

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Towards a next-generation measurement of the fine-structure constant

Permalink

<https://escholarship.org/uc/item/5qm678c2>

Author

Pagel, Zachary

Publication Date

2023

Peer reviewed|Thesis/dissertation

Towards a next-generation measurement of the fine-structure constant

By

Zachary Pagel

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Holger Mueller, Chair

Professor Karl Van Bibber

Professor Normal Yao

Fall 2023

Abstract

Towards a next-generation measurement of the fine-structure constant

by

Zachary Pagel

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Holger Mueller, Chair

Atom interferometry is a powerful metrological tool that has been developed over the last few decades. Large momentum transfer (LMT) methods manipulate atomic trajectories with tens or hundreds of photon momenta in order to increase sensitivity. This thesis furthers progress towards using LMT methods in next-generation atom interferometers. One main result establishes symmetric Bloch oscillations as a new, viable technique for LMT. Theory and numerics are used to show how the process is coherent and adiabatic, and experimentally we demonstrate coherence in an interferometer with up to $240\hbar k$, where $\hbar k$ is the momentum of a single photon of 852nm light. This was the second largest coherent momentum splitting demonstrated at the time of publication. The rest of the thesis focuses on design and construction of a new atomic fountain to measure the fine structure constant α . Discrepancies in recent measurements of α [67, 55] are currently limiting theory predictions for the electron gyromagnetic ratio [25] - an improved measurement of α is therefore highly motivated and would enable an improved test of the consistency of the Standard Model. Previously, our group published a measurement of α at the 0.2ppb level in 2018 [67]. We built a new experiment with a goal of improving the measurement by a factor of 3-10. Much of the thesis focuses on systematic effects related to spatial intensity inhomogeneities on the laser beam, which are some of the hardest to characterize systematic effects looking forward. A large clear aperture vacuum chamber accommodates larger waist laser beams without clipping on chamber walls. In addition, a high-speed, user-friendly Monte Carlo simulation package was made to predict experimental systematic shifts in the measured value of α due to laser beam inhomogeneities.

First, I'd like to dedicate this thesis to people working on physics because they believe their work is truly important. I fell in love with our experiment to measure the fine-structure constant because it's one of few experiments in atomic physics that has profound implications for cutting-edge fundamental. Too many experiments in our field are done because we physicists *can* do it, and not because we *should* do it. Our experiment to measure the fine-structure constant doesn't use sexy new atomic physics at this point, but we're pushing hard because of the value to fundamental science.

Of course, I didn't do this research all myself. Thanks to all the labmates over the years who I've worked with or learned from. I'm very proud to have worked in the Mueller group and with all of the researchers in our lab - it's a fun bunch of people working on very cool research. Thanks to Holger for assembling this team of researchers and for giving us loads of money to develop these experiments.

Next, I'd like to thank my family and friends for all the support over the years, for having to listen to me talk about boring physics and for hearing all of my complaints. Thanks to my friends that adventure outdoors with me, these adventures have come to be very central to my life and help keep me going. Thanks to Sara Campbell for helping me get an internship at Quantinuum last summer, which allowed me to experience the quantum industry and confidently decide on industry as my future career path. I'm very thankful to have started climbing and playing guitar, among many other hobbies, during grad school. I couldn't have done it without the balance in my life from family, friends and fun things to do.

Thanks to my partner Paige for helping me through the last year, and for giving me the excuse I needed to move to Colorado. And to our cute new puppy Huckleberry whose puppy energy and puppy kisses helped me get through writing my thesis.

Contents

1	Introduction	1
1.1	Overview of the thesis	1
1.2	What is atom interferometry?	2
1.2.1	Atomic beamsplitters	3
1.2.2	Atom interferometer phase	5
1.2.3	Different types of interferometers	9
1.3	Measuring the recoil frequency of Cesium	10
1.4	From the recoil frequency to dark matter	11
2	Bragg diffraction and Bloch oscillations - theory and numerics	16
2.1	Introduction	16
2.2	Optical lattice Hamiltonian	17
2.3	Bragg diffraction Hamiltonian	22
2.3.1	Single frequency Bragg diffraction	23
2.3.2	Multi-frequency Bragg diffraction	24
2.4	Different regimes of the Bragg Hamiltonian	25
2.5	Bloch oscillations Hamiltonian	26
2.6	Bloch numerical simulation	30
2.6.1	Landau-Zener Tunneling	31
2.6.2	Role of excited states during Bloch oscillations	32
2.6.3	Systematic effects from Bloch oscillation excited states	36
2.7	Phase evolution during Bloch oscillations	36
2.7.1	Potential systematics from Bloch oscillations	38
2.8	Dimensionless units	40
2.9	Laser power to lattice depth	43
3	Symmetric Bloch oscillations of matter waves	46
3.1	Pre-introduction	46
3.2	Abstract	47
3.3	Introduction	47
3.4	Theory	50
3.4.1	Hamiltonian and unitary transformation	50
3.4.2	Limits on ramp rate from the rotating wave approximation	55
3.4.3	Limits on ramp rate from Landau-Zener tunneling and higher-order transitions	56

3.4.4	Comparison of limits on the ramp rate	56
3.4.5	Crossing through velocity degeneracy	59
3.4.6	Experimental considerations	60
3.5	Experiment	61
3.5.1	Atom optics with symmetric Bloch oscillations	61
3.5.2	Mach–Zehnder interferometer	64
3.5.3	Gradiometer	64
3.6	Conclusions and Outlook	66
3.7	Supporting material	67
3.7.1	Unitary transformation for single-lattice Bloch Hamiltonian	67
3.7.2	Symmetrized Hamiltonian	68
3.7.3	Rotating wave approximation condition	69
3.7.4	Higher-order loss mechanisms	71
3.7.5	Crossing through velocity degeneracy	71
3.7.6	Diffraction phase	73
3.7.7	Higher-order generalization of the dual-lattice methods	74
3.7.8	Application to recoil measurements	75
4	Experimental Overview	77
4.1	Introduction	77
4.2	Motivation and design decisions	77
4.2.1	Beam related effects	79
4.2.2	In-vacuum optics and beam delivery	80
4.2.3	Acceleration Gradient	81
4.2.4	Statistical sensitivity	82
4.3	The rest of the vacuum chamber design	83
4.3.1	Magnetic shielding	83
4.3.2	Solenoid winding	88
4.3.3	In-vacuum baffles for stray light	89
4.3.4	Hexapod, support structure and vibration isolation	93
4.3.5	In-vacuum shutter for beam imaging	94
4.3.6	Designing for UHV	95
4.3.7	Summary of design decisions	101
4.4	Installing chamber and getting vacuum	102
4.5	The rest of the atomic fountain	105
4.5.1	Control system	105
4.5.2	Compact 2D MOT design	106
4.5.3	3D MOT	107
4.5.4	Raman Sideband Cooling	108
4.5.5	Detection	109
4.5.6	Interferometry	110
4.6	New ideas	111
4.6.1	Systematic effect from velocity selection light shifts	111
4.6.2	In-vacuum beam imaging	114
4.6.3	Systematic effect from parasitic lattices	117

4.6.4	Lattice launch	121
5	Monte Carlo overview	124
5.1	Introduction	124
5.2	Motivation	125
5.2.1	Monte Carlo Design Philosophy	127
5.3	Monte Carlo physics	127
5.3.1	Laser Electric Field Models	128
5.3.2	Physical effects from laser beam intensity inhomogeneities	132
5.3.3	From wavefunction amplitude to ellipses	137
5.4	Overview of the Monte Carlo code	139
5.4.1	Introduction to classes	141
5.4.2	Atom Class	142
5.4.3	TimeEvolution Classes	143
5.4.4	Ellipse and EllipsePoint Classes	146
5.4.5	Sweep and Sweep2D Classes	147
5.4.6	Precompute Class	147
5.4.7	Beam class	150
5.5	Typical Monte Carlo Workflows and Intended Use	151
5.6	Current Monte Carlo Speed	153
5.6.1	Potential Monte Carlo Speed-ups	155
5.7	Monte Carlo convergence	156
5.7.1	Integration convergence	156
5.7.2	Precompute table convergence	157
5.7.3	Statistical convergence	157
5.8	Monte Carlo Results	158
5.9	Initial experimental results	161
5.10	Future Monte Carlo Work	163
5.10.1	Monte Carlo Assumptions	163
6	Conclusions and next steps	167
6.1	Conclusions	167
6.2	Next steps	168

List of Figures

1.1	Optical versus matterwave interferometers	3
1.2	Bragg diffraction and Bloch oscillations matterwave evolution	4
1.3	Common atom interferometer geometries	8
1.4	Simultaneous conjugate interferometry	10
1.5	Feynmen diagrams relevant to $g_e - 2$	14
2.1	Bragg diffraction energy level diagram	22
2.2	Quasi-Bragg and Bragg diffraction regimes	25
2.3	Diffraction in the Raman-Nath regime	26
2.4	Bloch band structure	29
2.5	Simulation versus theory, Landau Zener tunneling	32
2.6	Excited state population with $10 E_r$ lattice depth	33
2.7	Excited state population with $20 E_r$ lattice depth	34
2.8	Optimizing lattice loading	35
2.9	Phase evolution during Bloch oscillations	37
2.10	Systematic effects from Bloch oscillations	39
3.1	Numerical and experimental symmetric Bloch oscillations	49
3.2	Energy level diagram for symmetric Bloch oscillations	51
3.3	Energy band structure for symmetric Bloch oscillations	54
3.4	Single lattice versus dual lattice Bloch oscillations	57
3.5	Symmetric Bloch oscillation losses	58
3.6	Experimental realization of various dual-lattice operations	62
3.7	Mach-Zehnder interferometer contrast with symmetric Bloch oscillations	63
3.8	Dual Mach-Zehnder interferometer with symmetric Bloch oscillations	65
3.9	Higher-order transitions during symmetric Bloch oscillations	70
3.10	Diffraction phase during symmetric Bloch oscillations	73
3.11	Diffraction phase versus ramp rate	74
3.12	Recoil-sensitive symmetric Bloch oscillations	76
4.1	Systematic shift from mass of vacuum flange.	81
4.2	Vacuum chamber overview	84
4.3	Cross section view of vacuum chamber	85
4.4	Systematic effects from magnetic field fringing	86
4.5	Rotisserie used to wind the in-vacuum solenoid	88
4.6	Estimated systematic effects from solenoid winding imperfections	90

4.7	Copper oxidation procedure	91
4.8	Image of the vacuum chamber insides	92
4.9	FEA simulation of support structure vibration modes	95
4.10	Vacuum dangers	97
4.11	Final vacuum pressure estimate	99
4.12	Vacuum parts assembly in physics breezeway	103
4.13	Vacuum installation steps	104
4.14	Images of vacuum installation	104
4.15	2D MOT design and realization	106
4.16	New photodiode array for detection	110
4.17	Raman transition light shifts	112
4.18	Light shift and Rabi frequency during Raman transitions	113
4.19	Achievable phase from in-vacuum beam imaging	115
4.20	Example population distrubition, in-vacuum beam imaging	116
4.21	Frequency pairs forming lattices in new experiment	117
4.22	Parasitic lattices interfering with Bloch oscillations	118
4.23	Limits to total phase from parasitic lattices	119
4.24	Lattice launch survival fraction	122
5.1	Gouy shift from Gaussian beam	125
5.2	Gaussian approximation to airy rings from dust	131
5.3	Propagation of a dusty beam	132
5.4	Monte Carlo wavefunction ordering	141
5.5	Precomputed integration data	149
5.6	Monte Carlo integration convergence	156
5.7	Contrast loss versus beam modulation frequency	159
5.8	Predicted phase shift versus spatial modulation frequency	159
5.9	Simulated effects of dust on optics	160
5.10	Square wave modulated beam propagation	161

List of Tables

1.1	Phase contributions for common atom interferometers	9
1.2	Current known accuracies of various relevant constants	13
2.1	Dimensionless unit system	41
4.1	Error table from 2018 α measurement	78
5.1	Profiling current speed of Monte Carlo code	154

Chapter 1

Introduction

1.1 Overview of the thesis

There are pros and cons to working on a measurement with a several-decade history - on the one hand, the path to doing good science is laid out for you very clearly, but on the other hand, the problems you're left to solve are nuanced and difficult. This thesis presents the current state of Berkeley's atom interferometry measurement of the fine-structure constant, which is a continuation of Steve Chu's experiment from Stanford. Once I found a lab notebook high up in the corner of one of our labs in Berkeley from 1992 at Stanford. Reading through the pages was comically similar to the day-to-day issues we still deal with in our lab. The technologies haven't really changed much since then, nor has our lab culture of having a lot of hand-built circuits and analog knobs.

