooke et al. 2024* [2].

For holonomies of this form (Eq. 5.55), the trace-commutator is,

$$\mathcal{W}_{abc} - \mathcal{W}_{bac} = \text{tr}\left\{\hat{\Gamma}_A(\ell_c)\left[\hat{\Gamma}_A(\ell_b), \hat{\Gamma}_A(\ell_a)\right]\right\},\tag{5.56}$$

for three loops labeled ℓ_a , ℓ_b , and ℓ_c . The commutator inside the trace is,

$$\left[\hat{\Gamma}_A(\ell_b), \hat{\Gamma}_A(\ell_a)\right] = -2i\sin\gamma_a\sin\gamma_b\left(\mathbf{q}_b\times\mathbf{q}_a\right)\cdot\hat{\boldsymbol{\sigma}}.\tag{5.57}$$

The trace-commutator is therefore,

$$\begin{aligned}\mathcal{W}_{abc} - \mathcal{W}_{bac} &= \text{tr}\left\{\left[\cos\gamma_c\hat{\mathbf{1}} + i\sin\gamma_c\left(\mathbf{q}_c\cdot\hat{\boldsymbol{\sigma}}\right)\right]\left[-2i\sin\gamma_a\sin\gamma_b\left(\mathbf{q}_b\times\mathbf{q}_a\right)\cdot\hat{\boldsymbol{\sigma}}\right]\right\}, \\ &= \text{tr}\left\{-2i\sin\gamma_a\sin\gamma_b\cos\gamma_c\left(\mathbf{q}_b\times\mathbf{q}_a\right)\cdot\hat{\boldsymbol{\sigma}}\right. \\ &\quad \left.+2i\sin\gamma_a\sin\gamma_b\sin\gamma_c\left(\mathbf{q}_c\cdot\hat{\boldsymbol{\sigma}}\right)\left[\left(\mathbf{q}_b\times\mathbf{q}_a\right)\cdot\hat{\boldsymbol{\sigma}}\right]\right\}.\end{aligned}\tag{5.58}$$

We may ignore terms that contain only a single Pauli matrix, since they are traceless. This leaves only the terms $\propto \hat{\mathbf{1}}$, yielding,

$$\mathcal{W}_{abc} - \mathcal{W}_{bac} = 4\sin\gamma_a\sin\gamma_b\sin\gamma_c\left[\mathbf{q}_c\cdot\left(\mathbf{q}_b\times\mathbf{q}_a\right)\right].\tag{5.59}$$

We can see from the form of the trace-commutator above, for it to not vanish the loops *must* generate transformations that satisfy $\mathbf{q}_c \cdot (\mathbf{q}_b \times \mathbf{q}_a) \neq 0$. Based on the holonomy Eq. 5.55 this requires each loop to generate transformations with the resulting vectors \mathbf{q} linearly independent. The loops ℓ_1 , ℓ_2 , and ℓ_3 (Tab. 5.1) meet this criteria, and fortunately the holonomies are known, (App. E). For these loops, we have that $\gamma_1 = \gamma_2 = \gamma_3 = \pi g$, and $\mathbf{q}_1 = -\mathbf{e}_y$, $\mathbf{q}_2 = \mathbf{e}_x$, and $\mathbf{q}_3 = -\mathbf{e}_z$ (using the conventional notation in atomic physics for unit vectors once again). The trace commutator becomes,

$$\mathcal{W}_{123} - \mathcal{W}_{213} = 4 \sin^3(\pi g) \neq 0, \quad (5.60)$$

hence, the Wilson loop is path-dependent for holonomies of this form, enabled by a Floquet-engineered degeneracy. This is a gauge-invariant manifestation of the non-Abelian character of the connection. For the parameters used in the simulations in Sec. 5.2, and experiments in Ch. 6, $g \approx 0.23$, so $\mathcal{W}_{123} - \mathcal{W}_{213} \approx 1.22$.

The loops 1-3 may therefore be used to verify the non-Abelian character of the connection in experiment, although from the standpoint of quantum simulation in which we take the motion in parameter space to be that of a particle in real space, it would not be possible to take these loops subsequently, because they do not share a common origin. After performing two of the loops, the position of the particle will not be located anywhere on the third loop. Another way of saying this is, at no point in parameter space \mathcal{M} do loops 1-3 intersect simultaneously; we would therefore have to move the particle some distance to then start the other loop, but such a move will impart a transformation of its own. It is for this reason that the loop ℓ_6 was studied here, and in experiments (Ch. 6); this loop produces the necessary \hat{F}_z phase, so that together with ℓ_1 and ℓ_2 it produces a non-vanishing trace-commutator, and unlike ℓ_3 shares a common origin with the other loops. Unfortunately, the connection does not commute with itself along the path, so only a numerical value for the trace-commutator may be found. For the relevant parameters, as above, we find that $\mathcal{W}_{126} - \mathcal{W}_{216} \approx 0.47$, for $F = 1/2$. Both sets of loops, either with ℓ_3 or ℓ_6 , may be used to demonstrate the non-Abelian character of the connection, but from a quantum simulation standpoint, ℓ_6 (or some other loop like it) is perhaps more *physical*.

5.4.1 Detuning

The results discussed above rely on the assumption that there is no detuning present in the holonomies, *i.e.* that the phase is purely geometric. In Sec. 4.1.4, it was pointed out that any dynamical contribution to the phase may also produce

a non-vanishing trace-commutator. The Wilson loop captures the nature of the transformation, but gives no indication as to the source of the phase being geometric or dynamical in nature. Therefore, to make any certain claims on the property of the geometric phase specifically, the dynamical components must either be strictly known, or entirely suppressed. In either case, if one wishes to measure the trace-commutator under the influence of detuning, its sensitivity to detuning must be analyzed.

Here, we provide some numerical simulations looking at the impact of each detuning component individually, for several relevant spins; the results are shown in Fig. 5.7. In all cases, the trace-commutator varies significantly over relatively small ranges of detuning. There are also regions in which the trace-commutator vanishes, meaning the transformations for these detunings are effectively Abelian. These simulations demonstrate the sensitivity of this quantity to detuning.

With dynamical phase, such as detuning here, the degree to which the Wilson loop is useful in discerning Abelian from non-Abelian gauge fields is diminished. The Wilson loop is the gauge-invariant sum of the eigenvalues of a transformation; it does not indicate whether such transformation was geometric or dynamical in origin. This issue is not specific to the Floquet-engineered holonomies considered here, but rather *any* situation in which there is a dynamical phase. Even so, the simulations shown here demonstrate yet another complication in applying Wilson loops for this purpose: their path-dependence is sensitive to detuning. Any attempted experiments must take both of these points into consideration.

5.5 Summary

In this chapter, we have investigated the use of Floquet-engineering as a means to produce adiabatically degenerate states from otherwise explicitly nondegenerate systems; the resulting manifold exhibits non-Abelian geometric phases. This is accomplished by the periodic driving of a parameterized Hamiltonian [217, 218]. In this case, we looked specifically at driving the Hamiltonian for a spin in an external magnetic field. The results were a set of holonomies that span $SU(2)$. We then demonstrated the evolution of the system under such driving in both the rotating and Floquet bases, through some numerical simulations. Overall, such a technique may be applied to any such Hamiltonian, producing similar results.

Later, we introduced an undriven term to the Hamiltonian, to study its impact on the geometric phase. We found that such terms result in a dynamical contribution to the phase that breaks the adiabatic degeneracy, and couples states in the manifold. This fundamentally changes the nature of the connection from that of Wilczek-




Figure 5.7: Numerical simulations in the Floquet basis of the trace-commutator, for loops 1-3 with detuned holonomies $\hat{\Gamma}_A^\Delta(\ell)$. Each detuning component is simulated individually, with (a) Δ_x , (b) Δ_y , and (c) Δ_z . For each, $\Omega_0/\omega = 1$. In each case, results are shown for $F = 1/2, 1, 2$, demonstrating the dependence on the particular spin-manifold. In all cases, results show large amplitude oscillations over relatively small ranges of detuning, and exhibit regions where the trace-commutator becomes arbitrarily small, within which the transformations become effectively Abelian.

Zee [65] to Anandan [66]. The impact of this *detuning* was investigated numerically.

Last, we described a means through which the non-Abelian character of the connection may be experimentally verified, through the use of Wilson loops. Three holonomies applied consecutively, in two non-cyclic permutations, may be used for this purpose. For the holonomies produced here, we computed the general form of the trace-commutator. Results indicate that for there to be a signature of a non-Abelian connection, the three chosen loops must produce linearly independent sets of spin matrices for the phases. The loops 1-3 in Tab. 5.1 meet this criterion. Next, we demonstrated through numerical simulations that such a signature of the non-Abelian connection is sensitive to detuning. Altogether, this chapter has covered a powerful technique for generating non-Abelian geometric phases, and hopefully provided some insight into such phases in practical settings.

CHAPTER 6

Periodically Driven Ultracold Ensembles

THOUGH it may seem that all of the “engineering” required to realize non-Abelian geometric phases, through the Floquet driving technique presented in Ch. 5, is done, for a practical realization there are still many missing ingredients. In this chapter we present the results of our experimental investigation of this technique [2], which is based on the previous proposal of *Chen et al.* 2020 [219], with several major modifications. In attempting to implement the initial proposal based on an optical Raman Λ -scheme [219], we encountered several unavoidable issues, which are discussed in App. D. Due to the similarities of laser- and RF-dressing, we were able to adapt the technique to instead use modulated RF coupling. This modulation scheme was then implemented in both ground state hyperfine manifolds of ultracold ^{87}Rb .

We begin in Sec. 6.1 by detailing the experimental procedure, following the production of the ultracold ensemble (see Ch. 2). We discuss the use of chirped microwave pulses to prepare states in the $F = 1$ manifold. The pulse sequences used to prepare each state are set up so that they can be iterated through in a randomized order, by synchronizing the primary control software (responsible for producing the ultracold ensembles) with the source of our microwave and RF-pulses (responsible for applying the control fields that effect holonomic transformations). We then finish this section by discussing the measurement procedures, to elucidate various pieces of information from the holonomic transformations that are the subject of study here.

Next, in Sec. 6.2 we discuss the details of how the holonomies themselves were engineered, through modulated RF-dressing. We begin with a Hamiltonian in the lab frame, with spins driven by modulated RF fields, and transform into the rotating frame. The necessary modulations to realize the periodically-driven parameterized-Hamiltonian in this basis are derived, demonstrating how such a coupling scheme

may be used to produce the necessary driven system. We then discuss some of the practical considerations to be made when selecting the various parameters, and how the corresponding RF waveforms may be produced programmatically by an [arbitrary waveform generator \(AWG\)](#). We finish this section by considering detuning, and how it may manifest itself in this coupling scheme; some techniques for calibrating the holonomies against this are described.

Our experimental results for the Floquet-engineered holonomies are then presented in [Sec. 6.3](#), with measurements compared to the numerical simulations described in [Sec. 5.2](#) and [Sec. 5.3.3](#). The time-evolution throughout loops are verified in both $F = 1$ and $F = 2$ spin manifolds, in addition to phase imparted on the states. These results demonstrate the path-dependence of these transformations, indicating the geometric, rather than dynamical nature. In each case, residual detuning produces significant error, following the theory described in [Sec. 5.3](#). We then present results for measurements of the holonomies themselves, which are characterized by the fidelity, that is, the operator inner product between measured and expected holonomies. Results are once again tarnished by detuning, but by accounting for it in the numerics we show, through the fidelity, that detuning is by and large the most significant source of error. The levels of detuning observed are too substantial to characterize the holonomies further, such as through the Wilson loop ([Sec. 4.1.4](#)).

Finally, in [Sec. 6.4](#) we discuss the results of our experiments. We start by addressing the impact of detuning, and how it may be suppressed in future implementations. We then discuss various other sources of error, including the quadratic Zeeman effect ([App. A](#)), and errors in imaging the ensembles. The experiment is then viewed through the lens of quantum simulation of artificial gauge fields, where we briefly propose some extensions of the technique to fields of other symmetries. Despite our inability to characterize the gauge fields produced here (the connection) in full, our results pave the way for this technique to be used in other systems, and our experiences may be invaluable in such efforts. We finish this section by framing the experiment in the context of [HQC](#), where the holonomies implemented here may be considered as single-qubit gate operations. Extensions of this protocol to universal [QC](#) are discussed, before concluding in [Sec. 6.5](#).

6.1 Experimental Procedure

Here we discuss the practical implementation of non-Abelian geometric phases in the ground state hyperfine manifolds of an ultracold ensemble of ^{87}Rb , enabled by the Floquet engineering scheme detailed in [Ch. 5](#). The entire experimental sequence following the [ODT](#) evaporation that produces the ultracold ensemble may be broken




Figure 6.1: Experimental pulse sequence, showing the various stages of the procedure with their respective durations, and externally applied fields; time axis is not to scale. After preparing the BEC, the bias field B_0 is turned on during a *trigger halt* step, in which the system awaits a trigger from the 60 Hz “wall” before proceeding. The bias field remains on to maintain the splitting between m_F levels until measurement. Afterwards, state preparation using microwave fields $\Omega_{\mu w}$ is performed. Next, the holonomy $\hat{\Gamma}_A(\ell)$ is applied via a modulated RF-field, Ω_{RF} ; the duration of these pulses changed over time due to alterations with the hardware, but always remained less than 1 ms. Last, atoms are released from the ODT in TOF in a magnetic field gradient ∇B , for spin-dependent projective measurements; this step may be preceded by an RF-field Ω_{RF} to change the readout basis.

down into several steps (see Fig. 6.1): state preparation, holonomic evolution, and measurement. The details of engineering the holonomies and implementing them are described in a later section, Sec. 6.2. Here, we describe the state preparation and measurement procedure in detail.

6.1.1 State Preparation

We begin by preparing the ensemble according to the techniques detailed in Ch. 2, Sec. 2.1. Following these procedures, we obtain an ultracold ensemble of ^{87}Rb with about 10^5 atoms, purely in the $|F = 2, m_F = +2\rangle$ ground state hyperfine level, held in the ODT.

Immediately following evaporation in the ODT, we apply a background bias field that splits the m_F levels in both F manifolds by 1.25 MHz, along what we define as the y -axis (Fig. 2.1). This bias field will remain present for the entire duration of the experiment, playing a role in each of the steps that come after. All of the steps in the experimental sequence that follow are sensitive to background magnetic fields, which shift the various resonances they rely on. One of the biggest sources of such fields are from the 60 Hz “wall” lines, the circuits that power almost every device in the lab. Since there is no magnetic shielding or feedback installed on our apparatus, this noise must be accounted for by triggering the entire sequence on the 60 Hz line. Therefore, while the 1.25 MHz bias field is turned on, we await a trigger from the 60 Hz line; this step thus has a variable duration, between 1-18 ms, which we refer to as the *trigger halt* step. After receiving the trigger signal, the control program continues the sequence. This protocol does *not* ensure that the field

from the various wall circuits are negligible over the course of the experiment, only that it is consistent. We may then calibrate all of our resonances, accounting for this background source in a consistent manner since we are in phase with it. Furthermore, any individual step should not exceed a few milliseconds, as resonances will drift notably over this period.

Following the 60 Hz trigger step, we may wish to prepare atoms in either the $F = 1$ or $F = 2$ hyperfine manifold. This may be accomplished by utilizing a magnetic-dipole transition between them. The splitting between manifolds is about 6.8 GHz, in the microwave regime. Depending on the polarization of the microwave field, atoms may be distributed among the various m_F levels; however, the linewidth of the transition is narrow, and so with even meager Zeeman splitting from a background bias field (about 1 MHz), the individual $|F, m_F\rangle \rightarrow |F', m'_F\rangle$ transitions may be distinguished through the frequency of the microwave field [27]. In this way, with an unpolarized microwave source, we may transfer atoms back and forth through the various magnetic sublevels by a series of pulses that differ in frequency by integer multiples of the Zeeman splitting. This permits us to prepare atoms in any of the available $|F, m_F\rangle$ states. As such, the state preparation is varied according to the specific experiment we employ, namely, depending on which particular $|F, m_F\rangle$ state we wish to apply the Floquet-engineered holonomies to. Regardless of this choice, the duration of the step is always the same, to keep the resonance calibrations during the holonomic evolution step consistent; all that changes are the particular microwave fields applied.

For the microwave source we employ a truncated wave-guide¹, which points in at an angle from the top of the cell, along the xz -plane. With respect to our bias field along the y -axis, the microwave field is comprised of a mixture of all polarization components, and may therefore be used to effect transitions from the initial m_F state to $m'_F = m_F, m_F \pm 1$. As described above, we change the frequency of the applied pulse to address a particular target transition; the atoms are therefore only resonant with a particular polarization component of the field at any given time, according to the carrier frequency. To drive these transitions, we use a microwave source which is tuned 100 MHz below resonance with the clock transition, $|F = 1, m_F = 0\rangle \rightarrow |F' = 2, m'_F = 0\rangle$. This signal is mixed with the output of an AWG, with a carrier frequency set around 100 MHz. The combination is then amplified, passed through a stub tuner, and transmitted through the waveguide towards the atoms. Altering the carrier of the AWG allows us to tune frequency as needed.

With the microwave source in place, we determined a series of pulses that prepares atoms in any of the magnetic sublevels of the $F = 1$ hyperfine manifold, with

¹A microwave waveguide that has been cut in half, allowing the directed waves to exit the guide in a fairly directional manner.

near perfect purity. In principle, a similar technique could be used to do the same for $F = 2$, but a full characterization of this manifold was deemed unnecessary². This procedure consists of a series of microwave pulses, in concert with resonant laser pulses that remove atoms remaining in a given hyperfine level after an attempted transfer. The full sequence is broken down in Fig. 6.2, as it differs for each state in the $F = 1$ manifold we wish to prepare. The microwave transfer technique used in each pulse is adiabatic rapid passage, where the frequency of the AWG is swept through resonance over 1 ms effecting near complete transfer from one state to the other; this technique is generally more robust to shifts in resonance than a typical resonant π -pulse. The intermediate laser pulses are essential to ensure the purity of prepared states, at the cost of a slightly reduced atom number. The procedure is the same for preparing both $m_F = \pm 1$ states, up until the very end where a resonant RF π -pulse is used to transfer atoms from $|m_F = +1\rangle$ to $|m_F = -1\rangle$; for this, and all other RF control described in this chapter, we use the RF-coils aligned along the x -axis (Fig. 2.1). The long duration of these steps may be attributed to opening and closing the laser shutters, which mechanically block the beams when not being used; even small amounts of leaked light are enough to completely destroy the atomic ensemble.

This state preparation sequence allows us to change between the various initial states, programmatically, via the waveforms sent to the microwave and RF sources. This may be done in an automated, and even randomized fashion by pre-programming the AWG. As such, a Python script is used to pre-generate all the necessary waveforms for a desired sequence of measurements. These waveforms are uploaded to the AWG, which then iterates through them with the experimental cycle. An ultracold ensemble is produced, state prepared, holonomy applied, and measurement performed, before switching to the next waveform uploaded, while the next ensemble is produced. This greatly reduces the amount of time each set of measurements takes, and allows us to perform scans over various parameters in a truly randomized fashion, to avoid any systematic related to the ordering of measurements. Crucially, the sequence always has the same duration, and the microwave pulses are staged such that we do not need to alter when the resonant laser pulses are applied; the software that controls the BEC apparatus may not be synchronized with the AWG in the manner required for this³. In this way, atoms may be prepared

²Preparing each individual m_F level is only necessary to reconstruct the holonomies in full (Sec. 6.3.3), and is not required for a more proof-of-concept demonstration. Furthermore, as discussed in Sec. 5.2, the results from one spin manifold are easily generalized to the others, due to the $SU(2)$ symmetry of the holonomies.

³The BEC control software allows one to set up a procedure broken down into steps. Each step has a duration, and a set of digital triggers and analogue outputs sent to various devices. We are able to perform automated scans of analogue outputs, even multiple channels at a time, but the digital channels and basic structure of the procedure are fixed. In this case, we would be required

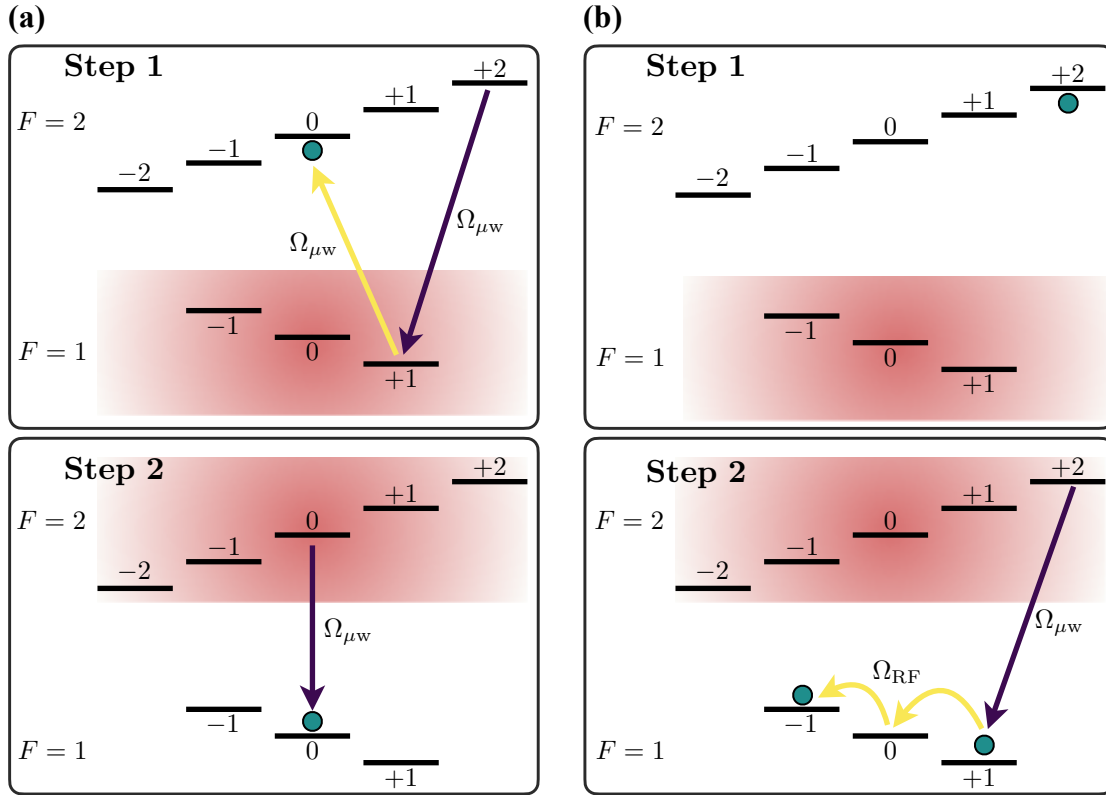


Figure 6.2: Diagram of pulse sequences for preparing atoms in any state in the $F = 1$ hyperfine manifold, starting from $|F = 2, m_F = +2\rangle$. States in both manifolds are labeled by their m_F values. Separation between manifolds is about 6.8 GHz. A bias field splits m_F levels by $\omega_Z = 1.25$ MHz, permitting specific transitions to be targeted. The atomic population is denoted by the shaded circle. Microwave pulses, $\Omega_{\mu w}$, utilize adiabatic rapid passage to transfer population between target states, with 1 ms pulse duration. Resonant intermediate laser pulses, denoted by the shaded boxes, remove atoms remaining in the indicated manifold; these pulses take 4 ms, due to the laser shutters taking 3 ms to close. The pulse sequences for each state must have the same duration, and laser pulses must occur at the same times, due to constraints in hardware control. (a) Procedure for the $|F = 1, m_F = 0\rangle$ state. (Top) two microwave pulses in series are used to transfer atoms to $|F = 2, m_F = 0\rangle$, through the $|F = 1, m_F = +1\rangle$ level. Immediately after, a laser pulse removes atoms remaining in the $F = 1$ manifold. (Bottom) atoms are then transferred to the target state by another microwave pulse, and all remaining atoms in the $F = 2$ manifold are removed by a laser pulse. (b) Pulse sequences for both $|F = 1, m_F = \pm 1\rangle$ states, which only differ in the final step. (Top) atoms left in the initial state, until after the first resonant laser pulse. (Bottom) next, they are transferred to $|F = 1, m_F = +1\rangle$ by a microwave pulse, and remaining atoms in $F = 2$ are removed by the second laser pulse. Atoms may either remain here if it is the target state, or are transferred to $|F = 1, m_F = -1\rangle$ by a resonant RF π -pulse, with a duration of about 10 μs .

in any of the magnetic sublevels in $F = 1$, or the $|F = 2, m_F = 2\rangle$ state specifically; the flexibility in the former will be used to measure the holonomies directly, while the latter is still useful in demonstrating the generalization of results to higher spin manifolds.

6.1.2 Measurement Procedure

After initializing some state and applying an **RF** control pulse to transform it (be it a holonomy, or some other calibration pulse), measurements must be performed, from which we may verify the nature of such transformations. We employed several measurement sequences, all of which relied on the **SG-TOF** technique described in Sec. 2.3.

Following the control pulse that produces a transformation of interest, the **ODT** holding the ensemble in place is turned off, allowing the ensemble to expand and fall due to gravity. In this time, a magnetic field gradient is applied, separating the atoms by their spin components. After about 20 ms, the various spin components have separated enough to be distinguished from each other in an absorption image (see Fig. 2.5). The number of atoms in each component is extracted through a Gaussian fit, so the relative populations can be determined. This is a projective measurement in the \hat{F}_z basis, where the advantage of working with an ensemble is made apparent in that we measure a statistically significant sample size in a single-shot. Furthermore, by applying an additional *readout* **RF** pulse prior to turning off the **ODT**, immediately after the transformation of interest, we are able to change the measurement basis in which the projections are made (Sec. 2.3).

To characterize the transformations fully, we require a series of measurements, but this measurement technique is destructive; after each projection, the entire experiment sequence must be repeated, including the production of a new ultracold ensemble, which altogether takes about 25 s. Measuring the time-dependence of transformations may be accomplished by interrupting the transformation of interest at various times, making a measurement without a readout pulse (looking in the \hat{F}_z basis). This allows us to ascertain how the spin populations vary over time for a given transformation. In order to infer the phase of a state following the transformation, we apply readout pulses with varied phase and pulse area, to project the state in different bases. Putting these results together allows us to fully characterize the transformed wavefunction, through state tomography (App. C).

to change digital channels with the scan, and synchronize such changes with the fields produced by the **AWG**, which is not possible with the current software.

6.2 Engineering the Floquet Hamiltonian

Section 6.1 described how to prepare the atomic ensemble in various m_F levels with a combination of microwave and RF pulses. Here we describe the details of applying various holonomic transformations to the prepared states with RF-dressing, though the modulation scheme presented here may be generalized to other similar coupling mechanisms, in cold atoms or other systems. We describe the modulation scheme that results in the driven Hamiltonian (Eq. 5.2), followed by the specific parameters and hardware used to generate such RF pulses in the lab. Together with the measurement procedure described in Sec. 6.1.2, the holonomies may be characterized.

6.2.1 Modulated RF-Dressing

In order to realize the Floquet-driven Hamiltonian in Eq. 5.2 in an ultracold ensemble, we used the RF-dressing methods discussed in Sec. 2.2.1. As with the theory description in Sec. 5.1, this modulation scheme is fairly general, and could be adapted to many other systems with similar level structure with adjacent level couplings. In the context of ultracold ensembles, the modulations derived here readily apply to a Raman M -scheme as well (Sec. 2.2.2). We will reiterate the important equations here, for convenience. The Hamiltonian in the lab frame is (Eq. 2.9) [2],

$$\hat{H}_{\text{Lab}} = \tilde{\Omega}(t) \sin [\omega_{\text{RF}}t + \tilde{\phi}(t)] \hat{F}_x + \omega_Z \hat{F}_z, \quad (6.1)$$

where a bias magnetic field is applied resulting in the Zeeman splitting ω_Z , therefore lifting the degeneracy between spin levels. An oscillating magnetic field is then applied, through a set of RF-coils (Sec. 2.1.3), resulting in the coupling above. Both the amplitude $\tilde{\Omega}(t)$ and phase $\tilde{\phi}(t)$ of the driving field are modulated in time.

In the frame rotating at $\omega_{\text{RF}} + \partial_t \gamma$ (Sec. 2.2.1), the Hamiltonian becomes (Eq. 2.19),

$$\hat{H}_{\text{Rot.}}(t) = \frac{\tilde{\Omega}(t)}{2} \left\{ \sin [\tilde{\phi}(t) - \gamma(t)] \hat{F}_x - \cos [\tilde{\phi}(t) - \gamma(t)] \hat{F}_y \right\} + \delta(t) \hat{F}_z, \quad (6.2)$$

where $\delta(t) = \omega_Z - \omega_{\text{RF}} - \partial_t \gamma(t)$, and $\gamma(t)$ is an arbitrary phase associated with the transformation to the rotating frame; here, it will help us derive the form of phase modulations needed to match the target Hamiltonian, Eq. 5.2.

Looking at the target Hamiltonian, Eq. 5.2, we can determine what the various parameters need to be in order to realize this Floquet-engineered system. First, the amplitude of the driving RF field should be,

$$\tilde{\Omega}(t) = 2\Omega_0 \sin [\Theta(t)] \cos \omega t, \quad (6.3)$$

where ω is the Floquet frequency, Ω_0 is the amplitude of a *fictitious* magnetic field acting on the spins⁴, and $\Theta(t)$ is the polar angle for the orientation of the fictitious field. Looking now at the \hat{F}_z term, we require

$$\delta(t) = \Omega_0 \cos [\Theta(t)] \cos \omega t. \quad (6.4)$$

Ideally, the driving RF field should be on resonance with the Zeeman level splitting, $\omega_Z - \omega_{\text{RF}} = 0$; if this condition is not met, there will be major dynamical consequences. Such effects would be described by a z -component of the detuning, discussed in Sec. 5.3. Assuming we are able to meet the resonance condition, this leaves

$$\partial_t \gamma(t) = -\Omega_0 \cos [\Theta(t)] \cos \omega t, \quad (6.5)$$

$$\Rightarrow \gamma(t) = -\Omega_0 \int_0^t dt' \cos [\Theta(t')] \cos \omega t'. \quad (6.6)$$

The RF driving phase must therefore be

$$\tilde{\phi}(t) = \Phi(t) + \gamma(t) + \frac{\pi}{2}, \quad (6.7)$$

where $\Phi(t)$ is the azimuthal angle of the fictitious field. The extra $\pi/2$ term must be included to shift the phases so that we have the correct signs in front of \hat{F}_x and \hat{F}_y . We therefore obtain the target Hamiltonian,

$$\hat{H}_\lambda(t) = \Omega_0 \boldsymbol{\lambda}(t) \cdot \hat{\mathbf{F}} \cos \omega t, \quad (6.8)$$

where,

$$\boldsymbol{\lambda}(t) = [\sin \Theta(t) \cos \Phi(t), \sin \Theta(t) \sin \Phi(t), \cos \Theta(t)]^\top. \quad (6.9)$$

Loops in parameter space are therefore controlled through the phase of the RF-carrier, $\Phi(t)$, and a parameter relating to both the amplitude and frequency modulations of the carrier, $\Theta(t)$.

From the form of $\tilde{\phi}(t)$ and $\delta(t)$, we can see that $\gamma(t)$ is accounting for the frequency modulations in the basis transformation. The instantaneous frequency of a wave is given by the derivative of its phase; therefore, to modulate the frequency of a field, one simply needs to modulate the phase as the integral of the desired modulation, as in Eq. 6.6. However, when picking a rotating basis we are left with

⁴The Hamiltonian Eq. 5.1 is that of a spin in a magnetic field $\boldsymbol{\Omega} = \Omega_0 \boldsymbol{\lambda}$, where the spins undergo Larmor precession around $\boldsymbol{\Omega}$. In our case, the field is artificial, as it results from representing the *real* fields from the Lab frame in the rotating frame. For this reason we refer to it as a fictitious field. This is only the first such artificial field here, as the Floquet representation produces yet another, different, artificial magnetic field: the curvature.

a choice: to rotate with the carrier frequency alone, in this case ω_{RF} , or to include the frequency modulations. The choice determines whether the frequency modulations appear in the detuning term, $\delta(t)$, or in the phase of the couplings; both are equivalent. In this particular case, to obtain the target Hamiltonian, we want the frequency modulations to appear in the detuning term; going through these steps has let us identify the correct basis for this to be the case.

Based on this **RF** modulation scheme, which as stated above may be readily applied to other systems with similar level structures and control capabilities, we are able to implement the holonomies discussed in Sec. 5.1. The form of the amplitude modulations $\tilde{\Omega}(t)$ (Eq. 6.3) and phase modulations (Eq. 6.7) prescribe how to generate the **RF** waveform necessary to effect an arbitrary loop ℓ in parameter space, parameterized through the angles $\{\Theta(t), \Phi(t)\}$. The waveform may be programmed into an **AWG**, directly computed from these relations. Next, we discuss the specific implementation of this technique in each of the ground state hyperfine manifolds of an ultracold ensemble of ^{87}Rb .

6.2.2 Holonomies

Following the state preparation step, we are able to implement the holonomies described in Sec. 5.1, Tab. 5.1, through the **RF**-dressing technique introduced above. The modulation scheme allows us to take any desired loop, parameterized through the parameters $\Theta(t)$ and $\Phi(t)$, and generate the corresponding **RF** waveform. As discussed in Sec. 5.2, there are several considerations to be made when choosing the holonomy parameters. In our case, we are primarily limited by Ω_0 , the effective Rabi frequency of the **RF**-driving. The relation between this and the Floquet frequency ω sets the magnitude of the geometric phase, through $g = 1 - J_0(\Omega_0/\omega)$, so while there is some degree of tunability in ω , it has an impact on the acquired phase. We found that setting $\Omega_0 = \omega$ produced transformations with close to $\pi/2$ pulse area, which is favorable for the application of this technique to **HQC** (Sec. 4.3.1). As such, in order to still satisfy the adiabatic condition (Eq. 5.32), the loop rate $\Omega \ll \Omega_0, \omega$. Therefore, the smaller Ω_0 is, the longer the loop durations are; at some point, drifts in background fields due to the 60 Hz line, and eventually even spin decoherence from collisions in the ultracold ensemble, or scattering losses from the **ODT** will become relevant. These errors are mitigated by making Ω_0 as large as possible, which is done by increasing the amplitude of the **RF** field.

Given this constraint in Ω_0 , we are left with setting Ω to be small enough to maintain adiabaticity, while avoiding unnecessarily long evolution times that build up large error. As described in Sec. 5.2, it is prudent to set Ω as a subharmonic of ω , so that measurements at the end of the holonomy are done directly in the Floquet

basis, due to its stroboscopic equivalence with the rotating frame [219]. As such, we chose $\Omega = \omega/10 = \Omega_0/10$; in our case, we found this to be sufficiently adiabatic, but this could easily be decreased if non-adiabatic corrections became important, at the expense of other errors potentially growing.

RF Waveforms

Due to the required coordination the microwave pulses used in state preparation and the RF pulses used to produce the holonomies and readout pulses, it was necessary to generate both waveforms with the same AWG. A single Python script was used to generate all the necessary waveforms, and coordinate the timings between them to produce the correct initial states, and apply the desired holonomies to them. Specifically, we used a Tektronix AWG5204, which has a sampling rate of 6 GSa/s; this is more than sufficient to realize the fastest frequency in the waveforms, $\omega_{\text{RF}} = 1.25$ MHz, in addition to the ≈ 100 MHz microwave carrier (mixed with the 6.8 GHz source). The AWG has four output channels, and two independent triggers, which permitted all the required operations. The device may be interfaced with over an Ethernet connection, allowing remote operation in tandem with the other control software⁵. Some example waveforms are shown in Fig. 6.3, for different loops.

The output of the AWG is limited to 750 V_{pp}, which is insufficient to generate a large enough field in the RF coils to perform the experiment within the coherence time of the ensemble. The output of the AWG was therefore sent through a 1.5 W amplifier⁶, before being transmitted to the coils. Additionally, a custom circuit was made⁷ to better match the impedance of this long transmission line (about 15 ft) with the RF-coils, resulting in more efficient use of the power (stronger fields). This produced sufficiently high RF Rabi frequencies to implement the desired holonomies well within the coherence time of the ensemble⁸, and the 60 Hz wall-circuit drift (Sec. 6.1.1). Together, these devices allowed us to obtain Rabi frequencies on the order of 10 kHz, making the loop durations about 1 ms.

The value of Ω_0 was measured by observing Rabi oscillations in the spins without any Floquet modulations, such as those simulated in Fig. 2.3(a); the trend in spin populations over time could be fit for the value of Ω_0 , with more oscillations leading to a more precise value. We found Ω_0 to be remarkably consistent, though

⁵Specifically, Joseph Lindon wrote a Python package to handle the serial communication protocols needed, in order to upload waveforms and cycle between them in a way that was coordinated with the software that controls the BEC production, and all other devices.

⁶Mini-circuits model ZHL-72A+.

⁷Design and fabrication of this device is credited to Joseph Lindon.

⁸An ≈ 50 W amplifier was installed at some point to increase this even further, but not only did it distort the RF-waveforms notably, the various cables radiated significant RF power throughout the lab, ultimately crippling the BEC production. The high RF power also produced AC Stark shifts, making atomic resonances unstable.




Figure 6.3: Sample RF-waveforms which may be loaded into an AWG, to produce Floquet-driven holonomic transformations of atomic spins. Both waveforms shown are parameterized by $\Theta(t) = \Omega t$, but differing values of $\Phi(t)$; they would produce the loops ℓ_1 and ℓ_2 depicted in Fig. 5.1(a). In this case the waveforms are $V(t) = \sin(\Theta) \cos(\omega t) \sin(\omega_{\text{RF}} t + \varphi)$, where $\varphi(t)$ is a function of both Θ and Φ (Eq. 6.7). (a) The RF-carrier frequency is 1.25 MHz, making the individual oscillations difficult to resolve, but the structure of the modulations can be discerned. The slow frequency envelope is a single period of Ω , while the faster envelopes are due to the Floquet drive $\omega = 10\Omega$. (b) Looking closer at the central zero-crossing, we can see that the waveforms differ only by a constant phase of the driving RF carrier.

it changed over several maintenance cycles of the apparatus⁹. It was, therefore, carefully calibrated often, especially between any maintenance cycles of the system or other such alterations.

Calibrating Against Detuning

In Sec. 5.3 we introduced a generic undriven term in the rotating basis, that resulted in significant dynamical contributions to the phase in the Floquet basis; this *detuning* broke the established adiabatic degeneracy, and coupled Floquet bands dynamically. In the rotating frame (Eq. 5.2), this term corresponds a background magnetic field Δ . The Floquet driving ideally results in a magnetic field that is zero on average over many Floquet periods, but since the detuning is undriven it shifts the average from zero, resulting in dynamical effects on top of the desired geometric ones.

For the experiments presented here, the magnetic field in the rotating frame is fictitious. The *real* fields are produced in the lab frame, with a static bias field and fast RF carrier. This basis is transformed under the rotating wave approximation (see Sec. 2.2.1) to rotate with the fast oscillations; this yields the Hamiltonian in the rotating frame, Eq. 6.8, a basis representation in which the field is roughly static. As such, the detuning terms in this basis are not necessarily caused by a static external field. Due to the time-dependence of the rotating basis with respect to the lab frame, the detuning terms are caused by more complicated fields.

In the current basis, the z -component of the detuning, Δ_z , corresponds to a mismatch between the level splitting and the frequency of the driving field. The modulations of the \hat{F}_z terms are ideally centered on resonance with the transition, but this unmodulated term shifts that drive off resonance; this was the motivation for calling the Δ term detuning, as it is literally a detuning of the RF carrier from atomic resonance (the z -component of Δ , specifically). This term could arise through a miscalibration of the resonance, or an external field outside of the experimenter's control that perturbs the spin state energies. The other components of the detuning, Δ_x and Δ_y , must originate from additional RF fields in the lab frame. A potential source could be from leaked control fields, or external oscillating-fields near resonance; generally, fields of this nature would produce all three detuning components since they are not likely to be directly on resonance, thus producing a Δ_z component as well.

In the experiments discussed here, we have excellent extinction of unwanted RF fields; there are multiple signals that must be dispatched to drive the various RF coils, including triggers to RF switches and mixers, and for gating amplifiers.

⁹Between ODT alignments, for instance, as the location of the ensemble within the vacuum cell changes with this.

Additionally, the RF carrier signals themselves are produced by digital devices, such as our AWG, which are also gated, and produce fields with pre-programmed durations. Altogether, we detected no leaked control fields from any of our sources. This was verified by observing the stability of spin states when undriven; if there were external RF fields near resonance, we would expect to see background-induced Rabi oscillations. No evidence of this was observed over durations much longer than the 1 ms holonomies. Additionally, we verified that the Rabi oscillations under RF driving were also consistent with expectations. This was done by applying Ramsey-like sequences in which multiple RF pulses with varied phase were applied, separated by a variable delay. Phase measurements (Sec. 6.1.2) were done to verify that the populations and relative phase match that expected for such pulses, with delay times on the order of 1 ms. Based on these calibrations, the only component of the detuning Δ considered in the experiments that follow are due to miscalibrations or drifts in the RF resonance, Δ_z .

The carrier frequency of the RF waveforms was fixed at 1.25 MHz throughout the experiments. As such, the resonance condition was controlled by adjusting the applied bias field, along the y -axis. Drifts in the RF resonance contributed the largest uncertainty in the results that follow, as quantified in Sec. 5.3. Prior to each set of measurements, the resonance was calibrated by observing a single RF π -pulse, applied to the $|F = 1, m_F = +1\rangle$ state; the bias field was manually tuned so that there was no remaining spin population in the initial state, with all of the atoms having been transferred to $|F = 1, m_F = -1\rangle$. This method served as only a rough calibration of the resonance since a π -pulse has a duration of less than 50 μ s, in comparison to the 1 ms holonomy duration. As such, after setting the field in this way, we applied the holonomy $\hat{\Gamma}_A(\ell_1)$ in Table 5.1, tuning the bias further until the expected populations were observed. We found that the multi-spectral nature of the modulated RF pulses produced results that are much more sensitive to detuning from resonance than standard unmodulated RF (as in the rough calibration step), even for the same pulse durations; this is demonstrated theoretically by the simulations shown in Fig. 6.4. A comparison between $\hat{\Gamma}_A^\Delta(\ell_1)$ and a standard RF-pulse of the same duration with detuning error-bands are shown in Fig. 6.4(a); the bands are barely visible on the unmodulated pulse, but significant on the holonomy. This was investigated further by plotting the fidelity (the operator inner product, see Sec. 6.3.3 for further details) of the detuned operators versus the undetuned; the fidelity of the holonomies falls much faster.

Attempts at using longer standard RF pulses with odd-integer pulse areas did not result in calibrations with similar fidelity. We believe that, in addition to the lack of sensitivity shown in Fig. 6.4, the cause of this method not working as well is likely due to the heating of the RF amplifier over the duration of the pulse. Unlike

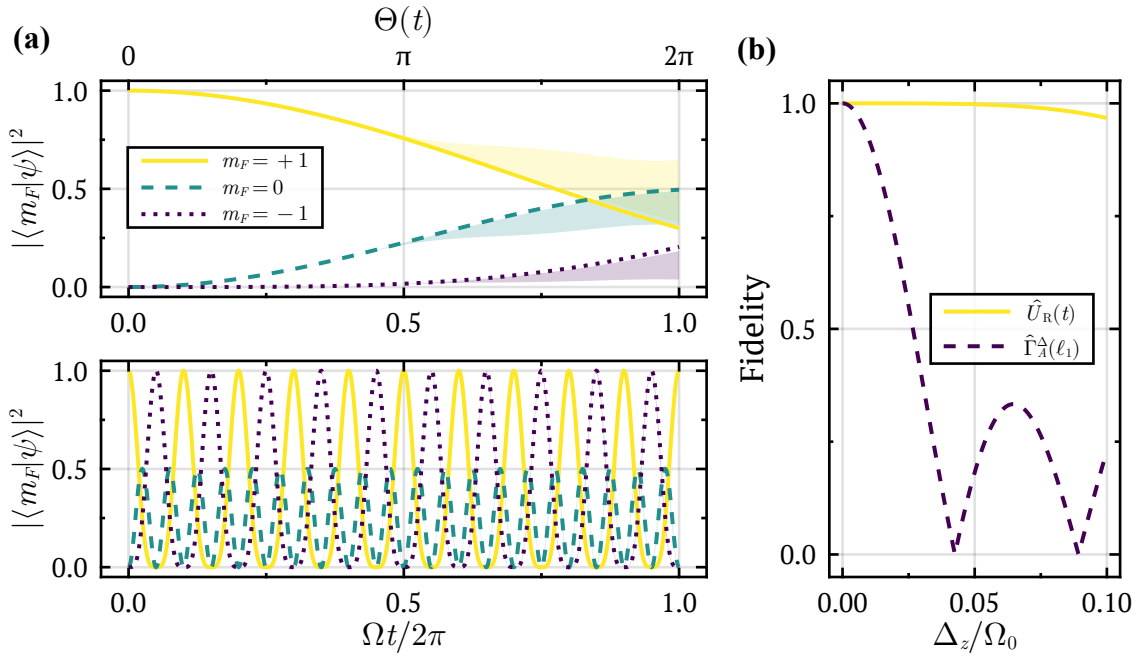


Figure 6.4: Numerical simulations comparing the detuning sensitivity of the holonomy $\hat{\Gamma}_A^\Delta(\ell_1)$ to an unmodulated RF-pulse \hat{U}_R of the same duration. (a) Simulations of the populations over time for the holonomy in the Floquet basis (top) and unmodulated pulse in the rotating basis (bottom). The pulse durations are a single period of $\Omega = \Omega_0/10$. For the holonomy, $\Omega_0/\omega = 1$, and for the unmodulated pulse, $\omega = 0$, $\Theta(t) = \pi/2$, $\Phi(t) = 0$. Detuning bands for both show the interquartile range of results from 5000 detunings sampled from a Gaussian with a mean of zero and standard deviation $\Omega_0/20$; bands are barely visible for the unmodulated pulse, but significant for the holonomy. (b) Fidelity (Eq. 6.13) of operators as a function of detuning for holonomy $\hat{\Gamma}_A^\Delta(\ell_1)$ (dashed) and unmodulated RF-pulse \hat{U}_R (solid), demonstrating the drastic difference in sensitivity between transformations.

with the holonomies, the RF carrier is at full amplitude for the entire pulse in this case. This likely caused heating in the amplifier, which changes the field amplitude over the pulse duration, causing inconsistent timing of the population transfer. As such, due to their sensitivity to detuning, we found that by setting the bias field based on the holonomy $\hat{\Gamma}_A(\ell_1)$, this produced an excellent calibration for resonance.