Since this thesis is building on decades of work from past students on our project, I refer the reader to other theses (mainly Brian Estey's [23] and Chenghui Yu's [99]) for material that they have already covered well. Some sections might require more context from Brian and Chenghui's theses in order to gain a complete understanding, but I do my best to make the thesis read as a single coherent document.

A main multi-year push of this thesis was designing and installing a giant vacuum chamber with a pressure at the 10^{-10} torr level. Another multi-year push was in writing a Monte Carlo code package for simulations of beam-related systematic effects. We had Monte Carlo code from the previous generation of the experiment, however the new code was written to be faster and more user-friendly. The hope is that future students can build on this instead of writing a new code package.

We've also worked out a method to use symmetric Bloch oscillations as a beamsplitter, mirror, and recombiner for making coherent atom interferometers. I really enjoyed incorporating theory, simulation, and experiment all into one project to give a fairly clear picture of the process in the first publication [65]. Moreover, we were able to demonstrate coherence of the process in an atom interferometer with up to $240\hbar k$ momentum splitting, which is the momentum of 240 photons from the 852 nm laser beam. This was the second-largest coherent momentum splitting achieved at the time of publication.

This thesis also presents an in-depth treatment of Bloch oscillations. The previous generation of the experiment assumed that the high degree of symmetry in our experiment meant

that systematic effects from Bloch oscillations could be assumed to be small [67]. I make these arguments more rigorous by outlining a more detailed theory of Bloch oscillations and the possible systematic effects. The high degree of symmetry in our simultaneous conjugate Ramsey-Borde interferometers does cancel potential systematic shifts to a high degree of accuracy. This should give a good starting point for treating Bloch oscillation systematic effects more carefully in the future.

With that all laid out, the outline of my thesis is as follows. This chapter will provide an introduction to atom interferometry, some basic atom interferometer geometries and what they measure, an overview of the SCI interferometer we use to measure the recoil frequency of Cesium, and lastly a summary of why we measure the recoil frequency of Cesium and how our measurement contributes to the most precise test of the Standard Model.

Chapter 2 will provide a detailed overview of Bragg diffraction and Bloch oscillations, deriving the Hamiltonians and discussing their different regimes. This includes a section on the phase evolution of wavefunctions during Bloch oscillations. Chapter 3 discusses the Symmetric Bloch oscillations paper [65]. Chapter 4 summarizes the experimental work in designing and building the new apparatus in Campbell Hall, the current status of the experiment, and some exciting new ideas for the near-future. Chapter 5 summarizes the physics and the design of the Monte Carlo code used to study beam-related systematic effects. It also discusses initial results in using the code to study the systematic effects from intentionally bad laser beam profiles, and some initial experimental data towards matching experimental results to the Monte Carlo. This Monte Carlo model will hopefully provide the foundation for bounding systematic effects related to laser beam inhomogeneities in the next measurement of α .

1.2 What is atom interferometry?

It turns out this is the answer to the Jeopardy prompt “Keeps you awake at night”. Atom interferometry is also a powerful technology that uses cold atoms to make very sensitive measurements. The heart of the technology is the particle-wave duality of matter - matterwaves can be interfered with themselves after travelling along two distinct spatial trajectories.

There is a very nice parallel between optical and matter-wave interferometry as illustrated in Fig. 1.1. Optical interferometers use physical mirrors and beam-splitters to split laser light along two spatial trajectories and then recombine the light. The electric fields of the recombined beams interfere constructively and destructively in the two output ports depending on the optical path length of the two interferometer arms. Varying the optical path length of one arm by half of a wavelength of light will cause the output ports to vary from constructive to destructive interference. Optical interferometers are therefore good at measuring nanometer scale length changes.

Matter-wave interferometers use optical pulses as beamsplitters and mirrors to send the matter-waves along distinct spatial trajectories. Note that the role of light and matter have been reversed here - beamsplitters and mirrors are now made of light instead of matter, and matter travels along the different trajectories instead of light.¹ Interference of the matter-waves causes the atom number population in the two output ports to vary, similar

¹Cool stuff




Figure 1.1: Schematic of an a) optical interferometer with output optical powers P_1 and P_2 , and b) an atom interferometer with output atom numbers N_1 and N_2 . In both diagrams, black lines indicate matter and red lines indicate light. The role of light and matter are inverted in the two cases.

to the optical power varying in the output ports of the optical interferometer. The analogue between optical and atom interferometers is satisfying and simple for generating an initial understanding of the physics. However, the details of what these instruments measure very quickly diverges, and atom interferometers are much more complicated instruments.

The matterwave interference in an atom interferometer depends on the phase difference of the matterwaves being interfered. Matterwaves are complex valued functions whose dynamics are governed by the Schrodinger equation - this will be introduced explicitly in Chapter 2, but for now we can loosely say that the phase of the matterwave wavefunction depends on the energy of the atomic states along the difference paths of the interferometer. For example, if one arm of the interferometer spends time higher up in the gravitational potential, the resulting energy difference can be used to measure gravity very precisely. Similarly, the kinetic energy can be different along the different trajectories if the atoms travel at different velocities, so some atom interferometers are sensitive to the velocity kicks that an atom receives during the various laser pulses (or any external forces that also change the velocity of the atom). The interferometers explored in this thesis are designed to very precisely measure the recoil kinetic energy that an atom receives upon absorption or emission of one photon from the beam-splitting laser.

1.2.1 Atomic beamsplitters

If you're not familiar with atom interferometry, the idea that you can use a laser to split an atom along multiple trajectories is a novel idea! The laser generates a quantum superposition between two different momentum states of the atom, and these different momentum states time-evolve in a way that they spatially separate. If you were to measure the position of the atom, you would only find the atom along one trajectory or the other. But if you don't measure the atoms, the quantum state evolves along both trajectories simultaneously. Because the relative phase of the matterwave wavefunctions is well-defined between the two trajectories, we call the beamsplitters 'coherent'.




Figure 1.2: Numerical simulation of matterwaves $\psi(x, t)$ undergoing a) Bragg diffraction and b) Bloch oscillations - the probability density $|\psi(x, t)|^2$ is plotted as a density plot. a) A moving optical lattice is pulsed with a Gaussian temporal profile about a third of the way through the time evolution. The sharp diagonal lines show the matterwave interacting with the (moving) standing wave of light. The pulse results in a 50/50 beamsplitter, where half of the wavefunction diffracts off of the optical standing wave. b) An optical lattice potential at zero velocity is adiabatically turned on and the wavefunction localizes inside the optical potential wells. Towards the middle of the plot, the lattice is accelerated. The matterwave is able to mostly remain stuck in the optical potential wells, though some of the matterwave tunnels through the potential barrier and leaks out.

Figure 1.2 illustrates what happens to an atomic matterwave during Bragg diffraction and Bloch oscillations. Fig. 1.2a) illustrates a matterwave being coherently split by a Bragg diffraction pulses. Bragg diffraction is a common beamsplitter used in atom interferometers. You may be familiar with Bragg diffraction from x-ray diffraction off of solid state crystal structures - this is the same physics but with the roles of matter and light reversed. Here, a matterwave diffracts off of a (moving) standing wave of light. If the optical pulse was slightly longer duration, or slightly higher intensity, you could instead fully diffracted the matterwave in a ‘ π -pulse’.

Fig. 1.2b) illustrates a matterwave undergoing Bloch oscillations, which is another technique used in this thesis to manipulate atoms using light. The atom is adiabatically loaded into an optical lattice in the first part of the image, and the matterwave can be seen localizing in the potential wells of the optical lattice. Then, the lattice is accelerated and most of the matterwave is able to remain in the accelerating lattice, although some of the matterwave tunnels out (and continues traveling horizontally in the image). Figure 1.2 is meant to give the reader a qualitative understanding of the processes - for detailed quantitative treatment, see Chapter 2.

1.2.2 Atom interferometer phase

With this general understanding of what an atom interferometer is and how light is used to manipulate the atomic trajectories, let’s look a little more carefully at what atom interferometers actually measure. The interference between matterwaves is governed by the phase difference between matterwaves traveling along different spatial trajectories. For more detailed derivations of the phase measured in atom interferometers, see Brian Estey’s thesis [23] - here I’ll give a summary of the main results.

The first main contribution to the phase measured by an atom interferometer is the integrated energy shifts. The Schrodinger equation at it’s most basic level tells us that the phase evolution of a wavefunction $\psi(0)$ in an eigenstate at time $t = 0$ is proportional to the energy E of the eigenstate:

$$\psi(t) = e^{iEt/\hbar}\psi(0) \tag{1.1}$$

where \hbar is the reduced Plank’s constant. Since we’re not only interested in pure eigenstates of the Hamiltonian, we need to make this a little more general. The path-integral formulation of quantum mechanics does just that. You can find the derivation in Brian’s thesis [23] which is based on Robert Littlejohn’s Physics 221AB Quantum Mechanics notes [50]. Phase evolution ends up being governed by the integrated Lagrangian, instead of the integrated Hamiltonian:

$$\int dt \langle \psi_f | e^{iH(t)t/\hbar} | \psi_i \rangle |_{q(t)} = e^{iS(q)/\hbar} \tag{1.2}$$

In words, this equation says that the path-integral of the Hamiltonian $H(t)$ over the classical phase-space path q leads to a phase evolution given by the classical action $S(q)$ over the classical path. ψ_i and ψ_f are the initial and final wavefunctions at the beginning/ending of the trajectory being considered. The classical action $S(q)$ is the time-integrated Lagrangian,

$S(q) = \int L(q)dt$, and the Lagrangian is the kinetic energy *minus* the potential energy $L = K - V$. If you have taken a classical mechanics physics course, you will be very familiar with the classical Lagrangian and action integrals.