To produce a Zeeman splitting of 1.25 MHz in ^{87}Rb , a bias field of 1.79 G is required; this corresponds to a current of about 4 A in our y -bias coil. For the Rabi frequency Ω_0 (or equivalently, the Floquet frequency ω) used in experiment, on the order of 10 kHz, a detuning of 0.2 kHz would produce dynamical errors equivalent to those depicted in Fig. 5.5 and Fig. 5.6, which are significant. This corresponds to a field of 0.3 mG, or less than 1 mA of current in the bias coil. Therefore, not only is it difficult to detect such a detuning through the techniques available to us, but even controlling the resonance to this tolerance is challenging. The bias coil supply in our system¹⁰ operates in constant-current mode, with an internal stability reported to be < 0.5 mA, which is close to this level of sensitivity. We therefore expect the detuning to drift between each measurement by about 0.2 kHz, with other sources potentially producing even larger detunings; the detuning seen in measurements (Sec. 6.3) match these expectations, as they are all under about 0.8 kHz.

6.3 Results

There are three main ways in which we characterized the holonomies in Table 5.1: by observing their time-evolution, looking at the relative phase imparted by the transformations as it depends on the loop, and finally, by measuring the *fidelity* of the applied holonomies. The latter of these required a measurement of the holonomies in full. In each case, results were compared to numerical simulations (Sec. 5.2), performed in the lab frame¹¹.

Each of the parameters discussed above (Sec. 6.2) was set in a Python script, which generated all of the waveforms for the AWG (Sec. 6.2.2), including the state-preparation microwave pulses (Sec. 6.1.1). The measurement sequences were pre-programmed, with the ordering often randomized to avoid any bias. Each measurement sequence, hereafter referred to as a *scan*, therefore consisted of a pre-set combination of holonomies and readout RF pulses, applied to either the same initial spin state, or several. Aside from the shot-to-shot drift in detuning, likely caused by the bias supply itself, we found the resonance to be stable for periods of time > 10 minutes, with some exceptions. As such, the resonance calibration discussed

¹⁰KEPCO BOP 20-20M.

¹¹The lab frame was chosen so that the same RF waveform uploaded to the AWG could be integrated to predict results. This kept proper track of the phase throughout experiments.

in Sec. 6.2.2 was implemented before and after every scan; this way, we were able to detect any noticeable changes to resonance that may have occurred over a single scan, and re-adjust before the next. The scans were typically kept to less than 20 operations, to avoid letting the detuning vary too much between measurements.

6.3.1 Time Evolution

Observing the time-evolution of the spin populations, in the \hat{F}_z basis, can be done by applying a holonomy, and interrupting it at various times with a measurement. This was our initial method of verification for several reasons: first, the bare-spin basis is insensitive to the relative phase. If there are any errors in the phase of the generated RF waveforms, or in our calibrations of the readout RF pulses, they have no impact here. Second, this allows us to verify the stroboscopic equivalence between the Floquet basis (Eq. 5.34) and the rotating basis (Eq. 6.8). Observing the high-frequency oscillations provides strong evidence that the pulses we apply are correct, aside from the relative phase imparted. Last, and perhaps most importantly, this method allows us to observe the build up of detuning error over time. For short durations, the detuning error is negligible, but as the measurements get closer to the end of the loop, noticeable deviations from the ideal undetuned holonomy arise. Detuning has a particularly noticeable effect on the relative phases, as it is essentially an error in choice of rotating frame, leading to a different rate of phase accumulation than expected. Therefore, since the time scans are insensitive to phase, the impact of detuning is smaller.

Despite these advantages, the number of measurements required to fully characterize the time evolution is significant, due to the high frequency oscillations requiring many samples. As such, these scans were either taken with an insufficient number of points, or have many taken over a long period of time in which the detuning may vary dramatically. Neither case is ideal.

Here we include several examples for both $F = 1$ (Fig. 6.5), and $F = 2$ (Fig. 6.6). In both cases, the data was fit for a particular detuning, Δ_z , with this as the only free parameter; the lines are therefore the detuned holonomy (Eq. 5.54), computed with this value of detuning, $\hat{\Gamma}_A^\Delta(\ell)$. On top of this model, we also compute an estimate for the error in this detuning. We numerically sampled 1000 detunings from a Gaussian distribution with a mean of Δ_z , and a standard deviation¹² of $2\pi \times 0.4$ kHz. The time evolution of the spin populations was then computed for each of these sampled detunings, resulting in a set of trajectories for each spin population; the spread of these results is then shown by computing the interquartile range.

¹²This particular choice for the standard deviation was determined from the distribution of all measured detunings, and roughly coincides with the expected stability of our bias field supply.

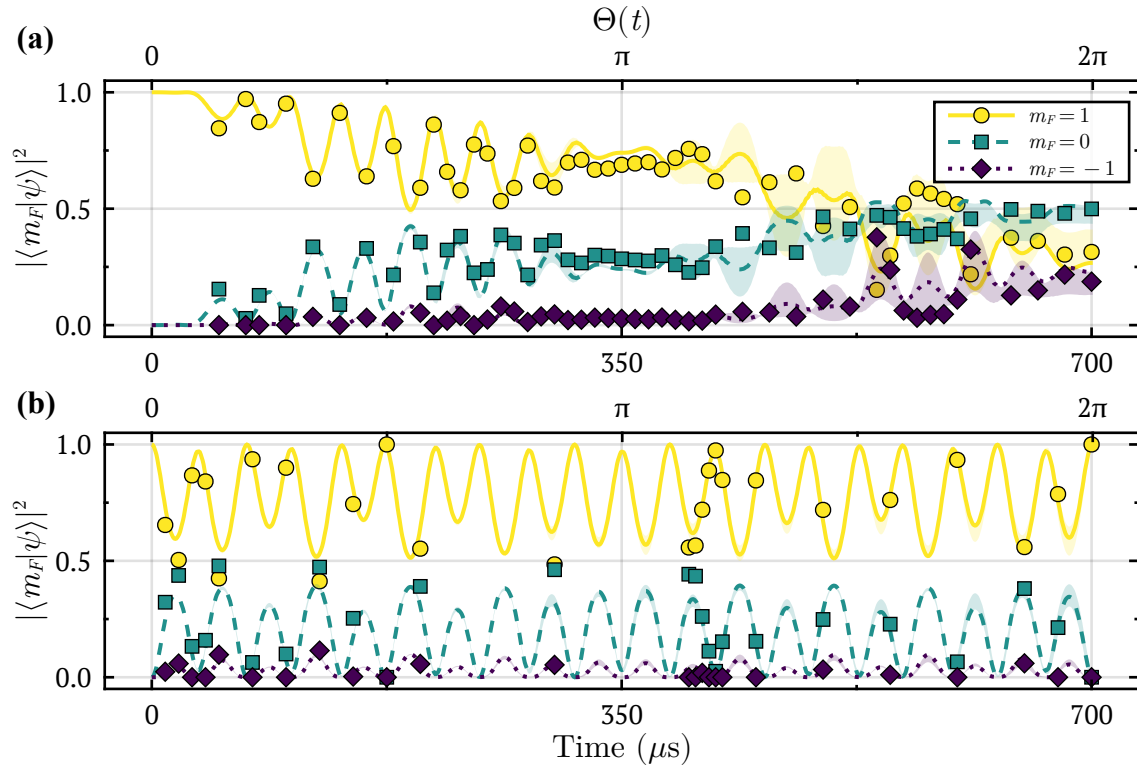


Figure 6.5: Measurements of spin populations (points) on top of numerical calculations (lines) as they vary in time throughout a loop, in $F = 1$. Theory lines were computed by fitting the population data for a value of detuning, Δ_z , which was then used to compute the detuned holonomy from Eq. 5.54. For each measurement, $\Omega_0/\omega = 1$, with $\Omega_0/2\pi = 14.27$ kHz, and $\Omega = \omega/10$. Colored bands display the interquartile range of populations subject to detuning, numerically computed from a Gaussian sample of detunings with mean Δ_z and standard deviation $2\pi \times 0.4$ kHz. Results are shown for ℓ_1 (a) and ℓ_3 (b). Simulations include the quadratic Zeeman shift (see Sec. 6.4.2 and App. A).

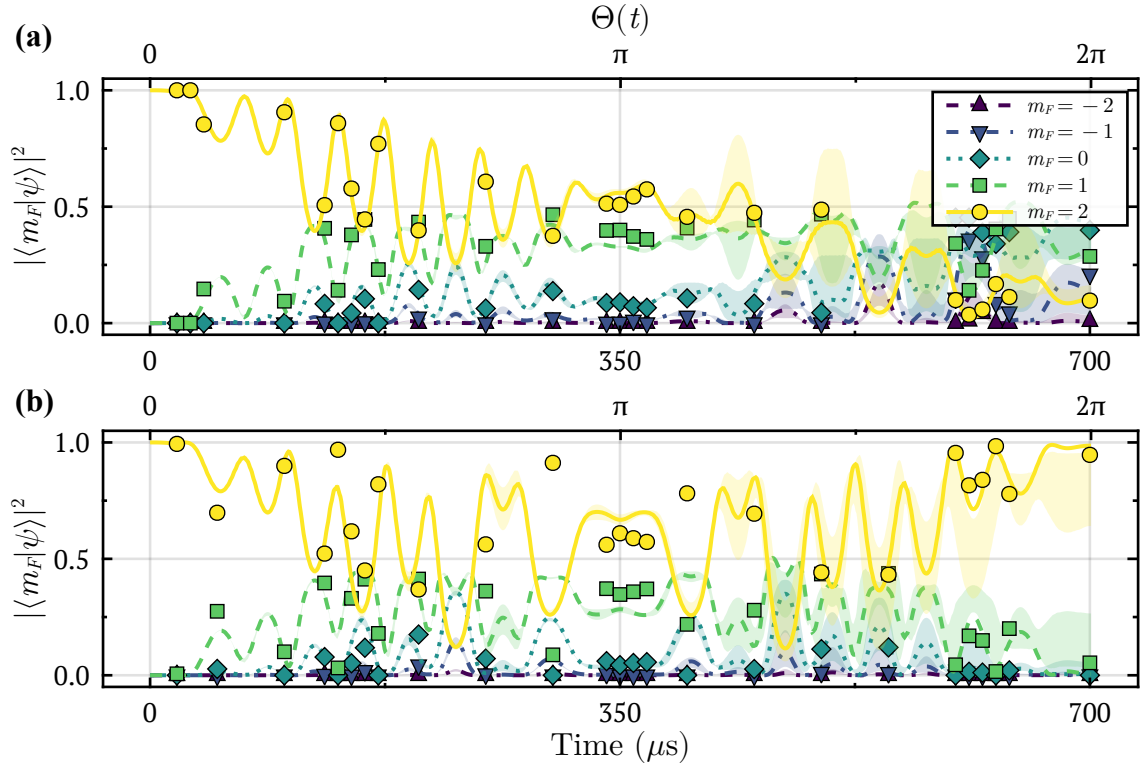


Figure 6.6: Measurements of spin populations (points) on top of numerical calculations (lines) as they vary in time throughout a loop, in $F = 2$. Theory lines were computed by fitting the population data for a value of detuning, Δ_z , which was then used to compute the detuned holonomy from Eq. 5.54. For each measurement, $\Omega_0/\omega = 1$, with $\Omega_0/2\pi = 14.27$ kHz, and $\Omega = \omega/10$. Colored bands display the interquartile range of populations subject to detuning, numerically computed from a Gaussian sample of detunings with mean Δ_z and standard deviation $2\pi \times 0.4$ kHz. Results are shown for ℓ_1 (a) and ℓ_6 (b). Simulations include the quadratic Zeeman shift (see Sec. 6.4.2 and App. A).

The results for the holonomies $\hat{\Gamma}_A(\ell_1)$, and $\hat{\Gamma}_A(\ell_3)$ are shown, in Fig. 6.5(a) and Fig. 6.5(b), respectively, both for $F = 1$. The initial state for both was $|F = 1, m_F = +1\rangle$. Loop ℓ_1 demonstrates the non-Abelian character of these transformations, as the initial state has been coupled to the other spins over the duration of the loop, ultimately ending in a superposition. This is not conclusive evidence of a non-Abelian connection, as discussed in Sec. 4.1.4, but provides strong evidence nonetheless, and gives some insight into how such phases may be used in the context of QIP (Sec. 4.3). The loop ℓ_3 only couples spins intermittently throughout the transformation, due to the micromotion of the spin basis, stroboscopically coinciding with the Floquet basis in which the state remains polarized for the entire loop.

For loop ℓ_1 the impact of detuning is apparent in the final populations, while loop ℓ_3 is somewhat impervious to the detuning. This may be understood by looking at the form of the detuning in the Floquet basis, Eq. 5.53. The detuning does not commute with the geometric phase for ℓ_1 . As such, both terms impact this loop, resulting not only in phase error, but also the dynamical coupling between Floquet states. Since these measurements, are insensitive to relative phase, we observe the latter source of error. For loop ℓ_3 , however, the detuning commutes with the geometric phase. This results only in phase error, which may not be observed through these measurements.

Similarly, the results for loops ℓ_1 and ℓ_6 are shown in Fig. 6.6(a) and Fig. 6.6(b), respectively, for the $F = 2$ manifold. The results are remarkably similar for ℓ_1 , as the impact of detuning results in dynamical coupling between spins, as it accumulates over the course of a loop. At first glance, the loop ℓ_6 seems to result in effectively the same transformation as ℓ_3 , only differing throughout the loop (but not at the end). Due to the connection not commuting with itself at various points along this path, however, there is a small amount of coupling from this holonomy, but it is eclipsed entirely by the effects of the detuning. Furthermore, unlike ℓ_3 , the detuning results in both phase and dynamical coupling errors here.

Altogether, the results from these time-scans serve as an initial verification that the holonomies work as expected, and they also emphasize the impact of what would otherwise be near negligible detuning. To better understand these transformations, we must also look at the relative phases imparted on the spin states, and how they too are impacted by detuning.

6.3.2 Phase Control

Having verified the time evolution of spins over the duration of a loop, the relative phases imparted on them must also be verified. After applying a holonomy, a readout RF $\pi/2$ -pulse was applied with a varied phase, to alter the basis in which

the projective measurement is made. Scanning the phase of the readout pulse in the range $[0, 2\pi]$ reveals a portion of the phase information¹³. For loops which result in the same spin populations, such as ℓ_1 , ℓ_2 , and ℓ_4 , this is the only way to discern the difference they have on the holonomies.

Similar to the time scans discussed prior, we collected a set of population measurements for varied readout RF phases. The results were fit for a single value of detuning, Δ_z , the only free parameter. This value was then used to numerically compute a detuned holonomy from Eq. 5.54, to which we applied the corresponding readout RF operators to model the sequence of measurements¹⁴. The uncertainty bands on these theory calculations are also computed in the same way, through sampling 1000 other detunings in the vicinity of Δ_z and plotting the interquartile range. In this case, due to the phase sensitivity of these measurements, the detunings were sampled from a Gaussian with a standard deviation of $2\pi \times 0.2$ kHz.

In Fig. 6.7 and Fig. 6.8, we demonstrate the difference in phase of the states resulting from holonomies $\hat{\Gamma}_A(\ell_1)$ and $\hat{\Gamma}_A(\ell_2)$, for $F = 1$ and $F = 2$, respectively. These holonomies, along with $\hat{\Gamma}_A(\ell_4)$, produce identical results in the \hat{F}_z basis, with states differing only in relative phase; this can be distinguished through the phase sensitive measurements discussed here, as shown by a shift in the population trend with readout RF phase. The phase imparted onto the states may therefore be controlled just by altering the loop taken. Unfortunately, the same shifts may be caused by detuning as well, though with the additional effect of reducing the amplitude of the observed crests, and dynamical coupling of spins. For this reason it was essential that detuning be the only free parameter in fitting, as almost any other parameter could result in a similar phase shift, preventing us from discerning the true origin as arising from the path-dependence of the holonomy.

The scans shown in Fig. 6.7 and Fig. 6.8 are only a small representative sample of several hundred such scans, which were collected gradually over the course of the project to verify that the holonomies were implemented correctly. These scans, along with the time evolution scans, formed the backbone of our verification efforts, with the help of the numerical model (Sec. 5.2) that was developed concurrently. It is through the phase scans in particular that the impact of detuning was identified. Once these results match expectations, we were left with the task of performing a more complete characterization of the transformations, by measuring the holonomies in full.

¹³The redundancy of using readout pulses with both 0 and 2π phases was intentional, as any difference between these measurements potentially demonstrates a drift in detuning over the course of the scan.

¹⁴The evolution operator for a readout RF pulse can be computed analytically from Eq. 2.19.

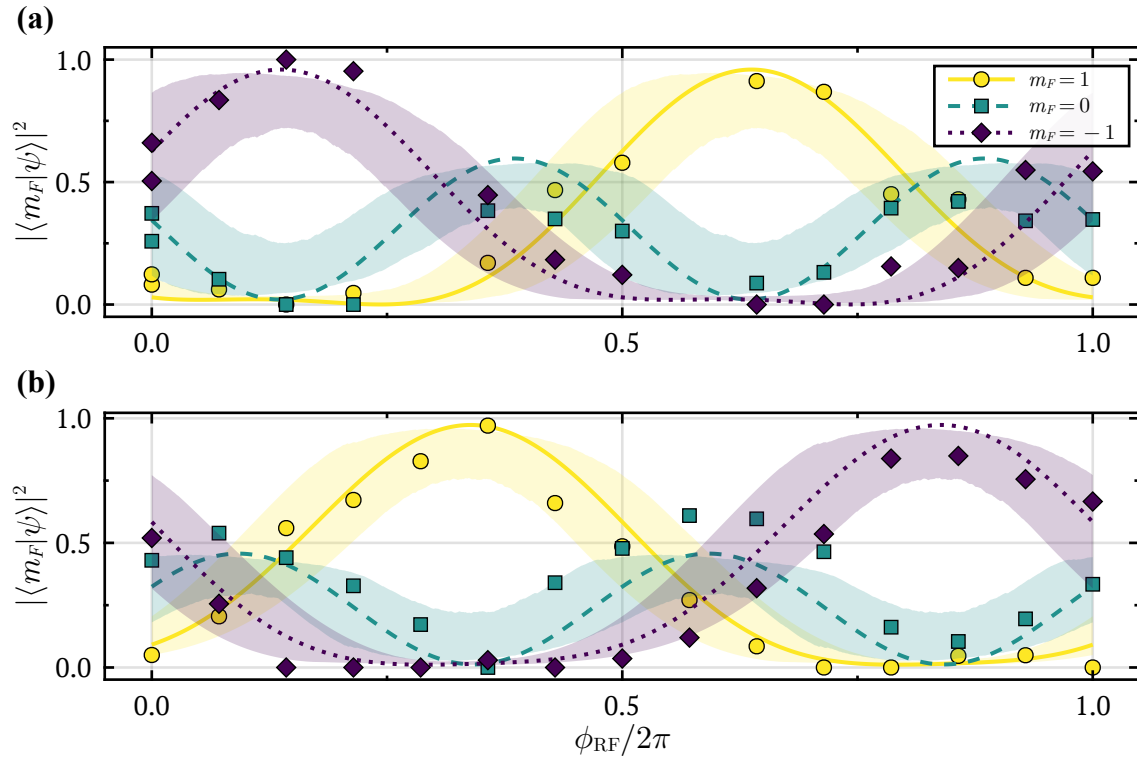


Figure 6.7: Measurements of spin populations (points) in different bases, as defined by the phase of a readout **RF** pulse following a holonomic transformation, in $F = 1$. Numerical calculations (lines) are overlaid, computed by fitting the population data for a value of detuning, Δ_z , which was then used to compute the detuned holonomy from Eq. 5.54, and analytically modeling the measurement sequence. For each measurement, $\Omega_0/\omega = 1$, with $\Omega_0/2\pi = 10.64$ kHz, and $\Omega = \omega/10$. Coloured bands display the interquartile range of populations subject to detuning, numerically computed from a Gaussian sample of detunings with mean Δ_z and standard deviation $2\pi \times 0.2$ kHz. (a) shows the results for holonomy $\hat{\Gamma}_A(\ell_1)$, while (b) shows that for $\hat{\Gamma}_A(\ell_2)$. The shift in population data with respect to the readout **RF** phase between (a) and (b) is indicative of the path-dependence of the holonomies. Simulations include the quadratic Zeeman shift (see Sec. 6.4.2 and App. A).

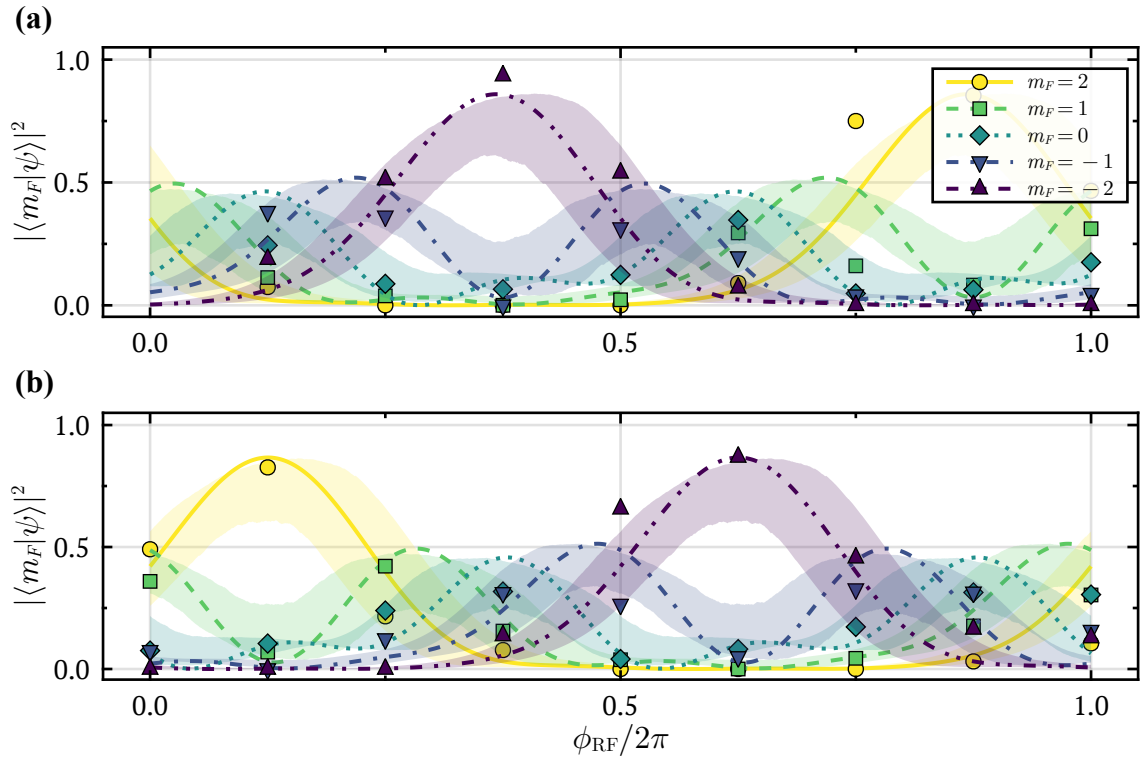


Figure 6.8: Measurements of spin populations (points) in different bases, as defined by the phase of a readout RF pulse following a holonomic transformation, in $F = 2$. Numerical calculations (lines) are overlaid, computed by fitting the population data for a value of detuning, Δ_z , which was then used to compute the detuned holonomy from Eq. 5.54, and analytically modeling the measurement sequence. For each measurement, $\Omega_0/\omega = 1$, with $\Omega_0/2\pi = 14.27$ kHz, and $\Omega = \omega/10$. Colored bands display the interquartile range of populations subject to detuning, numerically computed from a Gaussian sample of detunings with mean Δ_z and standard deviation $2\pi \times 0.2$ kHz. (a) shows the results for holonomy $\hat{\Gamma}_A(\ell_1)$, while (b) shows that for $\hat{\Gamma}_A(\ell_2)$. The shift in population data with respect to the readout RF phase between (a) and (b) is indicative of the path-dependence of the holonomies. Simulations include the quadratic Zeeman shift (see Sec. 6.4.2 and App. A).

6.3.3 Fidelity of Holonomies

Motivated by computing the Wilson loops to characterize the holonomies in a gauge-invariant manner (Sec. 4.1.4), we measured the holonomies in full. Performing state tomography (App. C) on a transformed basis state yields a single column of the evolution operator (holonomy). Therefore, if we perform a series of tomographic measurements on the results of each basis state after a transformation, we obtain the full evolution operator.

For each holonomy in Tab. 5.1 we performed such a scan, wherein each basis state was prepared, transformed, and a series of tomographic measurements were performed. We repeated these scans many times for each holonomy, due to the presence of detuning in each trial. Furthermore, the measurements in each scan were made in a randomized order, including which basis state and measurement pulse was applied, to avoid potential bias from ordering. More measurements than explicitly necessary for tomographic reconstruction of the states were also taken, for better fitting, and to account for occasional mis-triggers of our AWG in which the pulses were not applied¹⁵. From these scans we reconstructed the holonomies.

Holonomy Fitting

Each holonomy scan consists of a series of population measurements from several readout RF-pulses, on the state resulting from a holonomy applied to each spin basis state. The results for each independent basis state may be used to reconstruct the wavefunction and corresponding column of the holonomy, but performing the fits in this way would make the residuals in each scan have no impact on the others; for a more comprehensive analysis, the entire matrix may be fit all at once.

To fit a matrix to each scan, we need a parameterized evolution operator from which the various measurements may be simulated. The holonomy produced here may be decomposed into the basis of spin matrices as,

$$\hat{\Gamma}_A(\ell) = \exp\left(-i\mathbf{q} \cdot \hat{\mathbf{F}}/\hbar\right), \quad (6.10)$$

where the components of the vector \mathbf{q} parameterize the transformation. For some \mathbf{q} the holonomy may be computed, and the various readout RF-pulses are simulated to compute the resulting set of populations. We may then fit for \mathbf{q} by comparing these simulated measurements with the data. This choice of decomposition (Eq. 6.10) is problematic in higher spin systems, however, as it neglects potential effects outside

¹⁵Occasionally the AWG did not properly trigger, therefore not applying any control pulses; due to the intermediate resonant laser pulses in the state-preparation step, this would result in no atoms whatsoever. These points are easy to remove, but the resulting scan would be missing a data point, which inhibits state tomography.

of $SU(2)$. Applying such a decomposition would only serve to validate our initial assumptions that the holonomies induce $SU(2)$ transformations. Therefore, to remain as general as possible, the decomposition used for spin manifold F should be in terms of the generators of $SU(2F + 1)$.

In the case of our measurements, in $F = 1$, we expressed the holonomy as,

$$\hat{\Gamma}_A(\ell) = \exp\left(-i\mathbf{q} \cdot \hat{\mathbf{\Lambda}}/\hbar\right) \quad (6.11)$$

with $\hat{\mathbf{\Lambda}}$ the vector of Gell-Mann matrices, generators of $SU(3)$. This decomposition describes the most general evolution operator in $F = 1$, and therefore captures any potential coherent sources of phase¹⁶. The norm-squared difference between the simulated measurements and the data were then used as the objective function for the Nelder-Mead nonlinear optimization algorithm, which we implemented with the `Optim.jl` package in Julia [224]. Since the holonomy is generated in this way, it has no information about the actual Hamiltonian that produced the data, and so it is inherently unbiased.

Before implementing these fits on the data, we numerically simulated a set of measurements for each holonomy, and added measurement noise in varying amounts. The efficacy of this fitting routine was then characterized against this data. We found that for moderate levels of noise in the populations, even less than the deviations caused by detuning throughout our scans, the initial “guess” for the vector \mathbf{q} (thus, the holonomy) was important in having the fit converge properly. This is likely due to the number of parameters versus data points; there are eight Gell-Mann matrices, making these fits consist of eight parameters fit to about 15 points in each scan. A good initial guess may be found by taking the natural logarithm of the expected holonomy, and computing the inner products of the result with each Gell-Mann matrix, resulting in the corresponding components of \mathbf{q} . When this routine was applied to the data, we used the detuned holonomy $\hat{\Gamma}_A^\Delta(\ell)$ to generate the initial guess, following a fit of detuning Δ_z to the scan.

To summarize, we collected a series of scans for each holonomy in Tab. 5.1 that consisted of an informationally complete set of measurements, which may be used to reconstruct the holonomies in full. For each scan, we fit for a single value of detuning, Δ_z , as in the other scans described in Sec. 6.3.1 and Sec. 6.3.2. We then fit each scan with a single holonomy, through an $SU(3)$ decomposition to avoid introducing any bias. From this we are able to analyze the characteristics of the each holonomy, for which we have a set of measurements with various detunings.

¹⁶This still assumes that the evolution is unitary.

Fidelity

To compare the measured holonomies with expectations, we used a common quantitative method used in QIP applications, inspired by the overlap of these experiments with HQC. A simple way to quantify the success (or failure) of our quantum control is to compute the inner product between the operator we measured, and the one we attempted to implement (the target). In the QIP community this has been formalized into the *fidelity* of an operation, defined through the Frobenius inner product between matrices.

Supposing we wish to apply the transformation \hat{A} to a state, and measure the resulting transformation \hat{B} , the fidelity of this operation may be defined as [225],

$$\mathcal{F} = \frac{|\text{tr}(\hat{A}^\dagger \hat{B})|}{\sqrt{\text{tr}(\hat{A}^\dagger \hat{A})} \sqrt{\text{tr}(\hat{B}^\dagger \hat{B})}}, \quad (6.12)$$

where in this definition, the operators \hat{A} and \hat{B} need not be unitary. This is the general form of fidelity as applied in various QIP applications, to determine the quality of quantum control protocols in the presence of systematic noise, and decoherence. For our purposes, the latter effects may be ignored; in this case, the operators \hat{A} and \hat{B} are unitary, so the *pure-state* fidelity becomes,

$$\mathcal{F} = \frac{|\text{tr}(\hat{A}^\dagger \hat{B})|}{2F + 1}, \quad (6.13)$$

where the denominator is now the subspace dimension. This is just the Frobenius inner product, normalized by the size of the matrices. If $\mathcal{F} = 0$, then the two operators \hat{A} and \hat{B} are orthogonal, and if $\mathcal{F} = 1$ then $\hat{A} = \hat{B}$. This gives a simple metric for success.

Conveniently, the definition of fidelity is also reminiscent of the Wilson loop (Sec. 4.1.4). In this context, we may imagine the Wilson loop for the transformation $\hat{A}^\dagger \hat{B}$, where \hat{A}^\dagger is the same as performing the transformation \hat{A} backwards; the fidelity is thus a gauge-invariant measure of the net distortions following these loops. In the case where $\hat{A} = \hat{B}$, then $\hat{A}^\dagger \hat{B} = \hat{\mathbb{1}}$, meaning there is no net distortion beyond identity; this is because we simply performed loop \hat{B} and immediately reversed the operation. Of course, the fidelity only includes two loops so it does not speak to the non-Abelian character of the loops, but this serves as useful insight into the physical meaning behind the quantity.

From the holonomies measured in experiment we computed the fidelities, comparing them with the holonomies $\hat{\Gamma}_A(\ell)$ without detuning; these fidelities are labeled

	ℓ_1	ℓ_2	ℓ_3	ℓ_4	ℓ_5	ℓ_6
N	17	13	13	8	8	11
$\bar{\mathcal{F}}$	0.4(3)	0.5(3)	0.3(2)	0.6(3)	0.5(2)	0.5(2)
$\bar{\mathcal{F}}^\Delta$	0.7(2)	0.8(2)	0.8(2)	0.90(6)	0.9(1)	0.8(1)

Table 6.1: Statistics on fidelities of measured holonomies. N is the number of measurements for each loop, $\bar{\mathcal{F}}$ is the mean fidelity without detuning considered, while $\bar{\mathcal{F}}^\Delta$ is the mean with detuning included. Results are shown with the uncertainty in the last digit given from the corresponding standard deviation. When detuning is properly accounted for in the target operator, fidelities are higher, and more narrowly distributed with respect to the mean.

as \mathcal{F} . The results are shown as a *raincloud* plot in Fig. 6.9, which consists of a density distribution (histogram) plotted over the individual points¹⁷ The resulting fidelities are broadly distributed, with some transformations that closely match predictions, and others that are entirely different. This indicates overall low-quality quantum control through the Floquet engineering approach, but does not suggest a reason for why. The statistics from Fig. 6.9 are summarized in Tab. 6.1.

To investigate the cause of our lack of control, each set of holonomy measurements were also fit for the detuning, Δ_z . We then re-computed the fidelities, comparing measurements with the detuned holonomy $\hat{\Gamma}_A^\Delta(\ell)$; these fidelities are labeled as \mathcal{F}^Δ . These results are more narrowly distributed, closer to higher fidelities, as summarized in Tab. 6.1. This indicates that detuning is the most significant source of error in our control, as when it is properly accounted for then the transformations mostly follow what is expected by our detuning model (Sec. 5.3). A histogram of the magnitudes of all detunings from holonomy fits is shown in Fig. 6.10; the distribution has a mean of $\mu = 0.45$ kHz, and standard deviation of $\sigma = 0.37$ kHz, which are consistent with expectations from the noise level in our bias-coils supply (Sec. 6.2.2).

It is important to note that the results for \mathcal{F}^Δ being generally higher than those for \mathcal{F} is only an indication of the impact of detuning; the detunings were only extracted after-the-fact, and so the results for \mathcal{F} represent the true limits to our control capability. We have merely used fidelity here to demonstrate, through our model of detuning (Sec. 5.3), that it accounts for the largest source of error in our control.

¹⁷This format was necessary due to the low number of measured operators, providing poor statistics.




Figure 6.9: Density distributions of the measured fidelities for each loop in Tab. 5.1. Black vertical bars represent the means. The purple (dark) distributions are results when measured fidelities are compared with the undetuned holonomy, $\hat{\Gamma}_A(\ell)$ (Eq. 5.37). Due to the significant impact of detuning on the transformations, the fidelities are generally low, and broadly distributed. Each set of holonomy measurements is fit for a single value of detuning Δ_z , which is then used to compute the detuned holonomy, $\hat{\Gamma}_A^\Delta(\ell)$ (Eq. 5.54). The corresponding distribution of fidelities are shown in yellow (light). These results are grouped much closer to unity, and are more narrowly distributed as the effects of detuning have been accounted for.




Figure 6.10: Magnitude of detunings extracted from fits to measured holonomies, displayed as a histogram. The mean detuning magnitude is $\mu = 0.45$ kHz, with a standard deviation $\sigma = 0.37$ kHz. These values are consistent with the level expected from the bias coil power supply (Sec. 6.2.2) Total number of measurements is 70.

6.4 Discussion

The results presented in the previous section (Sec. 6.3) demonstrate several important features of the holonomic transformations investigated here, enabled by Floquet engineering, but fall short of a full characterization due to the limitations imposed by detuning. Here we discuss several important aspects of the preliminary investigation presented here, starting by addressing the issue of detuning (and relatedly, the quadratic Zeeman shift), and how it may be circumvented in future implementations. We then discuss the generality of this approach, in regards to applying the Floquet-engineering technique in other systems, and for using it to generate other interesting artificial gauge fields. Finally, we comment on the relation of this scheme to QIP, as a means to perform HQC in nondegenerate systems.

6.4.1 Detuning Sensitivity

Utilizing the scheme as described, we were able to implement each of the holonomies in Tab. 5.1. Our results (Sec. 6.3) agree with simulations in both the time-domain, and in phase of the prepared states, if the significant impact of detuning is considered; specifically, a z -component of the detuning. Detuning was identified as the primary source of error in these measurements by measuring the holonomies, and computing the fidelities both with and without detuning considered; the latter resulted in significantly lower, and more broadly distributed fidelities. As such, this preliminary investigation has verified several important aspects of this technique, through the agreement with our model, but the presence of detuning has significantly hampered our full characterization of the connection; measurement of the trace-commutator would provide a gauge-invariant confirmation of the non-Abelian

character of the connection, but such a manifestation is undermined by dynamical effects beyond the experimenter’s control (Sec. 4.1.4).

In light of the role that detuning played in these experiments, it is prudent to address how one could attempt to overcome these difficulties. First and foremost, we will reiterate here that the current setup has no active feedback control of resonance, aside from that built-in to our bias coil supply¹⁸. One could envision an array of magnetic field sensors placed in various locations near the ensemble to actively sense ambient fields, forming a feedback loop with the current supply to adjust the coil current in real time. This is an active research direction in our group, but finding sensors and corresponding algorithms for determining fields with the sensitivity required here will prove to be challenging. Alternatively, using the ensemble itself as a magnetic antennae is a promising approach, as through Faraday magnetometry or other optical magnetometry techniques [226, 227]. Some such protocols are non-destructive and operate with excellent sensitivity. In either case, such an upgrade to the system, in addition to some magnetic shielding, could prove useful.

Relatedly, our means of resonance calibration as described in Sec. 6.2.2 is limited by the tools we have available. We found that the holonomies themselves are more sensitive to ambient fields than traditional RF pulses. Even still, results show that our resonance detection capability was insufficient to remove the dynamical contributions, despite every scan beginning “on resonance”, as far as this researcher could tell. It is only after a full scan, with many data points, that a fit to the set of measurements could reveal a detuning. In the absence of more sophisticated magnetometry techniques discussed above, perhaps a longer pulse, or ones curated more specifically for this purpose could be utilized to obtain better sensitivity; in any case, our bias field control would also need to be improved.

Altogether, we believe some selection of the improvements discussed here would be sufficient to eliminate the detuning to a degree high enough to enable application of the Wilson loops to these transformations. There are other similar ultracold atom platforms boasting much higher resonance fidelity [67, 228, 229]. For instance, the field stability reported in *Sugawa et al. 2021* [67] is 2.5 ppm for a bias field of 19.8 G, in a remarkably similar system to ours. This corresponds to 35 Hz stability in the m_F level splitting, in ^{87}Rb . With detuning reduced to this level, the fidelity expected for each of the loops in Tab. 5.1 would exceed 0.99, as verified through numerical calculations.

¹⁸The supply is operated in constant-current mode, which has provided notably more stable resonances than constant-voltage mode, as expected. The inner workings of the supply in either case are beyond our control.

6.4.2 Other Sources of Error

In addition to detuning, there is another important source of error that requires mention: the quadratic Zeeman shift. Briefly, this term arises as the next leading order correction to the splitting of m_F levels in a magnetic field, as described in App. A. Unlike the linear Zeeman effect (Eq. 2.6), the quadratic shift ϵ does not split states linearly, as implied by the name. Rather, it imparts an additional shift of an amount proportional to m_F^2 . Looking at the $F = 1$ manifold for instance, the additional splitting ϵ for levels $m_F = \pm 1$ is the same, and in the same direction. The $m_F = 0$ level, on the other hand, is unaffected. As such, the manifold no longer has uniform spacing, so the resonance condition differs by transition. The quadratic Zeeman Hamiltonian is therefore described by a $\mathfrak{su}(3)$ term in this case.

In regards to the impact of the quadratic shift in this experiment, its magnitude ϵ is similar in size to the detuning. Specifically, for a linear splitting of 1.25 MHz, $\epsilon \approx 0.228$ kHz (App. A). While it depends on the Zeeman splitting, it is insensitive to the detunings observed here. For instance, if a background detuning alters the level splitting from the nominal 1.25 MHz by 0.8 kHz, ϵ is only changed by 1 Hz. Therefore, for the sake of our experiments, it is constant for all measurements. Similar to the detuning, this term is also undriven by the Floquet envelope. The impact of such terms in the Floquet basis, as covered in Sec. 5.3, is a dynamical phase that breaks the degeneracy between states in the bands, and couples them. This result was computed for undriven terms proportional to the spin-matrices, but to provide a similar description for the quadratic shift this calculation must be carried out again for other generators. Specifically, the commutators computed to re-sum the expansion in Sec. 5.3.1, Eq. 5.43 - 5.45, would take on a different form, but overall it is expected to result in similar dynamical terms as in the case of detuning. We expect, therefore, that the effect of this term on the holonomies is comparable to the detuning, with one major difference; it does not vary substantially between measurements. As such, this term may be treated as systematic source of error.

Based on the above arguments, the quadratic Zeeman shift was considered in all analyses of measurements, numerically. The theoretically predicted value for $\epsilon = 0.228$ kHz was used in all fits to the data. This also is part of the motivation for the holonomy fits to be done with a decomposition in terms of the Gell-Mann matrices (see Sec. 6.3.3), since a spin-matrix decomposition would be insufficient to capture the impact of the quadratic shift¹⁹. Since the quadratic shift is a systematic

¹⁹While testing the fitting routine on numerically generated data, it was found that a spin-matrix decomposition fit to holonomies with a quadratic shift resulted in fidelities limited to a threshold that was determined by the strength of ϵ . Therefore, depending on ϵ , the relative $SU(3)$ contribution it generates was not captured by the spin-matrices, resulting in fidelities constrained by a related value.

error it may be treated in this way, unlike the detuning which varied between each measurements, making it a random source of error.

The static nature of the quadratic shift allows us to conveniently “factor it out” of our investigation, in the sense that it must absolutely be considered in analyses, but ultimately does not alter our ability to make quantitative conclusions on the nature of the holonomies, as the detuning does; its impact may be computed numerically with minimal uncertainty, though since it does not commute with the geometric phase in general, time-ordering becomes important. With a set value of ϵ , measurements of the Wilson loop could, for instance, be performed, and compared with numerical calculations for the given value of ϵ . This was not possible in the case of detuning, because of its random variation between scans resulting in significant uncertainty (see Fig. 5.7). The desired properties of the non-Abelian connection could therefore be deduced, having a thorough understanding of the dynamics imposed by the quadratic shift.

Aside from the interference of external fields or higher order corrections to the Hamiltonian, our attempts at high-fidelity quantum control are prone to several other sources of error. In particular, the measurements used in all analyses are based off absorption images of the ultracold ensembles in TOF. There are several ways in which the population data extracted from these images may encode systematic errors. To understand this, we will first describe the imaging procedure in more detail, though for a more complete understanding, see *Hrushevskyyi 2017* [6].

Absorption imaging is done by illuminating the atomic ensemble with a resonant laser, and imaging the beam on a camera. Atoms in the ensemble scatter light from the laser field, and so the beam wavefront encodes the spatial distribution of the atoms; the denser regions of atoms scatter more light, resulting in lower intensity in the corresponding regions of the wavefront. To account for the inherent intensity pattern of the beam from this result, a second image is taken once the atoms have fallen away, due to the significant momentum imparted onto them from the beam in the first image. In our case, we also capture a third image with no beam or exposure time, to estimate the contribution from dark counts in our camera. These three images are combined to yield a spatial distribution of the *optical depth*,

$$OD = -\ln\left(\frac{I_A - I_D}{I_B - I_D}\right), \quad (6.14)$$

where I_A , I_B , and I_D are the intensity fields from the images with atoms, without, and of the dark counts, respectively. This is computed element-wise, for each pixel in the set of images. For a BEC, it is expected that the density distribution of atoms, and thus the corresponding OD , approximately follows an inverted parabola²⁰ as

²⁰For a parabolic trapping potential, as in the case of our ODT.

discussed in Sec. 3.2.2 in regards to the TF-approximation. If atoms are imaged at long TOF, the distribution may be further approximated as Gaussian.

From these images of the optical depth we extract all of the information used in our analysis. Our measurements consist of images displaying several spatially separated distributions of atoms, due to the SG fields that split the spin-components. To each distribution, a 2D-Gaussian fit is performed; from the pixel size, the number of atoms in each may be extracted. These atom numbers are normalized with each other, resulting in the relative spin populations. This step is performed automatically by the custom imaging software employed in the lab.

Based on the imaging process as described above, there are several potential sources of error. First, the relative populations are extracted from Gaussian fits to the atomic density distributions; the error in these fits produce a corresponding error in the relative populations. It was found that these counting errors were insignificant in comparison to the detuning, accounting for at most a few percent in the populations; it becomes more significant the lower the atom number is in a given distribution, as in these cases the density distributions were closer to background noise than when there are many atoms. As such, this error would predominately affect measurements in which the populations are close to zero, or measurement sequences in which the overall atom number was low. The latter case was avoided by only performing measurements when the number of atoms was high; due to significant fluctuations and instabilities in the atom number throughout the day, this dramatically reduced the amount of data that could be collected, which also explains the relatively small number of holonomies measured (see Tab. 6.1).

The other potentially important effect to consider is that, for images taken at long TOF (about 20 ms in this case), the atoms have fallen close to the edge of our imaging beam, where its intensity is low. The slope of the beam intensity profile in this region is higher than at the beam center, and therefore the intensity used to image each spatially-separated spin component is different. While this effect may be at least partially accounted for in calculating the optical depth (Eq. 6.14), we found empirically that the atom number extracted from fits depended slightly on the location of the atoms within the beam, only near its edge. That said, this effect was difficult to confirm: in order to verify this, atoms with a known atom number must be imaged at various different locations in the beam (say, by varying the TOF). The imaging process is destructive though, so such a measurement sequence is subject to drifts in the atom number between measurements. For a quantitative understanding, some numerical modeling would likely suffice. Regardless, this source of error was also mitigated by tuning the strength of the SG field, making the spacing between spins as small as possible without them overlapping, so that they may still be distinguished. This error is therefore most significant for measurements in the

$F = 2$ manifold, since there are five spin components that must remain spatially separated.