The next main source of phase in atom interferometers is laser phase. The electric field of a plane-wave laser traveling in the $+z$ with wave-vector k and frequency ω is given by $E = E_0 \exp(kz - \omega t)$. Each time a photon from this laser is absorbed or emitted by an atom, the atomic wavefunction is shifted in phase by $\phi = \pm(kz - \omega t)$ where the sign \pm depends on whether the atom is absorbing or emitting the photon respectively. Brian's thesis has a very good section outlining the derivation of laser phase [23]. For each photon of light that the atom absorbs or emits, the complex-valued atomic wavefunction is multiplied by $e^{i\phi}$.

Typical atom interferometer experiments use two-photon transitions with counter-propagating beams. An atom absorbs a photon from one beam and receives a momentum kick in that direction, then stimulated emits a photon into the other beam and receives a momentum kick in the same direction as the first photon. For two-photon transitions, the laser phase imparted to the atom is the difference in optical phase between the up-going and down-going beams. In experiments like ours [67], we retro-reflect a laser beam onto itself using a mirror. The laser phase difference then only depends on the path length between the atom and the retro-reflecting mirror - any phase noise on the laser is common to the up-going and down-going beams and cancels out. The laser phase from the two photons is

$$\phi_{\text{laser}} = (k_1 - k_2)z - \delta t \tag{1.3}$$

where k_1 and k_2 are the wavevectors of the up-going and down-going laser beams, and δ is the frequency difference between up-going and down-going laser beams. This assumes the mirror is located at $z = 0$, though the total phase always ends up being independent of the absolute position and only relative changes are measured. If you are driving an n^{th} order process such as n^{th} order Bragg diffraction², the laser phase received is $n * \phi_{\text{laser}}$, as derived in Eq. 2.21 in Brian's thesis [23].

Another way to think about Eq. 1.3 is that counter-propagating laser beams form a standing wave $E_{\text{tot}} = E_0 \cos((k_1 + k_2)z - \delta t)$. Eq. 1.3 describes the position of the atom within this standing wave. Vibrations in the position of the retro-reflecting mirror cause the standing wave to vibrate, which causes noise in the laser phase imparted on the atom. Many atom interferometer avoid this noise source by making differential measurements between two interferometers driven by the same laser beam and laser pulses. This way, laser phase noise is common to both interferometers while differential phase shifts between the two interferometer can still easily be measured.

Some recent experiments have started using single photon transitions in alkaline-earth atoms such as Strontium because alkaline-earth atoms often have long-lived excited states. In this case, the imparted laser phase from the single photon is sensitive to phase noise on the laser itself, instead of just the path length between the atom and a retro-reflecting mirror. This phase noise can come from phase noise in the laser itself, or from any optical path variations in the entire optical path between the laser and the atoms including vibrating optical fibers or mirrors as well as index of refraction shifts in the air from air currents. Similar to two-photon transitions, differential measurements can still be made between two

²see Chapter 2 for more discussion of Bragg diffraction

interferometers that are addressed by the same laser propagating in-vacuum since the phase noise is common to both interferometers. There are plans to use this method to measure the stretching of space between the interferometers due to gravitational waves³, and tens of millions of dollars are being invested to develop proof-of-concept experiments.

The last contribution to an atom interferometer's measured phase is separation phase which comes from non-overlapping trajectories during the final interference pulse. See e.g. [43] for a good derivation of separation phase, including Figure 2.1 this reference which is very helpful for gaining an intuition of separation phase.

Explicitly, separation phase is given by:

$$\phi_{sep} = \bar{\mathbf{p}} \cdot \Delta \mathbf{x} \tag{1.4}$$

where \mathbf{p} is the average momentum between the two output ports of the atom interferometer, and $\Delta \mathbf{x}$ is the relative displacement between the wavefunctions.

This covers the standard phases measured in atom interferometers - free evolution phase and laser phase are the primary source of measured phase, and if trajectories don't overlap during final interference separation phase. If you're interested in using atom interferometers to measure physical quantities precisely, you'll find that there are many different systematic effects that also shift the phase measured by the atom interferometer. However, these effects are always captured by a careful analysis of free-evolution phase, laser phase and/or separation phase.

For example, there are effects that locally shift the potential energy of the atom such as magnetic field from nearby coils or gravitational potential from the mass of your optics table - these will be accounted for in the action integral in equation 1.2. Gradients in these potential energies result in forces that modify the wavefunction trajectories, which results in different integrated kinetic energy and different laser phases. If the forces on the two interferometer arms are not equal in magnitude, these forces can also lead to separation phase from non-overlapping trajectories when the interferometer closes.

As another example, the laser beam k-vector isn't exactly ω/c because we don't have perfect plane waves. Dust on optics, for example, causes intensity ripples in the beam and leads to perturbations of the k-vector of the laser. This causes different momentum kicks imparted to the atom, which in turn leads to different kinetic energy phases and different laser phases. The dust also perturbs the phase-front of the laser beam, which directly affects the laser phase imparted on the atom. If the k-vector kicks aren't the same on the two interferometer arms, the trajectories won't overlap and once again you also get a separation phase.

Much of the work in measuring the fine-structure constant, the main project of this thesis, is in chasing down systematic effects like this and doing detailed phase accounting. Actual effects from gradients in magnetic and gravitational potentials are analyzed in Chapter 4, and effects from laser beam inhomogeneities are analyzed in Chapter 5.

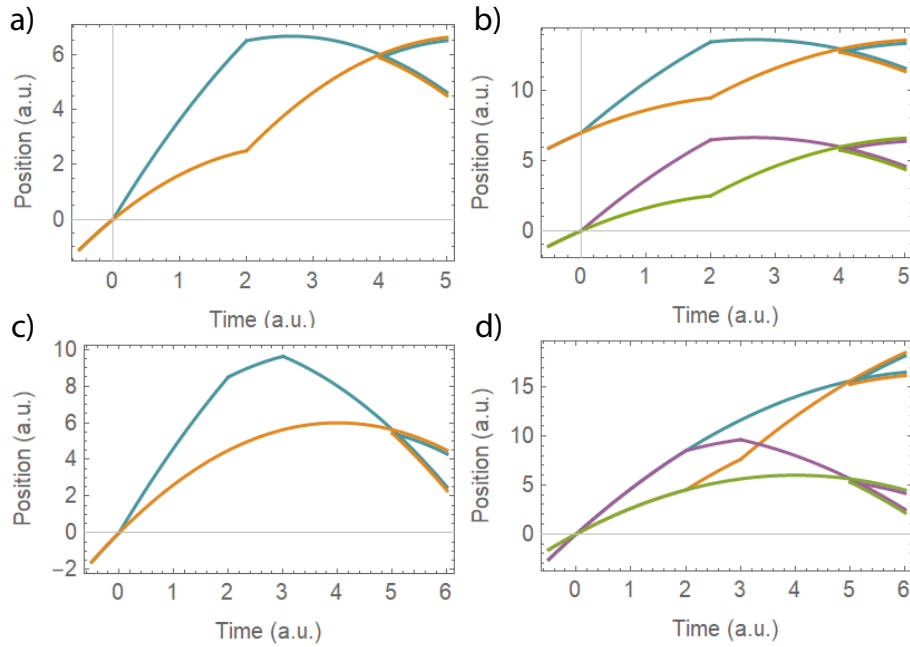


Figure 1.3: A number of common atom interferometer geometries. a) A Mach-Zehnder interferometer typically used to measure gravity, b) a dual Mach-Zehnder interferometer typically used to measure the difference in gravity between the upper and lower interferometers, ie. the gravity gradient. c) A Ramsey-Borde interferometer which measures a combination of gravity and recoil kinetic energy phase. d) A simultaneous conjugate Ramsey-Borde (SCI) interferometer measures the difference between two Ramsey-Borde interferometers, cancelling the gravitational phase and doubling the recoil kinetic energy phase.

Type	Free-evolution phase	Laser phase	Total phase
Mach-Zehnder	0	$-2 n k g T^2$	$-2 n k g T^2$
Dual Mach Zehnder	0	$2 n k T^2 \gamma \Delta x$	$2 n k T^2 \gamma \Delta x$
Ramsey-Borde	$4 k^2 n^2 T \hbar/m$	$-2 n k(g T (T+T1)+(4 k n T \hbar/m))$	$-2 n k(g T (T+T1)+(2 k n T \hbar/m))$
SCI	$-8 k^2 n^2 T \hbar/m$	$16 k^2 n^2 T \hbar/m$	$8 k^2 n^2 T \hbar/m$

Table 1.1: Phase derivations for the four common interferometers shown in Fig. 1.3. These are the first-order contributions to phase from free-evolution and laser phase - higher-order contributions are omitted for clarity.

1.2.3 Different types of interferometers

It's useful to practice going through the phase derivation of an entire interferometer to get a handle on how these calculations work in practice. I'll outline results for the free-evolution and laser phase ⁴ for the various common interferometer geometries in Fig. 1.3. This figure illustrates trajectories with arbitrary dimensionless units so the trajectories are not drawn to scale, but a negative acceleration is included to represent gravity in our experiment. The final interference pulse overlaps wavefunction amplitude from each interferometer trajectory into each output port in order to generate interference, and this is represented by multiple colored lines on the output ports.

Table 1.1 breaks down the phase contributions from free evolution and laser phase that each of these interferometer geometries measures. For example, the Mach-Zehnder (MZ) interferometer measures no free-evolution phase because the geometry is highly symmetric. This is easiest to see from a frame of reference freely-falling with the atoms. In a freely-falling frame, the trajectories are straight lines and each interferometer arm spends the same amount of time in each momentum state. The integrated phase from kinetic energy therefore cancels out upon the final interference. Since there is no gravitational potential energy in the freely-falling frame of reference, the total free-evolution phase is zero. This total phase is frame-independent - transforming back to the lab frame with trajectories shown in Fig. 1.3a) does not change the result. In total, the MZ interferometer is only really measuring laser phase, and for this reason people commonly refer to atom interferometry as nothing more than using a laser ruler to measure the position of your atoms. Similarly, the dual Mach-Zehnder (DMZ) interferometer measures the difference between the upper and lower MZs, and therefore measures the difference in gravity between the two regions as measured by the laser phase.

Although the MZ and DMZ interferometers only measure laser phase, more complicated interferometer geometries do measurable non-zero free-evolution phase. For example, the Ramsey-Borde (RBI) and simultaneous conjugate Ramsey-Borde interferometers (SCI) shown in Figs. 1.3c,d) both measure a non-zero free-evolution phase from kinetic energy. The laser phase contribution is exactly twice that of the kinetic energy, with an opposite sign. These types of cancellations or relationships between quantities are very common in analyzing atom interferometer geometries, and have led to much confusion... [60, 95, 59, 94].

³see e.g. [19] and much of the literature that this paper spawned

⁴Separation phase is typically a higher-order effect - here we're just trying to capture the first-order phase equation for different interferometer geometries.