Last, we have neglected the effects of interactions between atoms in the ensemble, which would have an impact on the fidelity. Most obviously, spin-exchange collisions are a decoherence mechanism for spin states; if they are significant, this would have to be accounted for in the fidelity. Additionally, if the intra-spin interaction strengths (see Sec. 3.2) vary with m_F , this would result in a mean-field shift, changing the splitting between spin levels. For instance, the intra-spin interactions for ^{87}Rb atoms in the $F = 1$ manifold differ slightly between the $m_F = \pm 1$ and $m_F = 0$ levels, which would produce such an effect. In either case, such an effect would present itself in a similar way to detuning, since the interactions are not modulated by the Floquet drive. In the experiments discussed here, the ensemble is dilute, so all interactions are suppressed over the relatively short time-scales of these experiments (< 1 ms). If other errors are reduced, the effects of interactions may eventually need to be considered.

Overall, we found that the sources of error from imaging discussed above were insignificant when compared with the detuning, since they could be suppressed by working with larger atom numbers, and small spacing between spin components. As such, the fluctuations in population between measurements was attributed in full to the detuning, stemming from our lack of control over background magnetic fields, and uncertainty in the bias power supply current. If the impact from detuning is reduced, such as by implementing some of the suggestions in Sec. 6.4.1, then these imaging errors must be revisited.

6.4.3 Floquet-Engineered Artificial Gauge Fields

The theory described in Ch. 5 demonstrates that, for a periodically modulated Hamiltonian, a degenerate set of states may be found in the Floquet basis, wherein non-Abelian geometric phases may arise. This notion was applied specifically to the Hamiltonian of a spin in a magnetic field (Eq. 5.1), resulting in a connection that generates $SU(2)$ transformations; this is expected, since the original Hamiltonian also produces such transformations. More generally, it is expected that *any* Hamiltonian driven in this way would result in similar physics. It is therefore worth investigating the dynamics of other, similarly parameterized Hamiltonians driven in a similar manner. For instance, a natural extension of the Hamiltonian for a spin in a magnetic field may be the $SU(3)$ equivalent, given by,

$$\hat{H}_{SU(3)} = \Omega_0 \boldsymbol{\lambda}(t) \cdot \hat{\boldsymbol{\Lambda}}, \quad (6.15)$$

with $\hat{\mathbf{A}}$ the vector of Gell-Mann matrices. In this case, the orientation of this field in time would be described by the eight-dimensional unit vector $\boldsymbol{\lambda}(t)$. The derivation of this Hamiltonian in the Floquet basis will follow a similar process as that in Sec. 5.1; the commutator expansion, Eq. 5.6, will converge differently for Gell-Mann matrices than for spin matrices. Supposing the Hamiltonian transforms in a similar manner to the case in $SU(2)$, the connection here would correspondingly generate transformations in $SU(3)$. This is, therefore, an interesting topic for future study, due to its application in high-energy physics [83]. Similar extensions to fields of other interesting symmetries may also be enabled by such the Floquet-engineering technique demonstrated here.

As with the theory of the Floquet-engineering technique, the modulations scheme presented in Sec. 6.2.1 is purposefully general, and may be extended to other systems with similar level structure and coupling control. For instance, in alkali-metal atoms the RF-coupling used there may be immediately replaced with a Raman M -scheme coupling, with minimal changes to the effective Hamiltonian (Sec. 2.2.2), though with much higher Rab -frequencies due to the relative strength of electric-dipole to magnetic-dipole transitions. This modulation technique may be readily adopted in other platforms with similar control capabilities, such as trapped ions, Rydberg atom arrays, and superconducting qubits.

In all cases, the present scheme enables $SU(2)$ transformations, regardless of the spin manifold F , but this too may potentially be extended to other symmetries by introducing more couplings (and subsequently more control parameters). For instance, in the case of the $SU(3)$ Hamiltonian Eq. 6.15, it may be possible to realize such a system in the $F = 1$ manifold of ultracold atoms. Such a Hamiltonian would require individually tuneable couplings between each of the three m_F levels, with amplitude and phase control, in addition to control over the energies of each state. In this circumstance, the quadratic Zeeman shift may act more as a resource than a source of error, since for large linear Zeeman splitting it will result in non-uniform spacing between states. If the non-linearity in this splitting is brought outside of the RF-transition linewidths, then the couplings between levels may be individually addressed through frequency selection, in a similar way to the microwave transitions we used for state-preparation (Sec. 6.1.1). The difficult part of this adaptation would be to engineer a coupling between the $m_F = \pm 1$ states, since they may not be directly coupled by a magnetic dipole transition due to selection rules [12]. As such, this coupling would require an intermediate level, such as in a Raman Λ -scheme (Fig. 2.4). As such, the modification of this coupling scheme to generate other Hamiltonians of interest also requires further investigation.

6.4.4 Floquet-Engineered HQC

It was pointed out in Sec. 4.3.1 that holonomies are the essential ingredient in HQC. In this context, the holonomies studied here may be interpreted as single-qubit gate operations, since they are members of $SU(2)$. A true single qubit is comprised of only two internal states, so with the additional levels in the hyperfine manifolds studied here, more information may be encoded, potentially permitting a d level qubit, or *qudit*. Despite this, the $SU(2)$ nature of the transformations restrict our ability to manipulate these levels in the manner required for full qudit control, effectively making it a qubit system. In fact, we can go one step further and claim that there is a representation of the higher spin systems used here in terms of several qubits; due to the $SU(2)$ symmetry of these gates, the evolution of the qubits are inextricably linked, making each qubit essentially a redundant copy of the others. Hence, these experiments quite literally demonstrate single-qubit gate operations.

An example way in which higher spin manifolds may be split into multiple spin-1/2 systems is through the Majorana-stellar representation [100]. In the Majorana-stellar representation, each unique spin state is decomposed into a set of spin-1/2 states, or qubits; these states may be represented on the Bloch-sphere, for instance, not as a single vector but as a “constellation” of the individual qubit vectors. Every unique spin state has a correspondingly unique constellation. Under arbitrary unitary transformations, each qubit may evolve independently of the others, but here the transformations have a more restricted symmetry. In this case, the evolution of each qubit is linked in such a way that they are not independently tuneable, so all qubits in the decomposition move together with a pre-defined relationship. Therefore, through the Majorana-stellar representation, we may consider the holonomies applied here to effect transformations on single qubit, with several redundant copies.

A natural extension of the Floquet-engineering technique would be to multi-qubit gate operations, which has already been proposed in Rydberg atom arrays [220]. As described in Sec. 6.4.3, the evolution of states by *any* modulated Hamiltonian may be represented by a similar non-Abelian geometric phase. The Hamiltonian for an entangling interaction between qubits may therefore be modulated in a similar manner, to be made holonomic. It may therefore be possible that the Floquet-engineering approach studied here can be used for universal QC, but specific gate architectures need to be developed, and fault-tolerance investigated, to determine if there are any advantages over other schemes. It is likely that this will depend on the system in which it is implemented.

The detuning observed in experiments also has interesting implications when viewed through the lens of HQC, namely in regards to fault-tolerance. As discussed in Sec. 4.3, geometric gates are thought to be inherently robust to various forms

of noise or error [72]. Here, we have identified detuning as a source of error that these gates seem to be sensitive to. While the specific form of the detuning in the Floquet basis is somewhat complicated, the result of adding a similar dynamical term to the desired geometric phase is fairly general to HQC overall. Geometric gates are typically resilient to errors in the rate of loop traversal, and to high-frequency fluctuations of the path since they tend to average out [124–131]; the detuning considered here is neither, and suffice to say, does not seem to have an apparent geometric interpretation as its purely dynamical in nature. These results seem to indicate one of the ways in which holonomies used in HQC are not fault-tolerant.

Relatedly, and similar to the arguments presented in Sec. 6.4.1, in the context of HQC, the effects from the quadratic Zeeman shift may be ignored here too. This is because in a two-level system, no such term may exist; the Hamiltonian is fully described by the spin matrices, and hence any error from unmodulated terms are captured by the generalized detuning (Sec. 5.3). Hence, if this protocol were implemented in a true two-level system then there is no quadratic shift to consider. As such, if the protocol presented here is implemented in qubits, it will have no impact on the resulting gate fidelities.

6.5 Conclusion

Our experimental results, presented in Sec. 6.3, substantiate aspects of the theory that these transformations arise from a non-Abelian connection in parameter space (Ch. 5) through their excellent agreement with our numerical models; we were able to verify that the state evolution in time, and phase imparted on states, both follow that predicted by holonomies as a function of the path chosen. Furthermore, we demonstrated that the holonomies measured, when detuning was accounted for, match those expected to greater than 70 %.

Overall, our experiments provide evidence that Floquet-engineering may indeed be used to generate non-Abelian holonomic transformations, but falls short of a gauge-invariant verification. This is due to our lack of control over the detuning, which hampered our ability to implement Wilson loops to truly verify the non-Abelian character of the connection. As such, more work is required. Despite this our experiments do demonstrate the efficacy of this technique as a means of quantum control, perhaps with applications in QIP. The practical technique to realize these driven Hamiltonians, as described in Sec. 6.2.1, may be generalized to other systems with similar level structure and coupling control. This work has therefore introduced another means of adiabatic quantum control, regardless of the underlying symmetry of the geometric phase. With more complicated parametrizations, it is possible that

such a scheme could be used to generate other non-Abelian geometric phases, for instance generating transformations in $SU(3)$, which would be of great interest to the artificial gauge-field and quantum simulation communities. We anticipate that the Floquet-engineering technique will become an important member of the quantum simulation toolbox for generating interesting artificial gauge fields, and potentially find its use in implementations of [HQC](#) as well.

CHAPTER 7

Conclusion

IN conclusion, we have described several ways in which artificial gauge fields may be produced in ultracold ensembles of alkali-atoms, using various forms of spin-dressing by external radiation. The thesis was organized into three main parts: first, we summarized the procedure used throughout the thesis work to produce BECs of ^{87}Rb , along with the primary techniques for manipulation of their ground state spin levels, and measurement thereof. Second, in Ch. 3 we investigated the use of such artificial gauge fields in understanding the many-body states of the interacting ensemble through numerical simulations. Third, Chs. 4-6 looked at how geometric phases may be interpreted through the lens of gauge theory, and how more complicated gauge fields may therefore be studied in this manner; we focused specifically on the production of non-Abelian gauge fields through Floquet engineering, which was then implemented in experiment. Here we will independently summarize these primary contributions, before concluding overall.

The methods discussed in Ch. 2 describe a procedure for reliably producing large ultracold ensembles of ^{87}Rb , which may then be manipulated in various different ways to perform quantum simulation experiments. Of these techniques, we described both RF- and laser-dressing in detail; together, they establish a means to manipulate spins within the ground state hyperfine levels of alkali-metal atoms, with readily available hardware. While the techniques introduced here produce transformations in $SU(2)$, they quite generally describe the interactions between magnetic sublevels through either source of radiation. In each case, they may be extended to more complicated transformations, such as through the quadratic Zeeman shift (App. A). This chapter was completed by describing the SG-TOF technique used to readout the relative spin populations from an ultracold ensemble, in various different bases obtainable through RF-pulses. Altogether, Ch. 2 establishes the necessary ingredients to perform complicated manipulations of spinor BECs, for use in a broad range of quantum simulation experiments.

The second primary part of this thesis was contained within Ch. 3, which discussed artificial gauge fields that were generated by engineering a Hamiltonian with Raman dressing (Sec. 2.2.2). The engineered Hamiltonians generated transformations in spinor BECs, coupling the spin and momentum degrees of freedom. This coupling could be described by the presence of Abelian vector and scalar potentials in the Hamiltonian, for which there are corresponding electric and magnetic fields. We concerned ourselves with finding the many-body ground states of this system in various configurations, using our intuition from the behavior of charged particles in classical fields to explain the results. The goal of these studies was to learn something about quantum many-body systems, which remain intractable even within mean-field theory due to the nonlinear interactions. As such, we described a numerical algorithm for finding ground state solutions of spinor BECs, by integrating the GPE; this algorithm was implemented in Python, and accelerated on GPU hardware, which drastically improved its performance. Equipped with the GPU accelerated code, we were able to simulate a variety of different artificial gauge fields at low computational cost; we included the example of an artificial magnetic field, which permitted the nucleation of quantized vortices in the BEC. Chapter 3 describes a powerful numerical technique for predicting ground states and dynamics of BECs under the effects of complicated external fields, such as those used to produce artificial gauge fields, which provides an indispensable tool for researchers looking to find exciting new experiments to perform, or to explain existing ones.

The final major part of this thesis was split between Chs. 4-6, which described how the study of geometric phase in quantum systems is synonymous with that of artificial gauge fields. We introduced the concept of geometric phase in Ch. 4 as resulting from the condition for parallel transport of states through an adiabatically evolving Hilbert space, with curvature; the transformations of states are described by a holonomy. We placed particular emphasis on non-Abelian geometric phases, which have traditionally been difficult to synthesize in experiment. The differences between Abelian and non-Abelian phases were discussed, as well as a means to experimentally distinguish them, by the measurement of a gauge-invariant quantity called the Wilson loop. We then described some of the close relations between geometric phase and the gauge fields of interest to the quantum simulation community, in addition to its application to QIP, where it serves as the primary resource for several forms of QC, most notably HQC.

An effort to circumvent many of the challenges faced by the community in producing non-Abelian geometric phases is discussed in Ch. 5. Here, we describe a means to use periodic driving of a parameterized Hamiltonian in order to engineer an adiabatically degenerate set of states, comprised of a mixture of explicitly non-degenerate spin states. Through the adiabatic evolution of parameters, non-Abelian

geometric phases may be observed. We also found that if the periodic modulations are shifted from a zero time-average, this *detuning* leads to a significant alteration of the transformations; the geometric phase remains unchanged, but there are additional dynamical contributions resulting in a broken degeneracy, and dynamical coupling between states. The holonomy takes on a new, more generalized form, in terms of the Anandan connection [66]. Chapter 5 establishes Floquet engineering as a powerful quantum control technique for surpassing the explicit need for degeneracy in producing non-Abelian geometric phase, and goes further to describe its sensitivity to time-independent shifts of the periodic driving; this latter point is of substantial interest to those who aim to implement such a technique in experiment.

Finally, we conclude the third major part of the thesis in Ch. 6, where we introduced a way to realize such transformations in the ground state hyperfine manifolds of ultracold alkali-metal atomic ensembles, through RF-dressing (Sec. 6.2.1). We then described a particular implementation in ^{87}Rb , and a means for measuring the non-Abelian holonomies through a set of informationally complete measurements. Our results verify the time-dependence of the holonomies, and the imparted phase, both in the presence of detuning, caused by a mismatch between the RF-resonance and level splitting. The holonomies were further characterized in terms of their fidelity, which were all low due to the significant impact of this detuning; through modeling, we verified that detuning was the largest source of error. Due to the level of detuning observed, we were unable to perform a gauge-invariant characterization of the geometric phase, as through the Wilson loop.

Despite these limitations, the preliminary investigation in Ch. 6 has identified many important practical considerations when implementing such a technique in ultracold ensembles. While we are not able to rigorously conclude that the geometric phase is non-Abelian, as through the Wilson loop, our results provide extensive evidence, through the path-dependence of the holonomies, and the observed relation between the Floquet basis and rotating frame; together with the analysis of the detuning against our numerical model, we have shown a detailed understanding of the system and our quantum control capabilities. Thinking ahead to future implementations, the RF-modulation scheme is purposefully general, in terms of spin matrices, and we anticipate that such a scheme could be adapted to a plethora of other systems with similar control capabilities, such as trapped ions, superconducting qubits, Rydberg atoms, or liquid NMR systems; the insights established here would prove invaluable in such efforts. We expect that the Floquet-engineering approach will be extended to produce artificial gauge fields beyond $SU(2)$, and perhaps may also find application in HQC.

Overall, this thesis establishes several techniques for producing and probing artificial gauge fields that manifest in different ways in ultracold ensembles. This work

covers all of the necessary components, from the production of such ensembles, to their manipulation and measurement. We have also provided a means to numerically simulate such systems, including the effects of interactions, with low computational cost. Finally, we have introduced a path towards generating non-Abelian artificial gauge fields, through geometric phase. Each aspect provides substantial value to the community, for each part of the research pipeline, all the way from theory to experiment. This author hopes that the thesis will serve as a useful resource to the interested researcher seeking to produce their own artificial gauge fields, or to aid in other quantum simulation endeavors.

References

1. Smith, B. D., Cooke, L. W. & LeBlanc, L. J. GPU-accelerated solutions of the nonlinear Schrödinger equation for simulating 2D spinor BECs. *Computer Physics Communications* **275**, 108314 (2022).
2. Cooke, L. W., Tashchilina, A., Protter, M., Lindon, J., Ooi, T., Marsiglio, F., Maciejko, J. & LeBlanc, L. J. Investigation of Floquet engineered non-Abelian geometric phase for holonomic quantum computing. *Physical Review Research*. (arXiv preprint: 2307.12957) (2024). Forthcoming.
3. Zee, A. *Group theory in a nutshell for physicists* (Princeton University Press, 2016).
4. Scully, M. O. & Zubairy, M. S. *Quantum optics* 1999.
5. Feynman, R. P. Simulating physics with computers. *International Journal of Theoretical Physics* **21**, 467–488 (1982).
6. Hrushevskiy, T. *Quantum gas apparatus for Bose-Einstein condensation of ^{87}Rb* PhD thesis (University of Alberta, 2017).
7. Anderson, M. H., Ensher, J. R., Matthews, M. R., Wieman, C. E. & Cornell, E. A. Observation of Bose-Einstein condensation in a dilute atomic vapor. *science* **269**, 198–201 (1995).
8. Davis, K. B., Mewes, M.-O., Andrews, M. R., van Druten, N. J., Durfee, D. S., Kurn, D. & Ketterle, W. Bose-Einstein condensation in a gas of sodium atoms. *Physical review letters* **75**, 3969 (1995).
9. Dalfovo, F., Giorgini, S., Pitaevskii, L. P. & Stringari, S. Theory of Bose-Einstein condensation in trapped gases. *Reviews of Modern Physics* **71**, 463–512 (1999).
10. Becker, D. *et al.* Space-borne Bose–Einstein condensation for precision interferometry. *Nature* **562**, 391–395 (2018).
11. Metcalf, H. J. & Straten, P. v. d. Laser Cooling and Trapping. *Graduate Texts in Contemporary Physics* (1999).
12. Foot, C. J. *Atomic physics* (OUP Oxford).

13. Tretiakov, A. *Versatile apparatus for ultracold atomic hybrid systems* PhD thesis (University of Alberta, 2017).
14. Rastogi, A. *Light-matter interactions for quantum simulation and quantum memory experiments* PhD thesis (University of Alberta, 2018).
15. Steck, D. A. *Alkali D Line Data* Sept. 2001.
16. Chandrashekar, C. M. Dipole trap for 87 Rb atoms using lasers of different wavelength. *arXiv* (2005).
17. Mudrich, M., Kraft, S., Singer, K., Grimm, R., Mosk, A. & Weidemüller, M. Sympathetic Cooling with Two Atomic Species in an Optical Trap. *Physical Review Letters* **88**, 253001 (2002).
18. Kasper, V. *et al.* Universal quantum computation and quantum error correction with ultracold atomic mixtures. *Quantum Science & Technology* **7**, 015008 (2021).
19. Sakurai, J. J. S. & Napolitano, J. J. *Modern Quantum Mechanics* (Pearson, 2014).
20. Fewell, M. Adiabatic elimination, the rotating-wave approximation and two-photon transitions. *Optics Communications* **253**, 125–137 (2005).
21. Loy, M. M. T. Observation of Population Inversion by Optical Adiabatic Rapid Passage. *Physical Review Letters* **32**, 814–817 (1974).
22. Warren, W. S. & Silver, M. S. The Art of Pulse Crafting: Applications to Magnetic Resonance and Laser Spectroscopy. *Advances in Magnetic and Optical Resonance* **12**, 247–384 (1988).
23. Melinger, J. S., Gandhi, S. R., Hariharan, A, Goswami, D & Warren, W. S. Adiabatic population transfer with frequency-swept laser pulses. *The Journal of Chemical Physics* **101**, 6439–6454 (1994).
24. Tashchilina, A., Cooke, L., Moiseev, E., Lindon, J., Ooi, T., Milson, N. & LeBlanc, L. Suppressing and tuning-out Raman transitions in multilevel alkali atoms via multi-path interference. *Bulletin of the American Physical Society* (2023).
25. Goldman, N, Juzeliūnas, G, Öhberg, P & Spielman, I. B. Light-induced gauge fields for ultracold atoms. *Reports on Progress in Physics* **77**, 126401 (2014).
26. Juzeliūnas, G & Spielman, I. B. Flux lattices reformulated. *New Journal of Physics* **14**, 123022 (2012).
27. Lindon, J., Tashchilina, A., Cooke, L. W. & LeBlanc, L. J. Complete Unitary Qutrit Control in Ultracold Atoms. *Physical Review Applied* **19**, 034089 (2023).

28. Lin, Y.-J., Compton, R. L., Perry, A. R., Phillips, W. D., Porto, J. V. & Spielman, I. B. Bose-Einstein Condensate in a Uniform Light-Induced Vector Potential. *Physical Review Letters* **102**, 130401 (2009).
29. Jimanéz-García, K. *Engineering Hamiltonians for Ultracold Neutral Atoms* PhD thesis (Joint Quantum Institute, 2012).
30. Lin, Y.-J. & Spielman, I. B. Synthetic gauge potentials for ultracold neutral atoms. *Journal of Physics B: Atomic, Molecular and Optical Physics* **49**, 183001 (2016).
31. Radić, J., Sedrakyan, T. A., Spielman, I. B. & Galitski, V. Vortices in spin-orbit-coupled Bose-Einstein condensates. *Physical Review A* **84**, 063604 (2011).
32. Lin, Y.-J., Compton, R. L., Jiménez-García, K., Porto, J. V. & Spielman, I. B. Synthetic magnetic fields for ultracold neutral atoms. *Nature* **462**, 628–632 (2009).
33. Lin, Y.-J., Compton, R. L., Jiménez-García, K., Phillips, W. D., Porto, J. V. & Spielman, I. B. A synthetic electric force acting on neutral atoms. *Nature Physics* **7**, 531–534 (2011).
34. Lin, Y.-J., Jiménez-García, K. & Spielman, I. B. Spin-orbit-coupled Bose-Einstein condensates. *Nature* **471**, 83–86 (2011).
35. Jiménez-García, K., LeBlanc, L. J., Williams, R. A., Beeler, M. C., Perry, A. R. & Spielman, I. B. Peierls Substitution in an Engineered Lattice Potential. *Physical Review Letters* **108**, 225303 (2012).
36. Yılmaz, F. & Oktel, M. O. Artificial magnetic-field quenches in synthetic dimensions. *Physical Review A* **97**, 023612 (2018).
37. LeBlanc, L., Jiménez-García, K., Williams, R., Beeler, M., Phillips, W. & Spielman, I. Gauge matters: observing the vortex-nucleation transition in a Bose condensate. *New Journal of Physics* **17**, 065016 (2015).
38. Galitski, V. & Spielman, I. B. Spin-orbit coupling in quantum gases. *Nature* **494**, 49–54 (2013).
39. Wu, Z. *et al.* Realization of two-dimensional spin-orbit coupling for Bose-Einstein condensates. *Science* **354**, 83–88 (2016).
40. Su, S., Gou, S., Liu, I., Spielman, I., Santos, L., Acus, A., Mekys, A., Ruseckas, J. & Juzeliūnas, G. Position-dependent spin-orbit coupling for ultracold atoms. *New journal of physics* **17**, 033045 (2015).
41. Chen, H.-R. *et al.* Spin-orbital-angular-momentum coupled Bose-Einstein condensates. *Physical Review Letters* **121**, 113204 (2018).