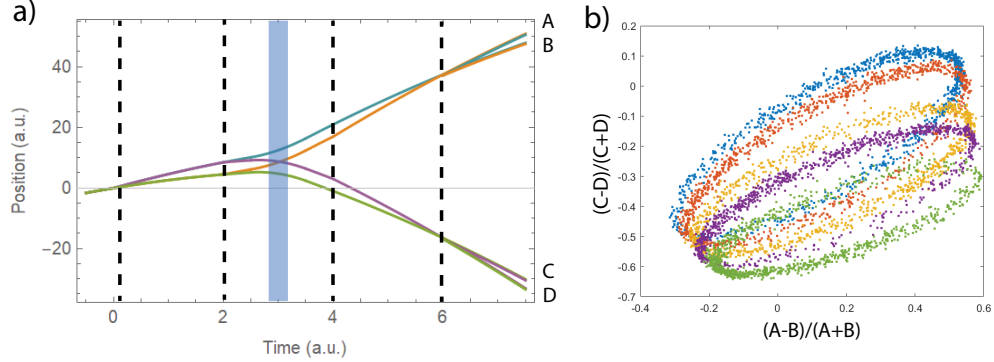


Figure 1.4: a) Illustration of the SCI interferometer with Bloch oscillations that is used in this thesis to measure the fine-structure constant α . The units are dimensionless and meant for illustration purposes only. Bragg diffraction pulses are indicated with black dotted lines, Bloch oscillations are indicated by the thick transparent blue line. b) Population in the output ports is plotted parametrically with A and B from the upper interferometer on the x-axis, and C and D from the lower interferometer on the y-axis. The different ellipses are experimental data from and SCI with $n = 3$, $N = 0$ and $T \approx 1$ ms. The different ellipses are taken at different heights in the chamber which is affecting ellipse contrast.

Note that the SCI is two RBIs inverted relative to one another - the inversion causes the gravitational phase to flip sign and cancel, whereas the kinetic energy phases add constructively. The gravitational phase is linearly proportional to k , whereas the kinetic energy is proportional to k^2 - the former therefore switches sign when $k \rightarrow -k$, whereas the latter remains the same.⁵

Another main takeaway from this analysis is that the RBI and SCI geometries both measure the recoil frequency $\omega_r = \hbar k^2 / 2M_{Cs}$. The SCI geometry is designed to cancel the gravitational phase so that all you measure is the kinetic energy phase. The following section discusses the SCI geometry in-depth and how we use it to measure the recoil frequency of Cesium.

1.3 Measuring the recoil frequency of Cesium

The focus of this thesis is in designing and building a new atom interferometer experiment to measure the recoil frequency in Cesium. This experiment is the successor to the results published by our group in 2018 [67]. Measuring the recoil frequency is the starting point for measuring the fine-structure constant α .

The actual interferometer geometry used for measuring α is shown in Fig. 1.4a). Four n^{th} -order Bragg diffraction pulses (dotted black lines) are used to split and recombine atoms

⁵This line of thinking about which terms switch sign upon inversion is very useful when trying to understand how to cancel terms experimentally and isolate the terms you are interested in. For example, averaging the gravity measurement from a MZ interferometer with momentum kicks up vs. down allows for a measurement of gravity while cancelling systematic effects from magnetic field gradients and light shifts [96].

as in a standard SCI interferometer. For additional momentum transfer, N Bloch oscillations are also used in the middle of the sequence (thick light blue line) to accelerate the upper interferometer upwards by $2N\hbar k$ recoil momentum, and the lower interferometer downwards by $2N\hbar k$. Multiple frequencies of light are used for Bloch oscillations and for the final two Bragg diffraction pulses so that we can interact with multiple velocity classes of atoms at the same time. The frequency splitting ω_m between two of these optical frequencies is referred to as the ‘modulation frequency’.

Figure 1.4b) shows example experimental data from an SCI interferometer. Population in the output ports is plotted parametrically in order to generate ellipses. Random vibrations of the experiment cause phase shifts that are common mode to both the upper and lower interferometers, causing data points to walk randomly about the ellipse. The differential phase Φ_{diff} between the interferometers causes the ellipse to rotate and is not affected by the vibrations. Note that the maximum contrast for our SCI ellipse is 50% - there are ‘junk’ ports not shown in the interferometer diagram that don’t coherently interfere, and half of the atomic population ends up in these junk ports. See Brian Estey’s thesis for a more detailed discussion of ellipse data and ellipse fitting [23].

If one includes effects from the modulation frequency ω_m when calculating laser phase, and also includes the effects of Bloch oscillations, the total phase equation for the SCI interferometer with Bloch oscillations can be written as

$$\Phi_{\text{diff}} = 16n(n + N)\omega_r T - 2n\omega_m T \quad (1.5)$$

where the recoil frequency is defined as $\omega_r = \hbar k^2 / (2m)$. Experimentally, we carefully tune the modulation frequency ω_m so that the second term in this equation nearly cancels the first. We alternate data taking between $\Phi_{\text{diff}} = \pm\pi/2$ so that the ellipses generated by the data are open ellipses that can be easily fit. The differential phases extracted from the two ellipses are used to determine the exact ω_m needed to give exactly zero phase in Eq. 1.5.

In reality, the phase equation 1.5 is only approximate. For example, the SCI geometry cancels phase from gravitational acceleration, but the SCI is still affected by gradients in gravitational acceleration because the two interferometers slightly different gravitational accelerations. Gravity gradient phase needs to be removed from the measurement in order to get an accurate measurement of the recoil frequency ω_r . Diffraction phase is another phase contribution that needs to be corrected for before getting an accurate value of ω_r . Diffraction phase is a phase shift from Bragg diffraction that is sensitive to the velocity of the atom relative to the optical lattice and the intensity of the Bragg laser pulse. Data is taken at multiple values of T in order to fit the diffraction phase and remove it from the measurement. See Chenghui’s thesis [99] for further discussion of both of these effects as well as a phase derivation including higher-order corrections that I have ignored here.

1.4 From the recoil frequency to dark matter

So why do we spend so much time trying to measure the recoil frequency of Cesium?⁶ It turns out that knowledge of the ratio of \hbar to the mass of some atom is the limiting quantity

⁶This is not meant to cause an existential crisis, it’s only a pedagogical question...

in the known accuracy of the fine-structure constant.⁷ This is in turn the limiting quantity in theorists’ prediction of the electron magnetic anomaly $g_e - 2$. Experimentalists are also measuring $g_e - 2$ at a similar level of accuracy, which means that our value of \hbar/M_{Cs} allows for a more precise experiment/theory comparison for $g_e - 2$ [25, 38]. The theorists’ calculation of $g_e - 2$ involves all known particles in the Standard Model of physics. As a result, the three-way comparison between our measurement of the fine-structure constant (as well as the competing French measurement using \hbar/M_{Rb} [56]), Gerald Gabrielse’s measurements of $g_e - 2$ [38, 25], and the theorist prediction of $g_e - 2$ using our measurement of α [3], all combine to make one of the most precise tests of the Standard Model, and the most precise test of quantum electrodynamics (QED).

With recent improvements in both the $g_e - 2$ experimental and theoretical predictions, the discrepancy between the most recent measurements of α now limits the comparison by nearly an order of magnitude [67, 56, 25], and the error bars in the theoretical value of $g_e - 2$ are mostly from error vars in the value of α . It is therefore good timing for us to publish a new trustworthy value, since many people are waiting for a new measurement...

There is a very closely related experiment measuring the muon magnetic moment $g_\mu - 2$ and comparing the value to its theory prediction [17]. Although this measurement and theory prediction are only around the 10^{-7} level in accuracy, the QCD terms for the muon experiment are much much larger than in the electron $g_e - 2$ experiment. The muon experiment is *not* limited by knowledge in the fine-structure constants because the experimental and theoretical error bars are only at the 10^{-7} level, so our improved measurement of α does not directly contribute to this experiment, though the two experiments are very complementary tests of the Standard Model.

There’s a very detailed summary of these comparisons in the CODATA papers (see e.g. the most recent ‘2018’ paper published in 2021 [82]). I refer the reader there for a complete treatment of the subject, but I’ll summarize some results below.

First, let’s look closer at the α side. The definition of the Rydberg energy R_∞ is given by:

$$hcR_\infty = \frac{1}{2}\alpha^2 m_e c^2 \tag{1.6}$$

which can be rearranged to give an expression for α :

$$\alpha = \left(\frac{2hR_\infty}{m_e c} \right)^{1/2} \tag{1.7}$$

Penning traps can be used for precision mass ratio measurements, so the ratio h/m_e can be expanded into a product of mass ratios and the quantity \hbar/M_{Cs} , since the latter is easier to measure experimentally. Eq. 1.7 can therefore be re-written as

$$\alpha = \left(\frac{2R_\infty}{c} \frac{u}{m_e} \frac{M_{Cs}}{u} \frac{h}{M_{Cs}} \right)^{1/2} \tag{1.8}$$

⁷The fine-structure constant is famous for being nearly equal to 1/137. It is also one of the few dimensionless fundamental constants, and for this reason it is inscribed on the ‘Golden record’ that was placed on the two Voyager probes in order to (hopefully) communicate with aliens after the leaving the solar system.

Constant	Relative error
Rydberg constant	1.9 e-12
Electron mass u	2.9 e-11
Cesium mass in u	6 e-11
Speed of light	exact

Table 1.2: A table of the current know accuracies of other constants (besides \hbar/M_{Cs}) needed to determine the value of the fine-structure constant. The rydberg constant, electron mass in u, and speed of light are all taken from current CODATA recommended values [82]. Cesium mass taken from [5].

where the quantity limiting the overall accuracy of α is the ratio \hbar/M_{Cs} . A summary of the current level of accuracy in each of these quantities is summarized in Table 1.2. Note that the square root in Eq. 1.8 means that these errors are divided by two when translating to errors in the value of α .

One thing to note is that if our experiment measured \hbar/M_{Cs} much more accurately, the quantity limiting the accuracy of α would soon be M_{Cs}/u ! Klaus Blaum has promised Holger to publish a new measurement of M_{Cs}/u before we are limited by its value, and they've gotten very good at using their penning traps for these kinds of measurements - see e.g. [39] for a recent result. Note that if Klaus Blaum waited to publish his results after our new measurement was published, then he could steal our thunder and claim he is now measuring the fine-structure constant and searching for dark matter! Then we would only really be measuring the mass of Cesium very precisely, since \hbar is now a fixed quantity in the new SI unit system [82].⁸

So that's how our experiment relates to the overall value of the fine-structure constant. Next, let's look at the theory prediction for the anomalous magnetic moments, and see how the value of α fits in.

A lepton magnetic dipole moment is related to its spin via

$$\mu_l = g_l \frac{e}{2m_l} s \quad (1.9)$$

where l indicates the lepton type, which can be electron, muon or tau. The gyromagnetic ratio g_l is very nearly two, but deviates slightly from 2 due to the interaction between the lepton and a 'virtual sea' of Standard Model particles. The exact calculation of g_l leads to:

$$|g_l| = 2(1 + a_l) \quad (1.10)$$

where

$$a_l = a_l(QED) + a_l(QCD) + a_l(weak) \quad (1.11)$$

⁸Of course the community would understand that our measurement required orders of magnitude more dedicated work. Also, this would only be an issue if we improve our previous measurement [67] by an order of magnitude.