42. Campbell, D., Price, R., Putra, A., Valdés-Curiel, A., Trypogeorgos, D. & Spielman, I. Magnetic phases of spin-1 spin-orbit-coupled Bose gases. *Nature communications* **7**, 10897 (2016).
43. Raman, C, Abo-Shaer, J., Vogels, J., Xu, K & Ketterle, W. Vortex nucleation in a stirred Bose-Einstein condensate. *Physical review letters* **87**, 210402 (2001).
44. Anglin, J. R. & Ketterle, W. Bose-Einstein condensation of atomic gases. *Nature* **416**, 211–218 (2002).
45. Leggett, A. J. Bose-Einstein condensation in the alkali gases: Some fundamental concepts. *Reviews of modern physics* **73**, 307 (2001).
46. Stamper-Kurn, D. M. & Ueda, M. Spinor Bose gases: Symmetries, magnetism, and quantum dynamics. *Reviews of Modern Physics* **85**, 1191–1244 (2013).
47. Bao, W., Jaksch, D. & Markowich, P. A. Numerical solution of the Gross-Pitaevskii equation for Bose-Einstein condensation. *Journal of Computational Physics* **187**, 318–342 (2003).
48. Bao, W., Jin, S. & Markowich, P. A. Numerical Study of Time-Splitting Spectral Discretizations of Nonlinear Schrödinger Equations in the Semiclassical Regimes. *SIAM Journal on Scientific Computing* **25**, 27–64 (2003).
49. Bao, W., Jin, S. & Markowich, P. A. On Time-Splitting Spectral Approximations for the Schrödinger Equation in the Semiclassical Regime. *Journal of Computational Physics* **175**, 487–524 (2002).
50. Bao, W. & Cai, Y. Mathematical theory and numerical methods for Bose-Einstein condensation. *Kinetic & Related Models* **6**, 1–135 (2013).
51. Bao, W. & Shen, J. A Fourth-Order Time-Splitting Laguerre-Hermite Pseudospectral Method for Bose-Einstein Condensates. *SIAM Journal on Scientific Computing* **26**, 2010–2028 (2005).
52. Feynman, R. P., Hibbs, A. R. & Styer, D. F. *Quantum mechanics and path integrals* (Courier Corporation, 2010).
53. Owens, J. D., Luebke, D., Govindaraju, N., Harris, M., Krüger, J., Lefohn, A. E. & Purcell, T. J. A survey of general-purpose computation on graphics hardware in *Computer graphics forum* **26** (2007), 80–113.
54. Owens, J. D., Houston, M., Luebke, D., Green, S., Stone, J. E. & Phillips, J. C. GPU computing. *Proceedings of the IEEE* **96**, 879–899 (2008).
55. Lindholm, E., Kilgard, M. J. & Moreton, H. A user-programmable vertex engine in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques* (2001), 149–158.

56. Krüger, J. & Westermann, R. in *ACM SIGGRAPH 2005 Courses* 234–es (2005).
57. Galoppo, N., Govindaraju, N. K., Henson, M. & Manocha, D. *LU-GPU: Efficient algorithms for solving dense linear systems on graphics hardware* in *SC'05: Proceedings of the 2005 ACM/IEEE Conference on Supercomputing* (2005), 3–3.
58. Bolz, J., Farmer, I., Grinspun, E. & Schröder, P. Sparse matrix solvers on the GPU: conjugate gradients and multigrid. *ACM transactions on graphics (TOG)* **22**, 917–924 (2003).
59. Moreland, K. & Angel, E. *The FFT on a GPU* in *Proceedings of the ACM SIGGRAPH/EUROGRAPHICS conference on Graphics hardware* (2003), 112–119.
60. Purcell, T. J., Buck, I., Mark, W. R. & Hanrahan, P. in *ACM SIGGRAPH 2005 Courses* 268–es (2005).
61. Hable, J. & Rossignac, J. Blister: Gpu-based rendering of boolean combinations of free-form triangulated shapes. *ACM transactions on graphics (TOG)* **24**, 1024–1031 (2005).
62. Harris, C. R. *et al.* Array programming with NumPy. *Nature* **585**, 357–362 (2020).
63. Paszke, A. *et al.* Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019).
64. Berry, M. V. Quantal phase factors accompanying adiabatic changes. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* **392**, 45–57 (1984).
65. Wilczek, F. & Zee, A. Appearance of Gauge Structure in Simple Dynamical Systems. *Physical Review Letters* **52**, 2111–2114 (1984).
66. Anandan, J. Non-adiabatic non-abelian geometric phase. *Physics Letters A* **133**, 171–175 (1988).
67. Sugawa, S., Salces-Carcoba, F., Yue, Y., Putra, A. & Spielman, I. B. Wilson loop and Wilczek-Zee phase from a non-Abelian gauge field. *npj Quantum Information* **7**, 144 (2021).
68. Das, K. K. Measurement and significance of Wilson loops in synthetic gauge fields. *Physical Review A* **97**, 053620 (2018).
69. Pachos, J. & Zanardi, P. Quantum holonomies for quantum computing. *International Journal of Modern Physics B* **15**, 1257–1285 (2001).

70. Berry, M. The Geometric Phase. *Scientific American* **259**, 46–52 (1988).
71. Sjöqvist, E., Tong, D. M., Andersson, L. M., Hessmo, B., Johansson, M. & Singh, K. Non-adiabatic holonomic quantum computation. *New Journal of Physics* **14**, 103035 (2012).
72. Zhang, J., Kyaw, T. H., Filipp, S., Kwek, L.-C., Sjöqvist, E. & Tong, D. Geometric and holonomic quantum computation. *Physics Reports* **1027**, 1–53 (2023).
73. Aharonov, Y. & Bohm, D. Significance of Electromagnetic Potentials in the Quantum Theory. *Physical Review* **115**, 485–491 (1959).
74. Aharonov, Y. & Bohm, D. Further Considerations on Electromagnetic Potentials in the Quantum Theory. *Physical Review* **123**, 1511–1524 (1961).
75. Wilson, K. G. Confinement of quarks. *Physical Review D* **10**, 2445–2459 (1974).
76. Cheng, H. & Tsai, E.-C. Gauge invariance of the quantum Wilson loop. *Physical Review D* **36**, 3196–3202 (1987).
77. Aidelsburger, M., Nascimbene, S. & Goldman, N. Artificial gauge fields in materials and engineered systems. *Comptes Rendus Physique* **19**, 394–432 (2018).
78. Karp, R. L., Mansouri, F. & Rno, J. S. Product integral formalism and non-Abelian Stokes theorem. *Journal of Mathematical Physics* **40**, 6033–6043 (1999).
79. Uhlmann, A. Parallel transport and “quantum holonomy” along density operators. *Reports on Mathematical Physics* **24**, 229–240 (1986).
80. Galitski, V., Juzeliūnas, G. & Spielman, I. B. Artificial gauge fields with ultracold atoms. *Physics Today* **72**, 38–44 (2019).
81. Wiese, U. Ultracold quantum gases and lattice systems: quantum simulation of lattice gauge theories. *Annalen der Physik* **525**, 777–796 (2013).
82. Dalmonte, M. & Montangero, S. Lattice gauge theory simulations in the quantum information era. *Contemporary Physics* **57**, 388–412 (2016).
83. Zohar, E., Cirac, J. I. & Reznik, B. Quantum simulations of lattice gauge theories using ultracold atoms in optical lattices. *Reports on Progress in Physics* **79**, 014401 (2015).
84. Ohler, S., Kiefer-Emmanouilidis, M., Browaeys, A., Büchler, H. P. & Fleischhauer, M. Self-generated quantum gauge fields in arrays of Rydberg atoms. *New Journal of Physics* **24**, 023017 (2022).

85. Frölian, A., Chisholm, C. S., Neri, E., Cabrera, C. R., Ramos, R., Celi, A. & Tarruell, L. Realizing a 1D topological gauge theory in an optically dressed BEC. *Nature* **608**, 293–297 (2022).
86. Clark, L. W., Anderson, B. M., Feng, L., Gaj, A., Levin, K. & Chin, C. Observation of Density-Dependent Gauge Fields in a Bose-Einstein Condensate Based on Micromotion Control in a Shaken Two-Dimensional Lattice. *Physical Review Letters* **121**, 030402 (2018).
87. Barbiero, L., Schweizer, C., Aidelsburger, M., Demler, E., Goldman, N. & Grusdt, F. Coupling ultracold matter to dynamical gauge fields in optical lattices: From flux attachment to \mathbb{Z}_2 lattice gauge theories. *Science Advances* **5**, eaav7444 (2019).
88. Yang, Y., Zhen, B., Joannopoulos, J. D. & Soljačić, M. Non-Abelian generalizations of the Hofstadter model: spin–orbit-coupled butterfly pairs. *Light: Science & Applications* **9**, 177 (2020).
89. Hatano, N., Shirasaki, R. & Nakamura, H. Non-Abelian gauge field theory of the spin-orbit interaction and a perfect spin filter. *Physical Review A* **75**, 032107 (2007).
90. Schweizer, C., Grusdt, F., Berngruber, M., Barbiero, L., Demler, E., Goldman, N., Bloch, I. & Aidelsburger, M. Floquet approach to \mathbb{Z}_2 lattice gauge theories with ultracold atoms in optical lattices. *Nature Physics* **15**, 1168–1173 (2019).
91. Pinske, J. & Scheel, S. Symmetry-protected non-Abelian geometric phases in optical waveguides with nonorthogonal modes. *Physical Review A* **105**, 013507 (2022).
92. Neef, V., Pinske, J., Klauck, F., Teuber, L., Kremer, M., Ehrhardt, M., Heinrich, M., Scheel, S. & Szameit, A. Three-dimensional non-Abelian quantum holonomy. *Nature Physics* **19**, 30–34 (2023).
93. Bakke, K. Geometric quantum phase from the interaction of the permanent electric dipole moment of a neutral particle with an axial magnetic field. *The European Physical Journal Plus* **137**, 333 (2022).
94. Liberto, M. D., Goldman, N. & Palumbo, G. Non-Abelian Bloch oscillations in higher-order topological insulators. *Nature Communications* **11**, 5942 (2020).
95. Ruseckas, J., Juzeliūnas, G., Öhberg, P. & Fleischhauer, M. Non-Abelian Gauge Potentials for Ultracold Atoms with Degenerate Dark States. *Physical Review Letters* **95**, 010404 (2005).

96. Osterloh, K., Baig, M., Santos, L., Zoller, P. & Lewenstein, M. Cold Atoms in Non-Abelian Gauge Potentials: From the Hofstadter "Moth" to Lattice Gauge Theory. *Physical Review Letters* **95**, 010403 (2005).
97. Goldman, N., Kubasiak, A., Gaspard, P. & Lewenstein, M. Ultracold atomic gases in non-Abelian gauge potentials: The case of constant Wilson loop. *Physical Review A* **79**, 023624 (2009).
98. Phuc, N. T., Tataru, G., Kawaguchi, Y. & Ueda, M. Controlling and probing non-abelian emergent gauge potentials in spinor Bose-Fermi mixtures. *Nature Communications* **6**, 8135 (2015).
99. Hauke, P. *et al.* Non-Abelian Gauge Fields and Topological Insulators in Shaken Optical Lattices. *Physical Review Letters* **109**, 145301 (2012).
100. Hu, H. & Zhang, C. Spin-1 topological monopoles in the parameter space of ultracold atoms. *Physical Review A* **98**, 013627 (2018).
101. Zheng, W. *et al.* Measuring Quantum Geometric Tensor of Non-Abelian System in Superconducting Circuits. *Chinese Physics Letters* **39**, 100202 (2022).
102. Bermudez, A., Goldman, N., Kubasiak, A., Lewenstein, M & Martin-Delgado, M. A. Topological phase transitions in the non-Abelian honeycomb lattice. *New Journal of Physics* **12**, 033041 (2010).
103. Pietilä, V. & Möttönen, M. Non-Abelian Magnetic Monopole in a Bose-Einstein Condensate. *Physical Review Letters* **102**, 080403 (2008).
104. González-Cuadra, D., Zache, T. V., Carrasco, J., Kraus, B. & Zoller, P. Hardware Efficient Quantum Simulation of Non-Abelian Gauge Theories with Qudits on Rydberg Platforms. *Physical Review Letters* **129**, 160501 (2022).
105. Pastori, L., Olsacher, T., Kokail, C. & Zoller, P. Characterization and Verification of Trotterized Digital Quantum Simulation via Hamiltonian and Liouvillean Learning. *arXiv* (2022).
106. Weisbrich, H., Rastelli, G. & Belzig, W. Geometrical Rabi oscillations and Landau-Zener transitions in non-Abelian systems. *Physical Review Research* **3**, 033122 (2021).
107. Leroux, F., Pandey, K., Rehbi, R., Chevy, F., Miniatura, C., Grémaud, B. & Wilkowski, D. Non-Abelian adiabatic geometric transformations in a cold strontium gas. *Nature Communications* **9**, 3580 (2018).
108. H. M., B., Boguslawski, M., Barrios, M., Xin, L. & Chapman, M. S. Exploring Non-Abelian Geometric Phases in Spin-1 Ultracold Atoms. *Physical Review Letters* **123**, 173202 (2019).

109. Yang, Y., Peng, C., Zhu, D., Buljan, H., Joannopoulos, J. D., Zhen, B. & Soljačić, M. Synthesis and observation of non-Abelian gauge fields in real space. *Science* **365**, 1021–1025 (2019).
110. Huang, L., Meng, Z., Wang, P., Peng, P., Zhang, S.-L., Chen, L., Li, D., Zhou, Q. & Zhang, J. Experimental realization of two-dimensional synthetic spin-orbit coupling in ultracold Fermi gases. *Nature Physics* **12**, 540–544 (2016).
111. DiVincenzo, D. P. Quantum Computation. *Science* **270**, 255–261 (1995).
112. Ladd, T. D., Jelezko, F., Laflamme, R., Nakamura, Y., Monroe, C. & O’Brien, J. L. Quantum computers. *Nature* **464**, 45–53 (2010).
113. Steane, A. Quantum computing. *Reports on Progress in Physics* **61**, 117 (1998).
114. Benioff, P. The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines. *Journal of Statistical Physics* **22**, 563–591 (1980).
115. Benioff, P. Quantum Mechanical Models of Turing Machines That Dissipate No Energy. *Physical Review Letters* **48**, 1581–1585 (1982).
116. Deutsch, D. E. Quantum computational networks. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* **425**, 73–90 (1989).
117. Bravyi, S., Gosset, D. & König, R. Quantum advantage with shallow circuits. *Science* **362**, 308–311 (2018).
118. Bravyi, S., Gosset, D., Koenig, R. & Tomamichel, M. Quantum advantage with noisy shallow circuits. *Nature Physics* **16**, 1040–1045 (2020).
119. Daley, A. J., Bloch, I., Kokail, C., Flannigan, S., Pearson, N., Troyer, M. & Zoller, P. Practical quantum advantage in quantum simulation. *Nature* **607**, 667–676 (2022).
120. Devitt, S. J., Munro, W. J. & Nemoto, K. Quantum error correction for beginners. *Reports on Progress in Physics* **76**, 076001 (2013).
121. Zanardi, P. & Rasetti, M. Holonomic quantum computation. *Physics Letters A* **264**, 94–99 (1999).
122. Pachos, J., Zanardi, P. & Rasetti, M. Non-Abelian Berry connections for quantum computation. *Physical Review A* **61**, 010305 (2000).
123. Colmenar, R. K. L., Güngördü, U. & Kestner, J. P. Conditions for Equivalent Noise Sensitivity of Geometric and Dynamical Quantum Gates. *PRX Quantum* **3**, 030310 (2022).

124. Carollo, A., Fuentes-Guridi, I., Santos, M. F. & Vedral, V. Geometric phase in open systems. *Physical Review Letters* **90** (2003).
125. Chiara, G. D. & Palma, G. M. Berry Phase for a Spin 1/2 Particle in a Classical Fluctuating Field. *Physical Review Letters* **91**, 090404 (2003).
126. Fuentes-Guridi, I., Girelli, F. & Livine, E. Holonomic quantum computation in the presence of decoherence. *Physical Review Letters* **94** (2005).
127. Wu, L. A., Zanardi, P. & Lidar, D. A. Holonomic quantum computation in decoherence-free subspaces. *Physical Review Letters* **95** (2005).
128. Solinas, P., Zanardi, P. & Zanghì, N. Robustness of non-Abelian holonomic quantum gates against parametric noise. *Physical Review A* **70**, 042316 (2004).
129. Zheng, S.-B., Yang, C.-P. & Nori, F. Comparison of the sensitivity to systematic errors between nonadiabatic non-Abelian geometric gates and their dynamical counterparts. *Physical Review A* **93**, 032313 (2016).
130. Kang, Y.-H., Song, J. & Xia, Y. Error-resistant nonadiabatic binomial-code geometric quantum computation using reverse engineering. *Optics Letters* **47**, 4099 (2022).
131. Chen, T., Xue, Z.-Y. & Wang, Z. Error-Tolerant Geometric Quantum Control for Logical Qubits with Minimal Resources. *Physical Review Applied* **18**, 014062 (2022).
132. Pachos, J. Quantum Computation by Geometrical Means. *arXiv* (2000).
133. Chen, D., Luo, S., Wang, Y.-D., Chesi, S. & Choi, M.-S. Geometric manipulation of a decoherence-free subspace in atomic ensembles. *Physical Review A* **105**, 022627 (2022).
134. Florio, G., Facchi, P., Fazio, R., Giovannetti, V. & Pascazio, S. Robust gates for holonomic quantum computation. *Physical Review A* **73**, 022327 (2006).
135. Oreshkov, O., Brun, T. A. & Lidar, D. A. Fault-Tolerant Holonomic Quantum Computation. *Physical Review Letters* **102**, 070502 (2008).
136. Pachos, J. & Chountasis, S. Optical holonomic quantum computer. *Physical Review A* **62**, 052318 (2000).
137. Pinske, J., Teuber, L. & Scheel, S. Highly degenerate photonic waveguide structures for holonomic computation. *Physical Review A* **101**, 062314 (2020).
138. Faoro, L., Siewert, J. & Fazio, R. Non-Abelian Holonomies, Charge Pumping, and Quantum Computation with Josephson Junctions. *Physical Review Letters* **90**, 028301 (2003).

139. Bernevig, B. A. & Zhang, S.-C. Holonomic quantum computing based on the Stark effect. *Physical Review B* **71**, 035303 (2005).
140. Golovach, V. N., Borhani, M. & Loss, D. Holonomic quantum computation with electron spins in quantum dots. *Physical Review A* **81**, 022315 (2010).
141. Wang, Y., Zhang, J., Wu, C., You, J. Q. & Romero, G. Holonomic quantum computation in the ultrastrong-coupling regime of circuit QED. *Physical Review A* **94**, 012328 (2016).
142. Li, Y., Zhang, P., Zanardi, P. & Sun, C. P. Non-Abelian geometric quantum memory with an atomic ensemble. *Physical Review A* **70**, 032330 (2004).
143. Jones, J. A., Vedral, V., Ekert, A. & Castagnoli, G. Geometric quantum computation using nuclear magnetic resonance. *Nature* **403**, 869–871 (2000).
144. Duan, L.-M., Cirac, J. I. & Zoller, P. Geometric Manipulation of Trapped Ions for Quantum Computation. *Science* **292**, 1695–1697 (2001).
145. Kremer, M., Teuber, L., Szameit, A. & Scheel, S. Optimal design strategy for non-Abelian geometric phases using Abelian gauge fields based on quantum metric. *Physical Review Research* **1**, 033117 (2019).
146. Zu, C., Wang, W. B., He, L., Zhang, W. G., Dai, C. Y., Wang, F. & Duan, L. M. Experimental realization of universal geometric quantum gates with solid-state spins. *Nature* **514**, 72–75 (2014).
147. Sekiguchi, Y., Niikura, N., Kuroiwa, R., Kano, H. & Kosaka, H. Optical holonomic single quantum gates with a geometric spin under a zero field. *Nature Photonics* **11**, 309–314 (2017).
148. Recati, A., Calarco, T., Zanardi, P., Cirac, J. I. & Zoller, P. Holonomic quantum computation with neutral atoms. *Physical Review A* **66**, 032309 (2002).
149. Toyoda, K., Uchida, K., Noguchi, A., Haze, S. & Urabe, S. Realization of holonomic single-qubit operations. *Physical Review A* **87**, 052307 (2013).
150. Shen, P., Chen, T. & Xue, Z.-Y. Ultrafast Holonomic Quantum Gates. *Physical Review Applied* **16**, 044004 (2021).
151. Alves, G. O. & Sjöqvist, E. Time-optimal holonomic quantum computation. *Physical Review A* **106**, 032406 (2022).
152. Liu, B.-J., Huang, Z.-H., Xue, Z.-Y. & Zhang, X.-D. Superadiabatic holonomic quantum computation in cavity QED. *Physical Review A* **95**, 062308 (2017).
153. Xu, G. F., Zhao, P. Z., Xing, T. H., Sjöqvist, E. & Tong, D. M. Composite nonadiabatic holonomic quantum computation. *Physical Review A* **95**, 032311 (2017).

154. Setiawan, F., Groszkowski, P. & Clerk, A. A. Fast and Robust Geometric Two-Qubit Gates for Superconducting Qubits and beyond. *Physical Review Applied* **19**, 034071 (2023).
155. Su, S.-L., Sun, L.-N., Liu, B.-J., Yan, L.-L., Yung, M.-H., Li, W. & Feng, M. Rabi- and Blockade-Error-Resilient All-Geometric Rydberg Quantum Gates. *Physical Review Applied* **19**, 044007 (2023).
156. André, T. & Sjöqvist, E. Dark path holonomic qudit computation. *Physical Review A* **106**, 062402 (2022).
157. Kang, Y.-H., Xiao, Y., Shi, Z.-C., Wang, Y., Yang, J.-Q., Song, J. & Xia, Y. Effective implementation of nonadiabatic geometric quantum gates of cat-state qubits using an auxiliary qutrit. *New Journal of Physics* **25**, 033029 (2023).
158. Wang, S., Ni, W.-S., Chen, Z.-W., Feng, B., Ning, Y., Kang, Y.-H. & Xia, Y. Nonadiabatic geometric quantum computation of W-state codes using invariant-based reverse engineering. *Laser Physics Letters* **20**, 015203 (2023).
159. Xu, G. F., Zhang, J., Tong, D. M., Sjöqvist, E. & Kwek, L. C. Nonadiabatic Holonomic Quantum Computation in Decoherence-Free Subspaces. *Physical Review Letters* **109**, 170501 (2012).
160. Kang, Y.-H., Chen, Y.-H., Wang, X., Song, J., Xia, Y., Miranowicz, A., Zheng, S.-B. & Nori, F. Nonadiabatic geometric quantum computation with cat-state qubits via invariant-based reverse engineering. *Physical Review Research* **4**, 013233 (2022).
161. Zhao, P. Z., Li, K. Z., Xu, G. F. & Tong, D. M. General approach for constructing Hamiltonians for nonadiabatic holonomic quantum computation. *Physical Review A* **101**, 062306 (2020).
162. Xu, G. F. & Tong, D. M. Realizing multi-qubit controlled nonadiabatic holonomic gates with connecting systems. *AAPPS Bulletin* **32**, 13 (2022).
163. Liang, Y., Shen, P., Ji, L.-N. & Xue, Z.-Y. State-Independent Nonadiabatic Geometric Quantum Gates. *Physical Review Applied* **19**, 024051 (2023).
164. Tang, G., Yang, X.-Y., Yan, Y. & Lu, J. Fast evolution of single qubit gate in non-adiabatic geometric quantum computing. *Physics Letters A* **449**, 128349 (2022).
165. Li, S., Xue, J., Chen, T. & Xue, Z. High-Fidelity Geometric Quantum Gates with Short Paths on Superconducting Circuits. *Advanced Quantum Technologies* **4**, 2000140 (2021).