Figure 1.5: The two most important Feynman diagrams relevant to our fine-structure constant experiment. a) The fine-structure constant serves as the coupling constant between an electron interacting with a photon. b) The first radiative correction in Eq. 1.12, which is the first-order QED correction to $g_l - 2$.

ie. the calculation is broken into terms from quantum electrodynamics (QED), quantum chromodynamics (QCD), and the weak force separately. Further, the QED term can be broken into

$$a_l(QED) = \sum_{n=1}^{\infty} C_l^{(2n)} \left(\frac{\alpha}{\pi}\right)^n \quad (1.12)$$

where the $C_l^{(2n)}$ coefficients are calculated from summing over all $(2n)$ -order Feynman diagrams for the given Lepton. For reference, the C_e^2 diagram is shown in Fig. 1.5. You can see how it is the first so-called ‘radiative’ correction to the basic interaction between an electron and photon. The $(2n)$ in $C_l^{(2n)}$ refers to the fact that this ‘first-order’ correction actually involves two new nodes in the Feynman diagram. For higher orders, there are many possible diagrams with $(2n)$ nodes. By the 10th order diagrams ($n=5$), there are already 6,354 Feynman diagrams that need to be summed in the electron diagrams alone, not including the muon or tau contributions to these terms. See the CODATA paper and references therein for more details [82].

For both the electron and muon experiments, the vast majority of the magnetic anomaly comes from the $a_l(QED)$ term. The $g_e - 2/\alpha/theory$ comparison is at the 10^{-10} level, which means we are resolving Eq. 1.12 at the 10^{-10} level. The $g_\mu - 2$ experiment/theory comparison is only at the 10^{-7} level, and therefore the muon experiment is less sensitive to modifications to QED beyond the Standard Model.

For QCD, on the other hand, the $a_e(QCD)$ term is only a 10^{-10} contribution to the electron $g_e - 2$ - these terms are therefore only being resolved at $O(1)$ with our current sensitivities. The $a_\mu(QCD)$ term, however, is much larger - it is a 10^{-4} contribution to the total shift. So while we are just starting to resolve the QCD effects in the $g_e - 2$ comparison, the $g_\mu - 2$ experiment is resolving the QCD correction at the 0.1% level.

The $a_e(weak)$ term is a 10^{-11} contribution, and therefore is not yet resolveable in our experiment but will soon be. The $a_\mu(weak)$ term is a 10^{-6} correction, and is therefore already resolveable in the muon experiment at the $O(10\%)$ level.

In total, the electron and muon experiments are sensitive to different aspects of the

Standard Model, and are therefore nicely complimentary in creating a very broad test of the Standard Model. As difficult as it is experimentally for use to make an improved measurement of α , we're broadly testing the limits of the Standard Model by doing so. These experiments are among the most promising searching for physics beyond the standard model - physicists have been improving these measurements for decades and will continue to improve them for decades to come.

Chapter 2

Bragg diffraction and Bloch oscillations - theory and numerics

2.1 Introduction

Bragg diffraction [9] and Bloch oscillations [7] were first discovered at the beginning of the 20th. Bragg was 22 and Bloch was 24 when they published these groundbreaking papers.¹ Bragg diffraction and Bloch oscillations are quite similar solutions for how waves propagate in a periodic potentials, though they were first discovered in very different physical scenarios. Bragg diffraction was first identified in x-ray scattering off of a crystal, whereas Bloch oscillations were identified in how electrons interacted inside a crystal. 100 years later, Bragg diffraction and Bloch oscillations are used in our experiments in more of an engineering sense than a scientific one - the physics is fairly well understood, and we need to optimize the processes for the application of manipulating atomic trajectories in our SCI interferometer.

Compared to the original discoveries in solid state physics, there are some interesting things we can do only in our atomic physics experiments. For example, we can create two superposed lattices moving at different (and easily controllable) velocities, whereas in solid state there is no parallel to this. This is possible because we form optical lattices using counter-propagating lasers such that the velocity of the standing wave is determined by the frequency difference of the lasers. By superposing multiple frequencies of light, we can create multiple superposed standing waves at different velocities. We use this trick to control two interferometers as once, with a single laser pulse.

Another interesting case is when the velocities of the two superposed lattices become degenerate, or similarly when they start degenerate and then are accelerated in opposite directions. There is no analogue to this in solid state physics. This is explored in Chapter 3, where it turns out there is very interesting lattice physics that can be used in atom interferometers.

In this chapter, I outline derivations of the Bragg diffraction and Bloch oscillation Hamiltonians. Some different regimes of Bragg diffraction and Bloch oscillations are studied in order to gain an understanding of why our α experiment uses the regimes that it does. I'll

¹And what have you been doing with your life?

also show some useful results/takeaways regarding systematic effects from Bloch oscillations in our experiment - in particular, I outline a more rigorous justification of why systematic effects from Bloch oscillations in the previous generation experiment [67] were negligible. The chapter ends with a section summarizing the dimensionless unit system used to numerically integrate the Schrodinger equation, as well as a section on how to convert experimental laser power to lattice depth in our Hamiltonians.

2.2 Optical lattice Hamiltonian

Both Bragg diffraction and Bloch oscillations are dynamics that arise from an atom in an optical lattice. The term ‘optical lattice’ refers to the fact that two laser beams, far-detuned from any energy transitions in the atom, are interfering such that they create a standing wave of light. Atoms in our experiment see an energy shift depending on where they are in the optical standing wave, so the potential energy landscape that the atom sees is periodic - hence the term ‘optical lattice’.

For a good derivation of Bragg diffraction Hamiltonian, see Chenghui’s thesis [99], section 2.4. He walks through the rotating wave approximation, adiabatic elimination of the excited state, and then using a discrete momentum state basis for numerical simulation. I’ll outline some of the main results here, but see his thesis for further discussion.

Before making any assumptions, the Hilbert space of Cesium atoms in our experiment (really of the lone electron in the outermost shell) is represented by $|nLJFm_F, \vec{p}\rangle$ where n is the over-all energy level, L is the angular momentum quantum number, J is the total electron angular momentum after accounting for fine-structure coupling, F is the electronic angular momentum plus nuclear angular momentum, m_F is the magnetic sublevel, and \vec{p} represents the momentum of the atom.²

Fine-structure atomic notation is typically given by $n^{2s+1}L_J$, with $L = S, P, D, etc..$ In this notation, our 852 lasers in lab are near the Cesium ‘D2’ line ($6^2S_{1/2} \rightarrow 6^2P_{3/2}$). Our lasers are typically far-detuned from the ‘D1’ line around 895 nm ($6^2S_{1/2} \rightarrow 6^2P_{1/2}$), so we ignore effects from the D1 line.³ Transitions to other values of n or L are very very far-detuned from our laser frequencies and can be ignored. So, we have fixed values of $n, L, S,$ and J for D2 transitions.

In order to do coherent physics, we can’t populate excited D2 states of Cesium which only have lifetimes of 30 ns. Instead, we use two-photon transitions between ground states. For Bragg diffraction or Raman transitions, we typically operate our interferometry lasers on the order of $\Delta \approx 10$ GHz detuned from the D2 line, whereas the linewidths of states in the D2 line are around $\Gamma \approx 5.2$ MHz [77]. Since $\Delta \gg \Gamma$, populations in the excited states are therefore very small and we can adiabatically eliminate the excited state. See Chenghui’s thesis for more details about adiabatic elimination [99]. The only remaining possible (two-photon) transitions are between different F and m_F states within the ground state $6^2S_{1/2}$ -

²Note that the position and momentum cannot be known simultaneously due to the Heisenberg uncertainty principle. We use \vec{p} here because the momentum of the atom is more relevant for the following sections than position, but we could have equivalently used \vec{x} instead.

³The cavity interferometer experiment in our group needs to take the D1 line into account for their very far-detuned lattice laser [98].

all other internal degrees of freedom are effectively fixed.

Atom interferometers with alkali atoms typically use Raman transitions or Bragg diffraction for beamsplitters. Raman transitions change hyperfine state from $F = 3 \rightarrow F = 4$ or vice versa. Bragg transitions remain in the same hyperfine state. The difference between driving Raman transitions and Bragg diffraction comes down to 1) the frequency difference between the two laser frequencies used for the two-photon transition, and 2) the polarization of light. To drive Bragg diffraction, the frequency detuning of the lasers will be related to doppler detunings of how fast the atoms are moving in our lab (1 m/s velocity is on the order of 1 MHz frequency difference), and also on the Bragg resonance condition which requires frequency shifts on the order of 10's of kHz. Raman transitions require a frequency difference near $F = 3$ to $F = 4$ ground state transition, which for Cesium is about 9.8 GHz.

Here, I will consider only Bragg transitions - for further discussion of Raman transitions, as well as the polarization requirements for both Raman transitions and Bragg diffraction, see Chenghui's thesis [99]. For Bragg diffraction, we only need to consider transitions between m_F states within a single F state. These transitions are determined by selection rules related to the polarization of light. We initialize our atoms in the $m_F = 0$ state, and choose polarizations of light that only drive two-photon $m_F = 0 \rightarrow m_F = 0$ transitions. Effects of leakage to other m_F states would have to be derived from a more general Hamiltonian and Hilbert space that include other internal states. Chenghui's thesis outlines how to calculate these leakage transitions [99].

All of the above means that there are no free internal degrees of freedom left in the experimental situation that we set up - the internal state is exactly identical between the initial and final states, and the only degrees of freedom for the atom that remain are the position or momentum degrees of freedom. In a momentum basis, we are therefore solving a Schrodinger equation of the form:

$$i\hbar \frac{\partial}{\partial t} \psi(p, t) = \hat{H} \psi(p, t) \quad (2.1)$$

which describes time evolution of a wavefunction with only momentum degrees of freedom. This equation with a position representation instead of momentum is used in intro quantum courses to look at wavefunction solutions for particles in square well potentials, harmonic oscillator potentials, etc. and is most physicists' first introduction to the Schrodinger equation.

The Hamiltonian for an atom in this 'optical lattice' based on the two counter-propagating laser beams in our experiment is given by:

$$H = \frac{\hat{p}^2}{2M_{Cs}} + \hbar\Omega \cos((k_1 + k_2)\hat{x} - \delta t) + \hbar\omega_{Stark} \quad (2.2)$$

Here, Ω denotes the lattice depth and ω_{Stark} denotes the total energy shift from AC Stark shift/light shift. The frequency difference between the two laser beams is given by $\delta = \omega_2 - \omega_1$. The lasers in our experiment are very nearly counter-propagating, so $k_1 \approx k_2$ to a very high degree of accuracy.⁴

⁴I'm ignoring the sign of the k-vector here and only using the amplitude - in reality both k -vectors are vector quantities with directions as well as magnitudes.

In our experiments, we typically think of the Schrodinger equation as a one-dimensional problem along the laser propagation axis, and the transverse degrees of freedom are treated classically. Why is this assumption accurate? The system starts behaving quantum mechanically when there is spatial structure of the potential on a length scale comparable to or smaller than the spatial extent of the atomic wavefunction. In our experiment, our atoms have a velocity spread of around $\sigma_v \approx 0.05v_r$ recoil velocities, where the recoil velocity $v_r = \hbar k/M_{Cs}$ of Cesium with 852 nm light is approximately 3.5 mm/s. This corresponds to a thermal deBroglie wavelength of around:

$$\lambda_{dB} = \frac{h}{\sqrt{2\pi M_{Cs} k_B T}} = \frac{h}{\sqrt{2\pi M_{Cs}^2 \sigma_v^2}} \approx 6\mu\text{m} \quad (2.3)$$

We must treat the on-axis dynamics quantum mechanically in our experiment because the thermal deBroglie wavelength of our particles is much longer than the spatial period of the optical lattice, $\lambda_{dB} \gg \lambda/2$. In our experiment, the deBroglie wavelength of atoms along the laser propagation axis is about $6\mu\text{m}$, which is around a factor of 20 larger than the optical lattice periodicity.