166. Ding, C.-Y., Ji, L.-N., Chen, T. & Xue, Z.-Y. Path-optimized nonadiabatic geometric quantum computation on superconducting qubits. *Quantum Science & Technology* **7**, 015012 (2022).
167. Liang, Y., Shen, P., Chen, T. & Xue, Z.-Y. Composite Short-Path Nonadiabatic Holonomic Quantum Gates. *Physical Review Applied* **17**, 034015 (2022).
168. Ji, L.-N., Liang, Y., Shen, P. & Xue, Z.-Y. Nonadiabatic Holonomic Quantum Computation via Path Optimization. *Physical Review Applied* **18**, 044034 (2022).
169. Cheng, J.-j. & Zhang, L. Implementing conventional and unconventional nonadiabatic geometric quantum gates via SU(2) transformations. *Physical Review A* **103**, 032616 (2021).
170. Wei, J.-F., Guo, F.-Q., Wang, D.-Y., Jia, Y., Yan, L.-L., Feng, M. & Su, S.-L. Fast multiqubit Rydberg geometric fan-out gates with optimal control technology. *Physical Review A* **105**, 042404 (2022).
171. Zhao, P. Z., Wu, X. & Tong, D. M. Dynamical-decoupling-protected nonadiabatic holonomic quantum computation. *Physical Review A* **103**, 012205 (2021).
172. Li, S. & Xue, Z.-Y. Dynamically Corrected Nonadiabatic Holonomic Quantum Gates. *Physical Review Applied* **16**, 044005 (2021).
173. Wu, X. & Zhao, P. Z. Nonadiabatic geometric quantum computation protected by dynamical decoupling via the XXZ Hamiltonian. *Frontiers of Physics* **17**, 31502 (2022).
174. Liang, M.-J. & Xue, Z.-Y. Robust nonadiabatic geometric quantum computation by dynamical correction. *Physical Review A* **106**, 012603 (2022).
175. Guo, L.-J., Xu, H., Fang, Z.-Y., Chen, T., Wei, K. & Zhang, C. Optimizing nonadiabatic geometric quantum gates against off-resonance error in a silicon-based spin qubit. *Physical Review A* **107**, 012604 (2023).
176. Ji, L., Ding, C., Chen, T. & Xue, Z. Noncyclic Geometric Quantum Gates with Smooth Paths via Invariant-Based Shortcuts. *Advanced Quantum Technologies* **4**, 2100019 (2021).
177. Zhou, J., Li, S., Pan, G.-Z., Zhang, G., Chen, T. & Xue, Z.-Y. Nonadiabatic geometric quantum gates that are insensitive to qubit-frequency drifts. *Physical Review A* **103**, 032609 (2021).
178. Mommers, C. J. G. & Sjöqvist, E. Universal quantum computation and quantum error correction using discrete holonomies. *Physical Review A* **105**, 022402 (2022).

179. Li, K. Z., Xu, G. F. & Tong, D. M. Coherence-protected nonadiabatic geometric quantum computation. *Physical Review Research* **3**, 023104 (2021).
180. Liu, B.-J., Yan, L.-L., Zhang, Y., Yung, M.-H., Liang, E., Su, S.-L. & Shan, C.-X. Decoherence-suppressed nonadiabatic holonomic quantum computation. *Physical Review Research* **5**, 013059 (2023).
181. Kang, Y.-H., Shi, Z.-C., Song, J. & Xia, Y. Effective non-adiabatic holonomic quantum computation of cavity modes via invariant-based reverse engineering. *Philosophical Transactions of the Royal Society A* **380**, 20210279 (2022).
182. Zhou, J., Liu, B., Hong, Z. & Xue, Z. Fast holonomic quantum computation based on solid-state spins with all-optical control. *Science China Physics, Mechanics & Astronomy* **61**, 010312 (2017).
183. Yun, M.-R., Guo, F.-Q., Yan, L.-L., Liang, E., Zhang, Y., Su, S.-L., Shan, C. X. & Jia, Y. Parallel-path implementation of nonadiabatic geometric quantum gates in a decoherence-free subspace with nitrogen-vacancy centers. *Physical Review A* **105**, 012611 (2022).
184. Hong, Z.-P., Liu, B.-J., Cai, J.-Q., Zhang, X.-D., Hu, Y., Wang, Z. D. & Xue, Z.-Y. Implementing universal nonadiabatic holonomic quantum gates with transmons. *Physical Review A* **97**, 022332 (2018).
185. Chen, T., Zhang, J. & Xue, Z.-Y. Nonadiabatic holonomic quantum computation on coupled transmons with ancillaries. *Physical Review A* **98**, 052314 (2018).
186. Xue, Z.-Y., Zhou, J., Chu, Y.-M. & Hu, Y. Nonadiabatic holonomic quantum computation with all-resonant control. *Physical Review A* **94**, 022331 (2016).
187. Xue, Z.-Y., Gu, F.-L., Hong, Z.-P., Yang, Z.-H., Zhang, D.-W., Hu, Y. & You, J. Q. Nonadiabatic Holonomic Quantum Computation with Dressed-State Qubits. *Physical Review Applied* **7**, 054022 (2017).
188. Li, M., Li, J., Guo, F., Zhu, X., Liang, E., Zhang, S., Yan, L., Feng, M. & Su, S. Multiple-Qubit CkUm Logic Gates of Rydberg Atoms via Optimized Geometric Quantum Operations. *Annalen der Physik* **534**, 2100506 (2022).
189. Kang, Y.-H., Chen, Y.-H., Shi, Z.-C., Huang, B.-H., Song, J. & Xia, Y. Nonadiabatic holonomic quantum computation using Rydberg blockade. *Physical Review A* **97**, 042336 (2018).
190. Zhao, P. Z., Wu, X., Xing, T. H., Xu, G. F. & Tong, D. M. Nonadiabatic holonomic quantum computation with Rydberg superatoms. *Physical Review A* **98**, 032313 (2018).

191. Liu, B., Guo, F., Yan, L., Zhang, S., Feng, M. & Su, S. Realization of Deutsch–Jozsa Algorithm in Rydberg Atoms by Composite Nonadiabatic Holonomic Quantum Computation with Strong Robustness Against Systematic Errors. *Advanced Quantum Technologies* **4**, 2100093 (2021).
192. Zhao, P. Z., Cui, X.-D., Xu, G. F., Sjöqvist, E. & Tong, D. M. Rydberg-atom-based scheme of nonadiabatic geometric quantum computation. *Physical Review A* **96**, 052316 (2017).
193. Feng, G., Xu, G. & Long, G. Experimental Realization of Nonadiabatic Holonomic Quantum Computation. *Physical Review Letters* **110**, 190501 (2013).
194. Zhu, Z., Chen, T., Yang, X., Bian, J., Xue, Z.-Y. & Peng, X. Single-Loop and Composite-Loop Realization of Nonadiabatic Holonomic Quantum Gates in a Decoherence-Free Subspace. *Physical Review Applied* **12**, 024024 (2019).
195. Li, H., Liu, Y. & Long, G. Experimental realization of single-shot nonadiabatic holonomic gates in nuclear spins. *Science China Physics, Mechanics & Astronomy* **60**, 080311 (2017).
196. Ai, M.-Z., Li, S., He, R., Xue, Z.-Y., Cui, J.-M., Huang, Y.-F., Li, C.-F. & Guo, G.-C. Experimental realization of nonadiabatic holonomic single-qubit quantum gates with two dark paths in a trapped ion. *Fundamental Research* **2**, 661–666 (2022).
197. Nagata, K., Kuramitani, K., Sekiguchi, Y. & Kosaka, H. Universal holonomic quantum gates over geometric spin qubits with polarised microwaves. *Nature Communications* **9**, 3227 (2018).
198. Arroyo-Camejo, S., Lazariev, A., Hell, S. W. & Balasubramanian, G. Room temperature high-fidelity holonomic single-qubit gate on a solid-state spin. *Nature Communications* **5**, 4870 (2014).
199. Ishida, N., Nakamura, T., Tanaka, T., Mishima, S., Kano, H., Kuroiwa, R., Sekiguchi, Y. & Kosaka, H. Universal holonomic single quantum gates over a geometric spin with phase-modulated polarized light. *Optics Letters* **43**, 2380 (2018).
200. Shui, H., Jin, S., Li, Z., Wei, F., Chen, X., Li, X. & Zhou, X. Atom-orbital qubit under nonadiabatic holonomic quantum control. *Physical Review A* **104**, L060601 (2021).
201. Xu, J.-Z., Sun, L.-N., Wei, J.-F., Du, Y.-L., Luo, R., Yan, L.-L., Feng, M. & Su, S.-L. Two-Qubit Geometric Gates Based on Ground-State Blockade of Rydberg Atoms. *Chinese Physics Letters* **39**, 090301 (2022).

202. Abdumalikov, A. A., Fink, J. M., Juliusson, K., Pechal, M., Berger, S., Wallraff, A. & Filipp, S. Experimental realization of non-Abelian non-adiabatic geometric gates. *Nature* **496**, 482–485 (2013).
203. Li, Y. *et al.* Nonadiabatic Geometric Gates with a Shortened Loop in a Superconducting Circuit. *physica status solidi (b)*, 2200040 (2022).
204. Yan, T. *et al.* Experimental Realization of Nonadiabatic Shortcut to Non-Abelian Geometric Gates. *Physical Review Letters* **122** (2019).
205. Yang, X.-X. *et al.* Experimental Implementation of Short-Path Nonadiabatic Geometric Gates in a Superconducting Circuit. *Physical Review Applied* **19**, 044076 (2023).
206. Albash, T. & Lidar, D. A. Adiabatic quantum computation. *Reviews of Modern Physics* **90**, 015002 (2018).
207. Das, A. & Chakrabarti, B. K. Colloquium: Quantum annealing and analog quantum computation. *Reviews of Modern Physics* **80**, 1061–1081 (2008).
208. Nayak, C., Simon, S. H., Stern, A., Freedman, M. & Sarma, S. D. Non-Abelian anyons and topological quantum computation. *Reviews of Modern Physics* **80**, 1083–1159 (2008).
209. Stern, A. & Lindner, N. H. Topological Quantum Computation—From Basic Concepts to First Experiments. *Science* **339**, 1179–1184 (2013).
210. Lahtinen, V. & Pachos, J. A Short Introduction to Topological Quantum Computation. *SciPost Physics* **3**, 021 (2017).
211. Sarma, S. D., Freedman, M. & Nayak, C. Topological quantum computation. *Physics Today* **59**, 32–38 (2006).
212. Pachos, J. K. *Introduction to topological quantum computation* (Cambridge University Press, 2012).
213. Kalinowski, M., Maskara, N. & Lukin, M. D. Non-Abelian Floquet Spin Liquids in a Digital Rydberg Simulator. *Physical Review X* **13**, 031008 (2023).
214. Rahav, S., Gilary, I. & Fishman, S. Effective Hamiltonians for periodically driven systems. *Physical Review A* **68**, 013820 (2003).
215. Goldman, N. & Dalibard, J. Periodically Driven Quantum Systems: Effective Hamiltonians and Engineered Gauge Fields. *Physical Review X* **4**, 031027 (2014).
216. Sambe, H. Steady States and Quasienergies of a Quantum-Mechanical System in an Oscillating Field. *Physical Review A* **7**, 2203–2213 (1973).

217. Novičenko, V., Anisimovas, E. & Juzeliūnas, G. Floquet analysis of a quantum system with modulated periodic driving. *Physical Review A* **95**, 023615 (2017).
218. Novičenko, V. & Juzeliūnas, G. Non-Abelian geometric phases in periodically driven systems. *Physical Review A* **100**, 012127 (2019).
219. Chen, Z., Murphree, J. D. & Bigelow, N. P. SU(2) geometric phase induced by a periodically driven Raman process in an ultracold dilute Bose gas. *Physical Review A* **101**, 013606 (2020).
220. Wang, Y.-S., Liu, B.-J., Su, S.-L. & Yung, M.-H. Error-resilient Floquet geometric quantum computation. *Physical Review Research* **3**, 033010 (2021).
221. Bezanson, J., Edelman, A., Karpinski, S. & Shah, V. B. Julia: A fresh approach to numerical computing. *SIAM Review* **59**, 65–98 (2017).
222. Rackauckas, C. & Nie, Q. DifferentialEquations.jl – A Performant and Feature-Rich Ecosystem for Solving Differential Equations in Julia. *The Journal of Open Research Software* **5** (2017).
223. Curtright, T. L., Fairlie, D. B. & Zachos, C. K. A Compact Formula for Rotations as Spin Matrix Polynomials. *Symmetry, Integrability and Geometry: Methods and Applications* (2014).
224. Mogensen, P. K. & Riseth, A. N. Optim: A mathematical optimization package for Julia. *Journal of Open Source Software* **3**, 615 (2018).
225. Shukla, A., Sisodia, M. & Pathak, A. Complete characterization of the directly implementable quantum gates used in the IBM quantum processors. *Physics Letters A* **384**, 126387 (2020).
226. Budker, D. & Romalis, M. Optical magnetometry. *Nature physics* **3**, 227–234 (2007).
227. Kimball, D. J., Alexandrov, E. & Budker, D. General principles and characteristics of optical magnetometers. *Optical Magnetometry* **3** (2013).
228. Xu, X.-T., Wang, Z.-Y., Jiao, R.-H., Yi, C.-R., Sun, W. & Chen, S. Ultra-low noise magnetic field for quantum gases. *Review of Scientific Instruments* **90**, 054708 (2019).
229. Borkowski, M., Reichsöllner, L., Thekkepatt, P., Barbé, V., Roon, T. v., Druten, K. v. & Schreck, F. Active stabilization of kilogauss magnetic fields to the ppm level for magnetoassociation on ultranarrow Feshbach resonances. *Review of Scientific Instruments* **94**, 073202 (2023).

APPENDIX A

Quadratic Zeeman Effect

ZEEMAN shifts of magnetic sublevels in alkali-atoms are essential to all of the work presented in this thesis, although the linear form we utilized (Eq. 2.6) is only approximate. The next leading-order correction to the energies of the m_F levels is given by the quadratic Zeeman shift. The Hamiltonian for the quadratic shift ϵ is,

$$\hat{H}_\epsilon = \epsilon \left(\hat{\mathbf{1}} - \hat{F}_z^2 / \hbar \right), \quad (\text{A.1})$$

in terms of the total angular momentum operator \hat{F}_z , assuming the applied magnetic field is $\mathbf{B} = B_z \mathbf{e}_z$ [29, 46]. The magnitude of the quadratic shift in terms of the field strength is,

$$|\epsilon| = \frac{(g_s \mu_B - g_I \mu_N)^2}{\omega_{\text{HF}} (1 - 2I)^2} B_z^2, \quad (\text{A.2})$$

where $g_s \approx 2$ is the electron g -factor, $g_I \mu_N / \hbar$ is the nuclear gyromagnetic ratio, $\hbar \omega_{\text{HF}}$ is the hyperfine splitting, and I is the angular momentum quantum number of the nucleus.

From the matrix form of \hat{H}_ϵ , we can see that, unlike the linear Zeeman effect, the quadratic effect shifts m_F levels in the *same* direction; levels are shifted according to $|m_F|$. For this reason, \hat{H}_ϵ breaks the symmetric splitting between m_F levels, as some are shifted closer together than others. If $F = 1/2$ then the Hamiltonian vanishes, and hence there is no such shift in spin-1/2 systems.

In the case of ^{87}Rb , we may express ϵ in terms of the linear Zeeman shift $\omega_Z = \omega_B B_z$, with $\omega_B \approx 0.7$ MHz/G. Substituting in the other values, $\mu_B / \hbar = 1.4$ MHz/G, $\mu_N / \hbar \approx 7.7 \times 10^{-4}$ MHz/G, $g_I \approx 1 \times 10^{-3}$, $I = 3/2$, and $\omega_{\text{HF}} = 6.8$ GHz, we find,

$$|\epsilon| / \hbar = (1.5 \times 10^{-7} \text{ kHz}^{-1}) \omega_Z^2, \quad (\text{A.3})$$

which is useful in quickly estimating the size of this effect for a given Zeeman splitting.

APPENDIX B

Rotating Basis Transformations of Spin Operators

TRANSFORMATIONS between the *lab* and *rotating* frames are a crucial part of the work presented in this thesis, in order to remove rapidly rotating terms in the Hamiltonian due to the fast carrier frequencies of driving fields (see Sec. 2.2.1 and Sec. 2.2.2). A key piece in making these basis transformations easier to compute is utilizing the raising and lowering momentum operators, which have a simple transformation under the rotating basis unitary. Here we will compute these transformations from the Baker-Campbell-Hausdorff lemma,

$$e^{\hat{A}}\hat{B}e^{-\hat{A}} = \hat{B} + [\hat{A}, \hat{B}] + \frac{1}{2!} [\hat{A}, [\hat{A}, \hat{B}]] + \dots \quad (\text{B.1})$$

The rotating wave unitary transformation has the form,

$$\hat{U} = \exp\left[-\frac{i\omega(t)\hat{F}_z}{\hbar}\right]. \quad (\text{B.2})$$

We are interested in how \hat{F}_z and \hat{F}_\pm transform under this unitary. Starting with \hat{F}_z we have,

$$\hat{U}^\dagger \hat{F}_z \hat{U} = \hat{F}_z + \frac{i\omega}{\hbar} [\hat{F}_z, \hat{F}_z] + \dots, \quad (\text{B.3})$$

$$= \hat{F}_z, \quad (\text{B.4})$$

since \hat{U} commutes with \hat{F}_z .

For \hat{F}_\pm we recall that $[\hat{F}_z, \hat{F}_\pm] = \pm\hbar\hat{F}_\pm$. From this we have,

$$\hat{U}^\dagger \hat{F}_\pm \hat{U} = \hat{F}_\pm + \frac{i\omega}{\hbar} [\hat{F}_z, \hat{F}_\pm] + \frac{(i\omega)^2}{2\hbar^2} [\hat{F}_z, [\hat{F}_z, \hat{F}_\pm]] + \dots, \quad (\text{B.5})$$

$$= \hat{F}_\pm + \frac{i\omega}{\hbar} (\pm\hbar\hat{F}_\pm) + \frac{(i\omega)^2}{2\hbar^2} (\hbar^2\hat{F}_\pm) + \dots, \quad (\text{B.6})$$

$$= \hat{F}_\pm \left[1 + (\pm i\omega) + \frac{(\pm i\omega)^2}{2!} + \dots \right], \quad (\text{B.7})$$

$$= e^{\pm i\omega} \hat{F}_\pm, \quad (\text{B.8})$$

hence, we have the necessary relations.

APPENDIX C

Spin-State Tomography

STATE tomography is a process for reconstructing the states of a quantum system. In this case, we will consider spin states, which are superpositions of the eigenstates of angular momentum operators. For spins in a manifold with angular momentum quantum number F , an arbitrary state may be expressed as,

$$|\psi\rangle = \frac{1}{N} \sum_{f=-F}^F \tilde{c}_f |F, m_F = f\rangle, \quad (\text{C.1})$$

where \tilde{c}_f are complex coefficients, and N is the normalization factor. The goal in state tomography is to determine the coefficients \tilde{c}_f , typically in as few measurements as possible.

The measurements considered here will consist of projections onto the eigenstates of various spin operators; in each case, the relative populations in each eigenstate must be determined, which requires good statistics. The measurements must therefore be performed many times over, and the same state $|\psi\rangle$ must be prepared before each measurement, as they are destructive. This is not realistic for many systems, but spinor BECs have an advantage here because a single projective measurement may be applied to the entire atomic ensemble simultaneously, giving the relative populations directly. We are therefore only need to change measurement bases, and are not required to repeat each one many times to gather statistics.

Here we will describe the process for a spin-1/2 system, as it is the simplest to consider. Then, we will compute the results for spin-1, finding that more measurements are needed.

C.0.1 Spin-1/2

An arbitrary spin-1/2 superposition, expressed through the eigenstates of \hat{F}_z , $\{|\psi_1^z\rangle, |\psi_2^z\rangle\}$, is

$$|\psi\rangle = a |\psi_1^z\rangle + b e^{i\phi} |\psi_2^z\rangle, \quad (\text{C.2})$$

where $a, b \in \mathbb{R}$, and ϕ is the relative phase between states. We have placed all of the relative phase on the second eigenstate, which we are free to do by gauge symmetry. The notation used here is perhaps excessive for spin-1/2, but will be useful later for higher spin systems.

We start by taking the state $|\psi\rangle$ and projecting it into the eigenbases of \hat{F}_x , \hat{F}_y , and \hat{F}_z , which are denoted as $\{|\psi_1^q\rangle, |\psi_2^q\rangle\}$ with $q \in \{x, y, z\}$. We define the population in the j^{th} eigenstate of \hat{F}_q as $q_j \equiv |\langle \psi_j^q | \psi \rangle|^2$. The populations in the x -basis are,

$$x_1 = \frac{1}{2} (a^2 + b^2 + 2ab \cos \phi), \quad (\text{C.3})$$

$$x_2 = \frac{1}{2} (a^2 + b^2 - 2ab \cos \phi). \quad (\text{C.4})$$

From the normalization of $|\psi\rangle$, we can simplify, and taking the difference $x_1 - x_2$ we find,

$$\cos \phi = \frac{x_1 - x_2}{2ab}. \quad (\text{C.5})$$

Similarly for the y basis,

$$y_1 = \frac{1}{2} (a^2 + b^2 + 2ab \sin \phi), \quad (\text{C.6})$$

$$y_2 = \frac{1}{2} (a^2 + b^2 - 2ab \sin \phi). \quad (\text{C.7})$$

Once again, taking the difference $y_1 - y_2$ and simplifying we find,

$$\sin \phi = \frac{y_1 - y_2}{2ab}. \quad (\text{C.8})$$

Last, the z populations are,

$$z_1 = a^2, \quad (\text{C.9})$$

$$z_2 = b^2. \quad (\text{C.10})$$

The state coefficients are therefore known only up to a sign, $a = \pm\sqrt{z_1}$ and $b = \pm\sqrt{z_2}$. Thankfully, we may assume that $a, b \geq 0$, since if they have the same sign it is a global phase, and if they differ in sign this amounts to a relative phase of π ,

which we can capture in ϕ . Therefore, the state coefficients are just,

$$a = \sqrt{z_1}, \quad (\text{C.11})$$

$$b = \sqrt{z_2}. \quad (\text{C.12})$$

We can solve for the phase ϕ by taking the ratio of the Eqs. C.5 and C.8,

$$\phi = \arctan\left(\frac{y_1 - y_2}{x_1 - x_2}\right). \quad (\text{C.13})$$

Hence, from the populations measured in each basis, we may fully reconstruct each parameter of an arbitrary spin superposition. Here, we have chosen to project the state into the eigenbases of the typical spin operators, but this choice is arbitrary; we would find the same results with any three orthogonal spin matrices.