Transverse to the beam's propagation, our atoms are much hotter with a velocity spread of around $1.5 v_r$ and a thermal deBroglie wavelength of around $\lambda_{dB,transverse} \approx 200 \text{ nm}$. The length scale of changes in the optical potential transverse to the beam are much larger than this - for example the transverse beam profile of a typical Gaussian beam in our experiments changes on the order of several millimeters. Even for a laser beam with intensity noise from dust or other optical imperfections, the length-scale of changes in the optical potential must be greater than $\approx 50\mu\text{m}$ because any smaller length scale perturbations have a very short Rayleigh range and won't reach the atoms that are a meter or more from the perturbing optic. Similarly, the transverse effects from beam misalignment are still orders of magnitude too small to be treated quantum mechanically. If the counter-propagating beams have a 1 mRad angle between them, the transverse periodicity of the tilted lattice is still $\approx 1000\lambda \gg \lambda_{th,transverse}$. We co-align the beams much better than this experimentally. In total, transverse dynamics can safely be treated classically.

The k -vector is commonly written using $2k = (k_1 + k_2)$. Analysis is typically done in the frame of reference co-moving with the atoms, in which $k_1 + k_2$ would differ by 10^{-11} for recoil-level momentum kicks, and most experiments aren't actually making absolute measurements at the 10^{-11} level. From here on I will write $k_1 + k_2$ as $2k$, but for an accurate phase derivation of our interferometer these effects need to be taken into account.

Note also that when the two counter-propagating beams have equal intensity, the Stark shift and lattice depths in Eq. 2.51 are identical $\omega_{Stark} = \Omega$. However when the counter-propagating intensities are not equal $\Omega_1 \neq \Omega_2$, the stark shift and the lattice depth start to deviate.⁵

$$\omega_{Stark} = \frac{\Omega_1^2 + \Omega_2^2}{4\delta} \quad (2.4)$$

and

$$\Omega = \frac{\Omega_1 \Omega_2}{2\delta}. \quad (2.5)$$

⁵ Ω_i here denotes the single-photon resonant Rabi frequencies for a single beam

More progress can be made on the problem by using Bloch's theorem, since our potential is periodic in x with a period of $\lambda = \pi/k$, or equivalently as spatial frequency of $2k$. Since \hat{x} and \hat{p} don't commute, we can't find a diagonal/commuting set of basis states. However, plane-wave atomic states form a very natural basis for the problem, as you will see. Atoms in this experiment are much colder than a recoil velocity along the laser axis, so plane wave momentum eigenstates are a better choice of basis states than localized position eigenstates. For an analysis using position eigenstates instead, see reference [87] for treatment of 'Wannier-Stark' states. Wannier-Stark states are a better basis for e.g. the cavity interferometer in our group, in which atoms are around one recoil velocity temperature and have very deep lattice depths.

Bloch's theorem states that the wavefunction can be decomposed into Fourier components with multiples of $2k$ momenta:

$$\psi(x) = \sum_{n=0}^{\infty} e^{ik_q x} u_{n,k_q}(x) \quad (2.6)$$

where u_{n,k_q} is a Bloch wave for the n^{th} energy level eigenstate. $\hbar k_q$ is the quasi-momentum, which is essentially the momentum of the atomic state modulo $\hbar k$. Quasi-momentum will be discussed more in section 2.5.

This is a convenient basis for e.g. electrons in a crystal, especially when the periodic potential is complex within a period and the Bloch wave basis can capture a lot of that complexity. However, for our experiment with a sinusoidal potential, it ends up being much more convenient to work in a simple Fourier basis:

$$\psi(x, t) = e^{ik_q x} \sum_{m=-\infty}^{\infty} g_m(t) e^{2imkx} \quad (2.7)$$

where the time dependence is now explicitly shown. Eq. 2.7 describes a ladder of momentum states each separated by $2k$ momentum, a very convenient basis because the laser can only impart $\pm 2\bar{k}$ of momentum, or integer multiples thereof. The beauty of this basis can be seen when we plug this ansatz into the Schrodinger equation in Eq. 2.1. You get a system of coupled equations - isolating the equation for a single state coefficient g_m :

$$\begin{aligned} i\hbar \dot{g}_m(t) = & \frac{\hbar^2(k_q + 2mk)^2}{2M_{Cs}} g_m(t) + \hbar\omega_{Stark} g_m(t) \\ & + \frac{\hbar\Omega}{2} (e^{-i\delta t} g_{m-1}(t) + e^{i\delta t} g_{m+1}(t)) \end{aligned} \quad (2.8)$$

The Fourier basis states are a perfect choice for this problem because the optical potential only couples neighboring momentum states.

Since the basis used in our Schrodinger equation is now a vector of discrete momentum states, the Hamiltonian can be represented in matrix form as:

$$H = \begin{pmatrix} \ddots & \ddots & \ddots & & \ddots & & \\ \ddots & \alpha_{-2} & \beta & 0 & & \ddots & \\ \ddots & \beta^* & \alpha_{-1} & \beta & 0 & & \ddots \\ & 0 & \beta^* & \alpha_0 & \beta & 0 & \\ \ddots & & 0 & \beta^* & \alpha_1 & \beta & \ddots \\ & \ddots & & 0 & \beta^* & \alpha_2 & \ddots \\ & & \ddots & & \ddots & \ddots & \ddots \end{pmatrix} \quad (2.9)$$

where

$$\alpha_m = \frac{\hbar^2(k_q + 2mk)^2}{2M_{Cs}} + \hbar\omega_{Stark} \quad (2.10)$$

and

$$\beta = \frac{\hbar\Omega}{2}e^{i\delta t} \quad (2.11)$$

This is understood to be the Hamiltonian matrix that acts on a state vector of discrete Fourier states, where the index m denotes the momentum of the basis state $p_m = 2m\hbar k$. Tridiagonal matrices are extremely nice for numerical simulation because linear equations can be solved in $O(3n)$ time, where n is the number of basis states used, instead of $O(n^2)$ which would be the case for a general matrix equation.

Another brief aside - why can we treat our atoms as plane wave states, when experimentally we know they have finite size/temperature? Suppose we have an atom with some finite spatial extent. In order to write this state down mathematically using the formalism that we just outlined, one would need to integrate over all momentum states, with each momentum state having its own g_k coefficient:

$$\psi_{\text{finite}}(x) = \int_{k=-\infty}^{\infty} g_k e^{2ikx} dk \quad (2.12)$$

Because of Bloch's theorem and the spatial periodicity of our potential, we know that we can completely capture all dynamics within the momentum $[-k_q, k_q)$. In our Fourier basis, this equates to:

$$\psi_{\text{finite}}(x) = \int_{k=-k_q}^{k_q} e^{ik_q x} \sum_{m=-\infty}^{\infty} g_{m,k_q} e^{2ikx} dk \quad (2.13)$$

which is equivalent to Eq. 2.7 but integrated over all quasimomentum k_q . If we plug this into Eq. 2.1, we will still end up with Eq. 2.8, one for each k_q state. This occurs because every operation is linear, and because the optical potential can only couple states separated by integer multiples of $2\hbar k$ momentum. Even calculating the wavefunction interference at the end is linear, since these basis states of different k_q are orthogonal.

In total, it is entirely equivalent to treat atoms as an ensemble of randomly sampled k_q states according to the velocity spread of the atom ensemble, as it is to treat the atoms




Figure 2.1: Diagram illustrating a 3rd order Bragg diffraction process in terms of energy levels. We are far-detuned from any excited states, and the frequency difference between the two beams 2δ satisfies the $n = 3$ Bragg resonance condition. Figure adopted from Brian Estey's thesis [23].

more generally as finite size and temperature quantum wavefunctions. Since the former is so much easier, this is what is used for numerics and analytics. The above argument outlines one of the core justifications for why the Monte Carlo in Chapter 5 works as it does.

2.3 Bragg diffraction Hamiltonian

For atoms initialized in a single momentum state, Bragg diffraction will cause Rabi flopping between two momentum states separated by $2n\hbar k$ momentum, where n is the Bragg diffraction order. Fig. 2.1 illustrates what's happening to the atomic state during an $n = 3$ Bragg diffraction process.

Bragg diffraction occurs when the velocity of an optical lattice relative to a matter wave (very nearly) matches a Bragg resonance condition. Bragg resonance occurs when the absorption/emission of photon momenta conserves energy and momentum of the before and after atomic state. In words, the energy difference of the photons absorbed must equal the kinetic energy difference of the atom before and after the pulse, and similarly with the momentum of the photons and the atom. Since we know the photon momentum transferred is $2n\hbar k$, energy conservation becomes:

$$n\hbar\delta = \frac{M_{Cs}(v_i + 2n\hbar k/M_{Cs})^2}{2} - \frac{M_{Cs}v_i^2}{2} \quad (2.14)$$

where v_i is the initial velocity of the atom, and $\delta = \omega_1 - \omega_2$ is the frequency difference between the two lasers. Simplifying, we arrive at the Bragg resonance condition:

$$\delta = 2kv_i + 4n\omega_r \quad (2.15)$$

where $\omega_r = \hbar k^2 / (2M_{Cs})$ is the recoil frequency of cesium. Eq. 2.15 is very intuitive if you relate it to the velocity of the lattices by dividing by $2k$:

$$v_{latt} = v_i + nv_r \quad (2.16)$$

Since the atom's final velocity will be $v_i + 2nv_r$, Eq. 2.16 says the lattice velocity should be exactly halfway between the initial and final velocities for resonance. In other words, the atom is diffracting/reflecting off of a 'grating' formed by our optical lattice. This understanding is extremely useful when working with Bragg diffraction experimentally or numerically.

For a more in-depth introduction to Bragg diffraction, see Chenghui's thesis [99], or for a lot of details see Holger's old paper [57].

2.3.1 Single frequency Bragg diffraction

In our experiments, we adiabatically ramp the lattice depth on and off in a short amount of time such that the integrated effect of the pulse is roughly a ' $\pi/2$ pulse' between the two momentum states of interest. In order to arrive at the Bragg Hamiltonian starting from Eq. 2.8, the only changes we add are 1) there is a Gaussian temporal profile of our lattice depth, and 2) the optical lattice has some velocity offset from the initial atomic state to satisfy the Bragg resonance condition. The coupling term in the Hamiltonian therefore becomes:

$$\frac{\hbar\Omega}{2} e^{-t^2/2\tau^2} (e^{-i2kv_{latt}t} g_{m-1}(t) + e^{i2kv_{latt}t} g_{m+1}(t)) \quad (2.17)$$

where v_{latt} is the lattice velocity relative to the rest frame, and Ω is understood to be the peak lattice depth, independent of time, since the time profile is explicitly written separately. For n^{th} order Bragg, the lattice velocity at resonance for an atom starting with zero velocity is $v_{latt} = nv_r$, as was shown in Eq. 2.16.

In the Monte Carlo simulations discussed later in this thesis, it is much more convenient numerically to use a reference frame such that the initial velocity⁶ v_0 with respect to the initial momentum state is instead represented in the velocity of the optical lattice. We move to a frame of reference such that the atom velocity exactly lines up with one of our momentum ladder states, and instead the lattice moves at a slightly different velocity than the resonance condition. In this case, we use $v_{latt} = nv_r - v_0$ and set $k_q = 0$. The full Schrodinger equation therefore reads:

$$i\hbar\dot{g}_m(t) = \frac{\hbar^2(2mk)^2}{2M_{Cs}} g_m(t) + \hbar\omega_{Stark} g_m(t) + \frac{\hbar\Omega(t)}{2} (e^{-i2kv_{latt}t} g_{m-1}(t) + e^{i2kv_{latt}t} g_{m+1}(t)) \quad (2.18)$$

or in dimensionless units for simulation,

⁶i.e. the velocity by which you are off-resonant from a Bragg transitions, which is typically $\ll v_r$

$$i\dot{g}_m(t) = (2m)^2 g_m(t) + \omega_{Stark} g_m(t) + \frac{\tilde{\Omega}(t)}{2} (e^{-i4\tilde{v}_{latt}t} g_{m-1}(t) + e^{i4\tilde{v}_{latt}t} g_{m+1}(t)) \quad (2.19)$$

where $\Omega(t) = \Omega e^{-t^2/2\tau^2}$. All variables in the dimensionless equation are assumed dimensionless, and it is this dimensionless form that is used for simulation in the Monte Carlo code. See Section 2.8 for details about how to derive dimensionless equations.

2.3.2 Multi-frequency Bragg diffraction

In our SCI interferometer for the alpha measurement, we use so-called ‘multi-frequency’ pulses for the third and fourth Bragg diffraction pulses. For details of the experimental implementation, see Chapter 4. For our purposes here, we have one frequency of light traveling in one direction, and two frequencies traveling in the opposite direction. Interference between these beams forms two moving optical lattices superposed on one another, allowing us to address two velocity classes of atoms at the same time.⁷

If we want to describe two lattices moving at velocities $\pm v$, we can write the sum of two cosine waves as:

$$\cos(2kx - 2kvt) + \cos(2kx + 2kvt) = 2 \cos(2kx) \cos(2kvt) \quad (2.20)$$

From this you can see that a beat note that forms between the two interfering lattices. Experimentally, it is more convenient to work in frequency instead of velocity - the ‘splitting velocity’ between the two lattices is directly related to the all-important modulation frequency ω_m by $\omega_m = 2kv$. Note that the splitting velocity describes lattices moving at $\pm v$, so the velocity difference between the two lattices is $2v$. The modulation frequency ω_m is all-important because it is the knob we turn experimentally to actually measure the recoil frequency, as discussed in Chapter 1.

In terms of ω_m , the Schrodinger equation for multi-frequency Bragg is given by:

$$i\hbar\dot{g}_m(t) = \frac{\hbar^2(2mk)^2}{2M_{Cs}} g_m(t) + \hbar\omega_{Stark} g_m(t) + \frac{\hbar\tilde{\Omega}(t)}{2} \cos(\omega_m t) (e^{-i2kv_{latt}t} g_{m-1}(t) + e^{i2kv_{latt}t} g_{m+1}(t)) \quad (2.21)$$

or in dimensionless units for simulation,

$$i\dot{g}_m(t) = (2m)^2 g_m(t) + \omega_{Stark} g_m(t) + \frac{\tilde{\Omega}(t)}{2} \cos(\omega_m t) (e^{-i4\tilde{v}_{latt}t} g_{m-1}(t) + e^{i4\tilde{v}_{latt}t} g_{m+1}(t)) \quad (2.22)$$

⁷Actually, in the experiment we end up with many more (unwanted) lattices, but for now we will focus on the two that are intentionally interacting with the atoms. See Section 4.6.3 for details.




Figure 2.2: Both panels show wavefunction amplitude evolution during an $n = 3$ Bragg $\pi/2$ pulse. Two regimes are shown: a) the 'quasi-Bragg' regime that we use experimentally $\tau = 2, \mu\text{s}$, $\Omega_{\text{bragg}} = 11.335 \text{ E} \cdot \text{r}$, and b) the 'Bragg regime' with $\tau = 5 \text{ ms}$, $\Omega_{\text{bragg}} = 1.63 \text{ E} \cdot \text{r}$. The legend indicates which color corresponds to which momentum basis state.

where again, recall $\Omega(t) = \Omega e^{-t^2/2\tau^2}$. All variables in the dimensionless equation are assumed dimensionless, and it is this dimensionless form that is used for simulation in the Monte Carlo code. See Section 2.8 for how to derive dimensionless equations.

2.4 Different regimes of the Bragg Hamiltonian

The more slowly a Bragg diffraction process is driven, the more precisely momentum can be transferred between two states without populating intermediate states. For momentum processes driven in a characteristic time τ , the 'Bragg regime' corresponds to $\tau \gg n^2/\omega_r$. An intuitive way to think about this is based on the level diagram in Fig. 2.1. We're trying to coherently transfer amplitude within a 'sea' of neighboring momentum states. A pulse with duration τ will have a width of $1/\tau$ in frequency space - this Fourier width needs to be small relative to the spacing between neighboring momentum states in order to selectively drive higher-order processes without transferring population to any intermediate states. For a mathematical treatment of what I'm referring to, see [57] and references therein.

Fig. 2.2b) illustrates numerical integration/ time evolution of wavefunction amplitudes in the Bragg regime. We can see that intermediate momentum states are barely populated during the integration. This simulation uses $\tau = 5 \text{ ms}$ which satisfies the condition for the Bragg regime for $n = 3$: $\tau \gg n^2/\omega_r = 0.7 \text{ ms}$.

The regime we use in our experiments is the 'quasi-Bragg' regime. It turns out that you can still get near-ideal momentum transfer between two states with pulse durations $\tau \lesssim n^2/\omega_r$. Fig. 2.2a) shows time-evolution of the momentum state amplitudes in the 'quasi-Bragg regime'. There is a very large population in intermediate states during the Bragg pulse, but by the end of the pulse almost no amplitude remains in the intermediate states leaving us with a highly efficient beamsplitter. Details of the quasi-Bragg regime are discussed in [57] for the analytically curious reader. Personally, I find numerical integration to provide a more direct intuition for what works well and what doesn't.

On the other extreme, we have pulses with duration $\tau \ll n^2/\omega_r$ called the 'Raman-Nath'




Figure 2.3: Numerical integration of Bragg diffraction in the Raman-Nath regime. Pulse parameters are $\tau = 3 \mu\text{s}$, and $\Omega_{\text{bragg}} = 15 \text{ E.r.}$. The Bragg resonance condition is set to $n = 3$ but the $k = 3$ state is barely populated after the pulse.

regime. In this regime, you don't have well-resolved energy levels because of the Fourier smearing of the laser frequencies is broader than the separation between states. Fig. 2.3 illustrates numerical integration and wavefunction evolution in this regime. Even though the optical lattice is resonant with an $n = 3$ Bragg resonance condition, we can see that we drive almost no amplitude to the $k = 3$ momentum state. Instead, states neighboring the initial $k = 0$ state are populated. Deep into this regime, the output state population distribution is described by Bessel functions. Note that even though this is a messy process, it is still coherent - people have used this regime for interferometry, even though it's far from ideal for state-of-the-art precision measurements [73].

2.5 Bloch oscillations Hamiltonian

Historically, Bloch oscillations were first explored with electrons in a crystal with an applied uniform electric field. The electrons feel a constant force:

$$H = \frac{\hat{p}^2}{2m} + U_p(x) - Fx \quad (2.23)$$

where U_p is some periodic potential (for example the electric potential in a crystal), and F is the applied force.

In atomic physics experiments, we can accelerate the optical lattice instead of applying a force on the atoms. In the frame of reference of the atoms, there is a fictitious force that gives us the same physics as described by the Hamiltonian in Eq. 2.23. It is convenient to parameterize the sweep in terms of a frequency ramp rate $r = d\delta/dt$, since that is how we implement this acceleration experimentally. Note that to correctly keep track of the phase of the optical lattice as we ramp frequencies, we need to integrate in time: δt is replaced by

$\int \delta(t)dt = rt^2/2$. Plugging this into Eq. 2.51, we get the Hamiltonian in the freely-falling frame H_{FF} :

$$H_{FF} = \frac{\hat{p}^2}{2M_{Cs}} + \hbar\Omega \cos(2k\hat{x} - rt^2/2) + \hbar\omega_{Stark} \quad (2.24)$$

This form of the Hamiltonian ends up being good for numerical simulation of our experiment.

There are, however, other forms of this Hamiltonian that are useful analytically and numerically. This Hamiltonian can be boosted to different frames of reference by using a unitary transformation, which is the quantum equivalent of Galilean frame transformations [48, 65]. This unitary operator takes the form:

$$\hat{U} = e^{\frac{i}{\hbar}d(t)\hat{p}} e^{-\frac{i}{\hbar}mv(t)\hat{x}} e^{\frac{i}{\hbar}\theta(t)} \quad (2.25)$$

where positions are transformed as $\hat{x} \rightarrow \hat{x} + d(t)$, and momenta as $\hat{p} \rightarrow \hat{p} + mv(t)$. The $\theta(t)$ term transforms the overall energy of the Hamiltonian uniformly to all basis states, as a function of time. In general, a variable or operator \hat{O} should be transformed via:

$$\hat{O}' = U\hat{O}U^\dagger + i\hbar \frac{dU}{dt}U^\dagger \quad (2.26)$$

There are a couple of useful transformations of Eq. 2.24 that I'll discuss here. First, we can prove it is equivalent to Eq. 2.23.⁸ The Hamiltonian 2.24 can be transformed with the general unitary operator 2.25 by setting $d(t) = Ft^2/(2m)$, $v(t) = Ft/m$, and $\theta = 2F^2t^3/(3m)$. This results in the following transformed Hamiltonian in the lattice frame H_{latt} :

$$H_{latt} = \frac{\hat{p}^2}{2M_{Cs}} + \hbar\Omega \cos(2k\hat{x}) + \hbar\omega_{Stark} + Fx \quad (2.27)$$

where $+Fx$ potential energy indicates a force in the negative x direction (by convention $F > 0$). This transformation effectively boosted the position and velocity to a frame of reference co-moving with the accelerating lattice. The lattice accelerates in the $+x$ direction, so in the new frame of reference the atoms feel a fictitious force in the $-x$ direction. The $\theta(t)$ term corrects for the kinetic energy difference between the free-falling frame and the accelerating frame, ensuring that the ground state energy stays near zero at all times.