C.0.2 Spin-1

Following the same procedure as above (Sec. C.0.1), we may express an arbitrary spin-1 superposition as,

$$|\psi\rangle = a |\psi_1^z\rangle + b e^{i\phi_b} |\psi_2^z\rangle + c e^{i\phi_c} |\psi_3^z\rangle \quad (\text{C.14})$$

in the eigenbasis of \hat{F}_z . Once again, we have factored out the phase from the $|\psi_1^z\rangle$ term, leaving two relative phases remaining in this case. Projecting first into the x -basis, the populations are,

$$x_1 = \frac{1}{4} \left\{ 1 + b^2 + 2ac \cos \phi_c - 2\sqrt{2}b [c \cos(\phi_c - \phi_b) + a \cos \phi_b] \right\}, \quad (\text{C.15})$$

$$x_2 = \frac{1}{2} (1 - b^2 - 2ac \cos \phi_c), \quad (\text{C.16})$$

$$x_3 = \frac{1}{4} \left\{ 1 + b^2 + 2ac \cos \phi_c + 2\sqrt{2}b [c \cos(\phi_c - \phi_b) + a \cos \phi_b] \right\}. \quad (\text{C.17})$$

From these populations, we find two useful relations,

$$x_3 - x_1 = \sqrt{2}b [c \cos(\phi_c - \phi_b) + a \cos \phi_b], \quad (\text{C.18})$$

$$\cos \phi_c = -\frac{2x_2 - 1 + b^2}{2ac}. \quad (\text{C.19})$$

Looking at the y -populations,

$$y_1 = \frac{1}{4} \left\{ 1 + b^2 - 2ac \cos \phi_c + 2\sqrt{2}b [c \sin(\phi_c - \phi_b) + a \sin \phi_b] \right\}, \quad (\text{C.20})$$

$$y_2 = \frac{1}{2} (1 - b^2 + 2ac \cos \phi_c), \quad (\text{C.21})$$

$$y_3 = \frac{1}{4} \left\{ 1 + b^2 - 2ac \cos \phi_c - 2\sqrt{2}b [c \sin(\phi_c - \phi_b) + a \sin \phi_b] \right\}, \quad (\text{C.22})$$

which yield similar relations as before,

$$y_1 - y_3 = \sqrt{2}b [c \sin(\phi_c - \phi_b) + a \cos \phi_b], \quad (\text{C.23})$$

$$\cos \phi_c = -\frac{2y_2 - 1 + b^2}{2ac}. \quad (\text{C.24})$$

Combining the expressions for $x_3 - x_1$ and $y_1 - y_3$ we obtain,

$$\tan(\phi_c - \phi_b) = \frac{(y_1 - y_3) - \sqrt{2}ab \sin \phi_b}{(x_3 - x_1) - \sqrt{2}ab \cos \phi_b}. \quad (\text{C.25})$$

The z -populations are,

$$z_1 = a^2, \quad (\text{C.26})$$

$$z_2 = b^2, \quad (\text{C.27})$$

$$z_3 = c^2. \quad (\text{C.28})$$

As in the case of spin-1/2 above (Sec. C.0.1), we may assume that $a = \sqrt{z_1}$, taking only the positive solution, which amounts to factoring out a global phase of π . Unlike in the previous case, however, we are not permitted to do the same for the other coefficients since the various potential sign combinations result in different relative phases. Therefore, even if we determine ϕ_b and ϕ_c , we will only know them up to a factor of π , unless we can further constrain the populations b and c . Hence, for now,

$$a = \sqrt{z_1}, \quad (\text{C.29})$$

$$b = \pm\sqrt{z_2}, \quad (\text{C.30})$$

$$c = \pm\sqrt{z_3}. \quad (\text{C.31})$$

Looking at the remaining equations, there is not much more we can determine without adding additional measurements. We will thus add another set of projections, in the eigenbasis of the spin matrix $\hat{F}_v = (\hat{F}_x + \hat{F}_y)/2$, where we have used the subscript v for notation simplicity. As before, there are any number of combinations of spin matrices we could have chosen, but we found through trial-and-error

that this one is useful. The populations in this basis are,

$$v_1 = \frac{1}{4} \left\{ 1 + b^2 + 2ab (\sin \phi_b - \cos \phi_b) - 2c [b \{ \cos (\phi_c - \phi_b) - \sin (\phi_c - \phi_b) \} + a \sin \phi_c] \right\}, \quad (\text{C.32})$$

$$v_2 = \frac{1}{2} (1 - b^2 + 2ac \sin \phi_c), \quad (\text{C.33})$$

$$v_3 = \frac{1}{4} \left\{ 1 + b^2 - 2ab (\sin \phi_b - \cos \phi_b) + 2c [b \{ \cos (\phi_c - \phi_b) - \sin (\phi_c - \phi_b) \} - a \sin \phi_c] \right\}. \quad (\text{C.34})$$

Rearranging the expression for v_2 , we obtain,

$$\sin \phi_c = \frac{2v_2 - 1 + b^2}{2ac}. \quad (\text{C.35})$$

From this, and either form of $\cos \phi_c$ from above, we may solve for the phase ϕ_c as,

$$\phi_c = \arctan \left(-\frac{b^2 + 2v_2 - 1}{b^2 + 2x_2 - 1} \right) = \arctan \left(\frac{b^2 + 2v_2 - 1}{b^2 + 2y_2 - 1} \right). \quad (\text{C.36})$$

We now have each of the superposition amplitudes, a , b , and c up to a sign, and one of the relative phases ϕ_c . From these, we may determine the magnitude of the other phase ϕ_b by finding the root of,

$$\arctan \left[\frac{(y_1 - y_3) - \sqrt{2}ab \sin \phi_b}{(x_3 - x_1) - \sqrt{2}ab \cos \phi_b} \right] - (\phi_c - \phi_b) = 0. \quad (\text{C.37})$$

Generally, finding the roots of transcendental equations is not easy, even numerically. Furthermore, there is still some remaining ambiguity in the relative phase, due to us knowing the coefficients only up to a sign. Despite this, we are at least able to determine the magnitude of each parameter through these relations. A primary reason for this is that we have only considered eigenstates of the spin matrices, but for higher spin systems there are other more appropriate representations, such as the Gell-Mann matrices in the case of spin-1. This choice reflects our measurement capabilities in the experiments discussed in Ch. 6, where the RF-pulses we apply only produce transformations in $SU(2)$. If we were able to implement transformations in $SU(3)$, for states in spin-1, we expect that there are some much more useful projections we could measure, producing similar results to spin-1/2 where the state was well determined. Regardless, we still expect that more than three bases are required.

The projections considered here, in the very least, have informed us as to which projections constitute an informationally complete set of measurements; the four

projections here are enough for the amplitudes of the parameters, but more are necessary for complete information. We numerically verified the above analytical relations to determine the wavefunctions of arbitrary spin states (randomly sampled), testing the accuracy of our results. We found that due to the transcendental equation for ϕ_b , the results were unstable. Since spin-1 systems are not yet large enough to be difficult to simulate (a problem in the QC community due to the large dimensional Hilbert spaces considered there), we found that numerically simulating the measurements produced far better results, especially in the presence of noise (numerically simulated). The results improved drastically when adding additional measurement bases as well, such as $\hat{F}_u = (\hat{F}_x + \hat{F}_z)/2$. As such, when implemented in the analysis of experiments in Ch. 6, tomography was done by simulating the measurement process, rather than applying the analytical expressions above. It is also worth noting that the technique shown here works only on pure-states; in the presence of decoherence, other methods are required [27, 67].

APPENDIX D

Floquet-Engineered Λ -Scheme

THE original proposal for realizing the Floquet-engineered holonomies [219] in a BEC, which the experiments presented in Ch. 6 were based on, utilized a Λ -scheme Raman coupling. The level diagram for this scheme is depicted in Fig. 2.4(b). If applied to the $F = 1$ manifold of ^{87}Rb , only the outermost spins are coupled, $m_F = \pm 1$, resulting in a pseudospin-1/2 system. The desired Floquet-engineered holonomies may be obtained through this scheme if the coupling parameters are modulated in a similar manner as the RF-fields in Sec. 6.2.1. As such, our initial attempts at the experiment were to utilize such a Raman dressing approach, rather than RF. Here we will briefly discuss our progress towards this goal, which ultimately failed due to an unavoidable issue specific to the Λ -scheme, in which spontaneous emission from the lasers produced significant decoherence in the ensemble.

The experimental setup is essentially the same as that described in Ch. 6, aside from the replacement of RF fields with laser fields. As discussed in Sec. 2.2, we assume a bias magnetic field is applied to split the magnetic sublevels, creating a resonance condition when the frequency difference between the Raman beams is equal to this level splitting. When the beams are co-propagating, there is no net momentum transfer to the atoms from the absorption and subsequent stimulated emission events. As such, the Hamiltonian for this system is a modified version of Eq. 2.49; since the $F = 0$ level is left completely uncoupled, the system can be reduced to a pseudospin-1/2 subspace, consisting of only the $m_F = \pm 1$ states. The hyperfine operators in the Hamiltonian (Eq. 2.49) must be replaced by the spin matrices (Pauli matrices times $\hbar/2$), as they no longer represent total angular momentum in this case. Aside from these modifications, the system is effectively a spin-1/2 version of the unmodulated Hamiltonian Eq. 5.1.

The necessary Floquet driving may be accomplished by simultaneously modulating the amplitude and phase of the Raman beams through AOMs. Both beams, derived from the same laser, may each be brought through an AOM. The AOMs are

driven with frequencies that differ by the Zeeman splitting resulting from the applied bias field; the beams are therefore on resonance with the transition. Furthermore, the phase difference between the beams may be controlled by driving each AOM from a common source, with programmable phases; the same AWG described in Ch. 6 is sufficiently fast to provide the 80 MHz AOM carrier frequencies, with full phase and frequency control. The phase, frequency, and amplitude of the individual Raman beams may be modulated in this way. Counter to the original proposal [219], which suggests performing the frequency modulation through the background bias field, it was determined that in our setup, we would obtain better control if we accomplished this by frequency modulating the AOMs instead¹.

In this scheme, the relative amplitudes, phases, and polarizations of the Raman beams must be stable, maintaining the desired relationship. In order to implement this in such a stable manner, the path lengths of each laser through their AOMs must match as closely as possible. Furthermore, they must be close in proximity, along similar directions, so that any perturbations of the free-space optics are common to both beams, such as wind currents which disrupt the phase and polarizations. After recombination on a polarizing beam splitter, the beams will have orthogonal linear polarizations. If transmitted to the experiment through a single-mode optical fiber, the modes will be properly overlapped (and fluctuations from the fiber will be common to both beams). The orthogonal circular polarizations can then be obtained by passing the combined beams through a quarter wave-plate, prior to transmitting them through the ensemble.

In principle, setting up the beams and modulating them as described above should be sufficient to realize the Floquet Hamiltonian in Sec. 5.1, however, there are further practical challenges. We found that, even with the stable free-space optical setup, temperature fluctuations affecting the optical fiber were enough to make the beam intensities and polarizations unstable over short periods of time (several minutes). It would be prudent, therefore, to install some active feedback to stabilize the beams. Actively compensating the polarization directly is complicated, but thankfully for this scheme it should be sufficient to compensate with the beam intensity alone. Near where the beams are transmitted into the atomic ensemble, a small portion may be split off and analyzed through a polarizing beam splitter. A PID feedback system could then be tuned to maintain the intensity coming through each port of the polarizing beam splitter, each corresponding to one of the beams, controlling the driving AOMs. There would thus need to be two independent feedback systems for this.

¹This is due to two reasons: it is not necessarily easy to modulate coil currents, especially in a synchronized way, and the modulation depth matters here, which may be better controlled through careful calibration of the AWG channels.

Even with such an improvement, however, it is unlikely that the Λ -scheme is an appropriate method to realize the Floquet-engineered Hamiltonian; the spontaneous emission rates are too significant [24]. Raman coupling schemes rely on the adiabatic elimination of the excited state levels. In order to properly predict the Rabi frequencies, and the spontaneous emission rates, *all* of the excited state manifolds that may mediate the transition must be accounted for; the interference between transition paths plays a crucial role in the results. For Λ -transitions, one finds that the ratio between the Rabi-frequency of transitions Ω_0 and the spontaneous emission rate γ , Ω_0/γ , increases with detuning, and then saturates for large detunings. This means that no matter how large the detuning is, the number of Rabi-oscillations between levels permitted before a spontaneous emission event is fixed (on average). For ^{87}Rb the ratio saturates at $\Omega_0/\gamma = 13$, which is among the highest of all other alkali atoms [24]. This means that decoherence due to spontaneous emission cannot be avoided by increasing the Rabi-frequency, setting a fundamental “time limit” on applied pulses.

In the RF implementation of the Floquet-engineered Hamiltonian (Ch. 6), the duration of the operations was set by a subharmonic of the Rabi frequency; longer, more adiabatic gates, would be set by choosing a smaller subharmonic. This essentially permits more Rabi-oscillations per gate. In the Raman Λ -scheme described above, we are constrained by the fixed decoherence rate with respect to the drive. An example set of data, looking just at the how the populations evolve over time according to unmodulated Rabi-oscillations, may be seen in Fig. D.1. Over the pulse duration, a significant number of spontaneous emission events mix the spins between both ground state hyperfine manifolds, and lead to significant losses from the ODT, eventually making the atom number so small that they cannot be efficiently imaged. This may be seen directly in the images at two different pulse times in Fig. D.1(b-c). Near the end of the pulse, the populations tend towards being equally mixed, due to spin decoherence. We found this to be completely detrimental to the experiment, as there was no way to maintain the adiabatic condition, Eq. 5.32, while maintaining coherence, or retaining enough atoms to properly image. For these reasons, we adopted the RF dressing scheme utilized in Ch. 6.

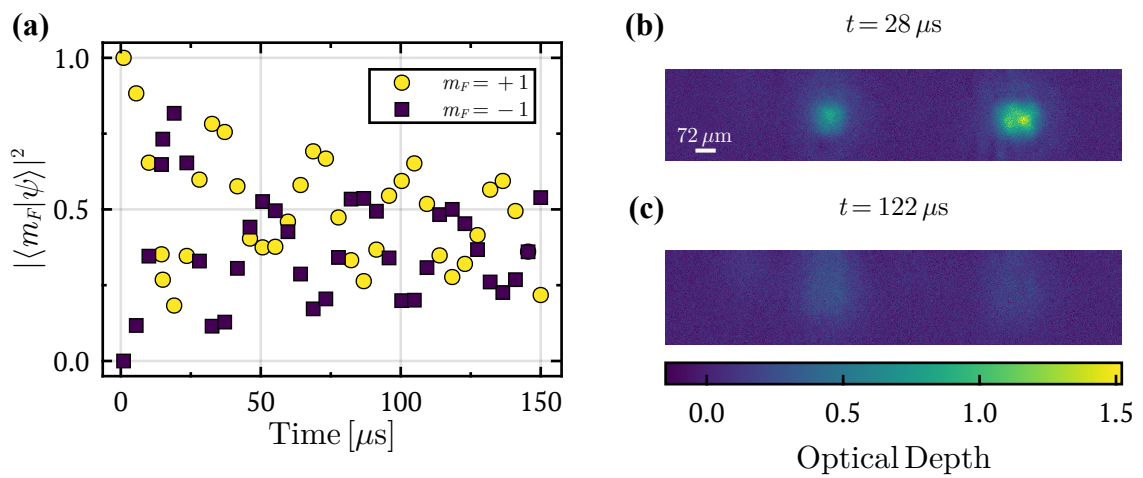


Figure D.1: Results of Rabi-oscillations under Raman Λ -scheme, with two co-propagating beams of equal power (about 33 mW each), opposing circular polarizations, and that differ in frequency by 2.5 MHz. The beams are red-detuned from the $F = 1 \rightarrow F' = 1$ D1 transition by 2.28 GHz. (a) populations in the $m_F = \pm 1$ levels over time, showing several Rabi-oscillations that are damped, due to decoherence effects from spontaneous emission events. (b) Sample image from the data set, with pulse duration $t = 28 \mu\text{s}$, showing relatively high overall atom number. (c) Sample image with pulse duration $t = 122 \mu\text{s}$, plotted on the same scale as (b), showing the significant decrease in atom number over time.

APPENDIX E

Computing Holonomies

HOLONOMIES play a crucial role in describing the evolution of quantum systems subjected to non-Abelian connections; despite this, they are often difficult to compute due to path-ordering, which ensures that, for loops over which the connection does not commute with itself at various points, the integration includes these non-commutative effects. Here, we consider the holonomies produced by periodically driving the Hamiltonian of a spin in a magnetic field (Eq. 6.8), studied in Ch. 5 and Ch. 6. We will set up the contour integral from the connection, and then carry out the integrals that result in the holonomies in Tab. 5.1.

Starting with the form of the connection (Eq. 5.35), we have,

$$\begin{aligned}
 \hat{\mathbf{A}} &= g \left(\hat{\mathbf{F}} \times \boldsymbol{\lambda} \right), \\
 &= g \left[\left(\lambda_z \hat{F}_y - \lambda_y \hat{F}_z \right) \mathbf{e}_x + \left(\lambda_x \hat{F}_z - \lambda_z \hat{F}_x \right) \mathbf{e}_y + \left(\lambda_y \hat{F}_x - \lambda_x \hat{F}_y \right) \mathbf{e}_z \right], \\
 &= g \left[\left(\cos \Theta \hat{F}_y - \sin \Theta \sin \Phi \hat{F}_z \right) \mathbf{e}_x + \left(\sin \Theta \cos \Phi \hat{F}_z - \cos \Theta \hat{F}_x \right) \mathbf{e}_y \right. \\
 &\quad \left. + \left(\sin \Theta \sin \Phi \hat{F}_x - \sin \Theta \cos \Phi \hat{F}_y \right) \mathbf{e}_z \right], \tag{E.1}
 \end{aligned}$$

where $\boldsymbol{\lambda} = (\sin \Theta \cos \Phi, \sin \Theta \sin \Phi, \cos \Theta)^\top$. Since the loops in this case are parameterized by the spherical coordinates $\{\Theta(t), \Phi(t)\}$, with the radius $r = 1$, we will change basis according to,

$$\mathbf{e}_x = \sin \Theta \cos \Phi \mathbf{e}_r + \cos \Theta \cos \Phi \mathbf{e}_\Theta - \sin \Phi \mathbf{e}_\Phi, \tag{E.2}$$

$$\mathbf{e}_y = \sin \Theta \sin \Phi \mathbf{e}_r + \cos \Theta \sin \Phi \mathbf{e}_\Theta - \cos \Phi \mathbf{e}_\Phi, \tag{E.3}$$

$$\mathbf{e}_z = \cos \Theta \mathbf{e}_r - \sin \Theta \mathbf{e}_\Theta. \tag{E.4}$$

Substituting these in to the above, we obtain,

$$\hat{\mathbf{A}}/g = \left(-\sin \Phi \hat{F}_x + \cos \Phi \hat{F}_y \right) \mathbf{e}_\Theta + \left(-\cos \Theta \cos \Phi \hat{F}_x - \cos \Theta \sin \Phi \hat{F}_y + \sin \Theta \hat{F}_z \right) \mathbf{e}_\Phi. \quad (\text{E.5})$$

where the \mathbf{e}_r component has canceled out.

The holonomies take the form,

$$\hat{\Gamma}_A(\ell) = \mathcal{P} \exp \left(-\frac{i}{\hbar} \oint_\ell d\boldsymbol{\lambda} \cdot \hat{\mathbf{A}} \right), \quad (\text{E.6})$$

where $d\boldsymbol{\lambda} = dr\mathbf{e}_r + rd\Theta\mathbf{e}_\Theta + r\sin\Theta d\Phi\mathbf{e}_\Phi$. Looking at the integral inside the exponential, together with the connection above, we can see that for loops in which $\Theta = \Theta(t)$ and $\Phi \equiv \text{const.}$, the \mathbf{e}_Θ component of the connection is constant over the loop. This may therefore be integrated explicitly, yielding,

$$\begin{aligned} \oint_\ell d\boldsymbol{\lambda} \cdot \hat{\mathbf{A}} &= g \left(-\sin \Phi \hat{F}_x + \cos \Phi \hat{F}_y \right) \int_0^{2\pi} d\Theta, \\ &= -2\pi g \left(\sin \Phi \hat{F}_x - \cos \Phi \hat{F}_y \right). \end{aligned} \quad (\text{E.7})$$

With the choices of Φ for the holonomies in Tab. 5.1 this leaves,

$$\Phi = 0 \Rightarrow \hat{\Gamma}_A(\ell) = \exp \left(-i2\pi g \hat{F}_y / \hbar \right), \quad (\text{E.8})$$

$$\Phi = \frac{\pi}{2} \Rightarrow \hat{\Gamma}_A(\ell) = \exp \left(i2\pi g \hat{F}_x / \hbar \right), \quad (\text{E.9})$$

$$\Phi = \frac{\pi}{4} \Rightarrow \hat{\Gamma}_A(\ell) = \exp \left[i\sqrt{2}\pi g \left(\hat{F}_x - \hat{F}_y \right) / \hbar \right], \quad (\text{E.10})$$

for loops ℓ_1 , ℓ_2 , and ℓ_4 , respectively.

Similarly, for loops in which $\Theta \equiv \text{const.}$ and $\Phi = \Phi(t)$, we are left only with an integral around Φ . Unlike the case before, the \mathbf{e}_Φ component does not commute with itself along the path, except when $\Theta = \pi/2$ specifically. Therefore, we may compute this loop, but the path-ordering must be taken into account for any other values of Θ . We have,

$$\begin{aligned} \oint_\ell d\boldsymbol{\lambda} \cdot \hat{\mathbf{A}} &= g \hat{F}_z \int_0^{2\pi} d\Phi, \\ &= 2\pi g \hat{F}_z. \end{aligned} \quad (\text{E.11})$$

The holonomy for this loop, ℓ_3 in Tab. 5.1, is therefore,

$$\hat{\Gamma}_A(\ell) = \exp \left(-i2\pi g \hat{F}_z / \hbar \right). \quad (\text{E.12})$$

For any other choices of $\Theta(t)$ and $\Phi(t)$ outside of the forms covered here, path-ordering must be considered.