Another useful unitary transformation puts the time dependence of the Hamiltonian into the kinetic energy term, instead of the potential energy term. This is accomplished using the unitary in Eq. 2.25 to transform the Hamiltonian H_{FF} in Eq. 2.24 where $d(t) = Ft^2/(2m)$, $v(t) = 0$, and $\theta(t) = -F^2t^3/(6m)$ leading to the Hamiltonian H' :

$$H' = \frac{(\hat{p} - Ft)^2}{2M_{Cs}} + \hbar\Omega \cos(2k\hat{x}) + \hbar\omega_{Stark} \quad (2.28)$$

In this case, we have boosted position to a co-moving frame with the lattice, but we did not boost the velocity/momentum operator. The quasimomentum is typically identified as $-Ft = +\hbar k_q(t)$, such that we can see the quasimomentum decreases linearly in time due to the negative fictitious force.

⁸I will leave the full derivation to the reader, but note that it is fairly annoying to keep track of commuting variables.

The Hamiltonian in Eq. 2.28 is useful because it allows us to readily plot the energy band structure. In our discretized momentum basis, it is equivalent to:

$$H' = \sum_{l=-\infty}^{l=\infty} \frac{(2l\hbar k + \hbar k_q(t))^2}{2m} |l\rangle\langle l| + \frac{\hbar\Omega}{2} (|l\rangle\langle l+1| + |l\rangle\langle l-1|) \quad (2.29)$$

where $|l\rangle$ is a momentum basis state with momentum $p = 2l\hbar k$. This equation can easily be written in matrix form similar to Eq. 5.40. The energy eigenstates as a function of k_q are shown in Fig. 2.4a) - the band structure⁹ is plotted over the range $k_q = (-3, 3)$ for illustration purposes, though quasimomentum is really only defined modulo k , i.e. in the first Brillouin zone only (indicated by dashed lines).

‘Bloch oscillations’ refers to a particle adiabatically remaining in the ground state of this Hamiltonian while experiencing an external force. As shown in Eq. 2.28, an external force is equivalent to a quasimomentum that is linearly ramping in time. One can imagine an atom with an initial quasimomentum $k_0 = 0$ in the ground state energy band. Eq. 2.28 describes a negative force, so the state starts moving left along the ground state. As the atom crosses the edge of the first Brillouin zone, the atomic state receives a $2\hbar k$ momentum kick from the lattice, shows up on the other side of the first Brillouin zone around $k_q = 1$, and then continues moving left.

Note that I use this ‘ $2\hbar k$ momentum kick’ wording very loosely - the process of momentum transfer is actually ‘smeared out’ in time. For shallow lattices such as the band structure shown in Fig. 2.4c), the momentum transfer is localized to the region around the level crossing where the free-space band-structure is distorted. For very deep lattices e.g. the band structure shown in Fig. 2.4b), the level structure is so highly distorted that you can no longer describe the dynamics as a discrete set of $2\hbar k$ momentum kicks. Instead, the various level crossings overlap one another and momentum is transferred more continuously in time. In this limit, the Wannier-Stark ladder is a more natural description of the dynamics, where atoms are trapped to a single lattice site [87, 28].

Some other things to note about interpreting band structure. First, the group velocity of the wavefunction is defined as $v_g = dE/dk$, so the slope of the band structure tells us how fast the atom is physically moving while in the lattice. In contrast, phase velocity is set by the quasimomentum, which can be seen directly in Eq. 2.7. For deep lattices, the ground state band flattens out which indicates that the group velocity is tending towards zero - atoms are becoming more localized in the deep potential wells at each lattice site. On, the other hand, as the lattice depth goes to zero the velocity smoothly transitions to the match the corresponding free-space velocity: $k_q = 1k$ quasimomentum corresponds to $v = 1v_r$ free-space velocity when the lattice is adiabatically unloaded.¹⁰

⁹This is a one-dimensional band structure - in solid state physics, three dimensional band structures are commonly plotted similar to this, by tracing a trajectory around the first (three-dimensional) Brillouin zone. If you take a solid state physics class, you’ve likely encountered 2D or 3D band structures.

¹⁰Many references discourage thinking of ‘quasimomentum’ as a real momentum, but quasimomentum maps directly to free-space momentum when the lattice is adiabatically unloaded. For people studying physics inside optical potentials that are always on, it might be poor form to think in this way, but since our experiments always adiabatically ramps optical lattice potentials on and then off again, it is very convenient to remember this connection.

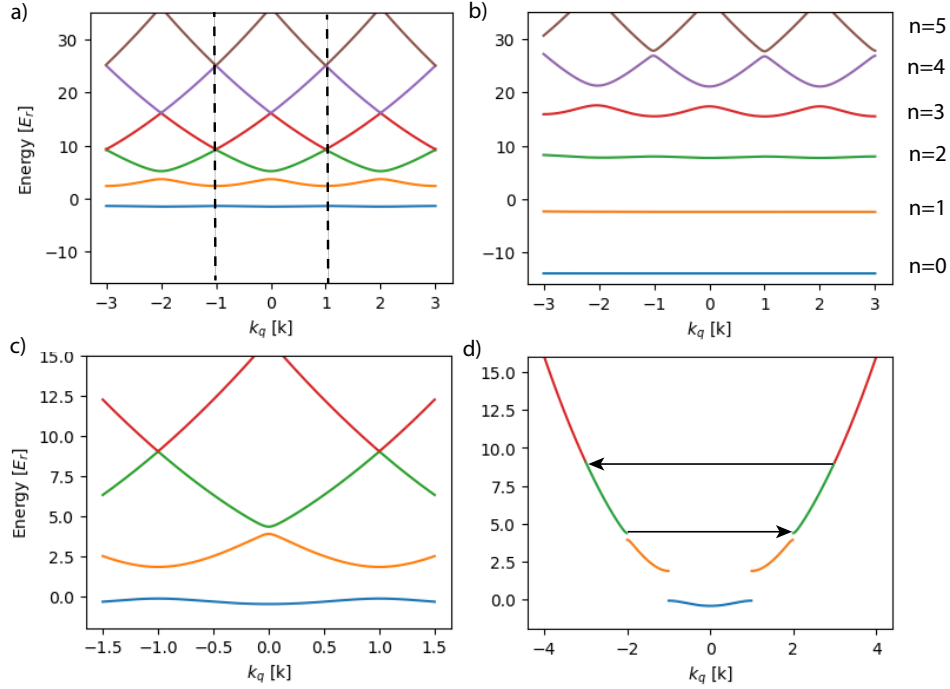


Figure 2.4: Band structure of the Hamiltonian in Eq. 2.28 as a function of the quasimomentum k_q . First two panels show band structure for a lattice depths a) $\Omega = 4E_r$ and b) $\Omega = 20E_r$. The dotted lines in a) indicate the edge of the first Brillouin zone. The labels on the right of b) indicate which line corresponds to which energy level. As the lattice depth gets very deep in b), the band structure approaches harmonic oscillator eigenstates (flat bands with evenly spaced levels) - atoms are localized to the bottom of each lattice site which locally looks harmonic. c) A zoomed-in plot to better illustrate lower energy levels in the first Brillouin zone, for $\Omega = 2E_r$. d) A plot showing the band structure 'unfolded' - this highlights the relationship between band structure and the free-space energies, since the unfolded bandstructure very closely follows the free-space kinetic energy parabola $p^2/2m$. Arrows indicate momentum transfer from Bloch oscillations in the second excited state.

2.6 Bloch numerical simulation

Before getting too much further, let's explicitly introduce the Schrodinger equation corresponding to the Bloch Hamiltonian. The best Hamiltonian to use for numerical simulation is Eq. 2.24 because it allows us to treat frequency ramps generally, including when there are multiple moving lattices. Eq. 2.24 is therefore re-written with the following form:

$$H_{FF} = \frac{\hat{p}^2}{2M_{Cs}} + \frac{\hbar\Omega}{2} \left(e^{i(2k\hat{x}-rt^2/2-\phi_0)} + e^{-i(2k\hat{x}-rt^2/2-\phi_0)} \right) + \hbar\omega_{Stark} \quad (2.30)$$

where ϕ_0 is a phase offset which is equivalent to a position offset $\Delta x = \phi_0/2k$ in the lattice at time $t = 0$.¹¹

Next, the ansatz from Eq. 2.7 is plugged into the Schrodinger equation to get the following set of coupled equations, one for each value of m :

$$i\hbar\dot{g}_m(t) = \frac{\hbar^2(k_q + 2mk)^2}{2M_{Cs}} g_m(t) + \hbar\omega_{Stark} g_m(t) + \frac{\hbar\Omega}{2} \left(e^{-i(rt^2/2+\phi_0)} g_{m-1}(t) + e^{i(rt^2/2+\phi_0)} g_{m+1}(t) \right) \quad (2.31)$$

or, in dimensionless units (see Section 2.8):

$$i\dot{g}_m = (2m + \tilde{k}_q)^2 g_m + \tilde{\omega}_{Stark} g_m + \frac{\tilde{\Omega}}{2} \left(e^{-i(r\tau^2/2+\phi_0)} g_{m-1} + e^{i(r\tau^2/2+\phi_0)} g_{m+1} \right) \quad (2.32)$$

which is the form used for numerical simulation. See Section 2.8 for how to derive dimensionless equations.

Two superposed Bloch lattices in a 'multi-frequency' pulse can be described by superposed optical lattices that are beating together, exactly analogous to the multi-frequency Bragg diffraction Hamiltonian in Eq. 2.21:

$$\cos(2k\hat{x} - rt^2/2 - \phi_0) + \cos(2k\hat{x} + rt^2/2 - \phi_0) = 2 \cos(2kx - \phi_0) \cos(rt^2/2) \quad (2.33)$$

In the multi-frequency case, the Schrodinger equation then becomes

$$i\hbar\dot{g}_m(t) = \frac{\hbar^2(k_q + 2mk)^2}{2M_{Cs}} g_m(t) + \hbar\omega_{Stark} g_m(t) + \frac{\hbar\Omega}{2} \cos(rt^2/2) \left(e^{-i\phi_0} g_{m-1}(t) + e^{i\phi_0} g_{m+1}(t) \right) \quad (2.34)$$

¹¹This ϕ_0 parameter is critical to understanding the symmetric Bloch oscillations results in Chapter 3.

and the dimensionless form used for numerical simulation:

$$i\dot{g}_m = (2m + \tilde{k}_q)^2 g_m + \tilde{\omega}_{Stark} g_m + \frac{\tilde{\Omega}}{2} \cos(\tilde{r}\tau^2/2) \left(e^{-i(r\tau^2/2+\phi_0)} g_{m-1} + e^{i(r\tau^2/2+\phi_0)} g_{m+1} \right) \quad (2.35)$$

2.6.1 Landau-Zener Tunneling

Continuing the discussion from section 2.5, let's look more closely at the avoided level crossings seen in the band structure in Fig. 2.4. Avoided level crossings are ubiquitous in quantum physics. We are approaching the 100-year anniversary of the Landau-Zener solutions to tunneling from an avoided level crossing [103].¹² The canonical picture of an avoided level crossing used to derive the Landau-Zener tunneling formula is two linearly sloped energy levels that cross one another as a function of some parameter in the Hamiltonian. An example of a typical avoided level crossing is between states $n = 1$ and $n = 2$ in Fig. 2.4c) near $k_q = 0$. For weak lattices¹³, tunneling through the level crossings is well-described by the Landau-Zener formula, since locally around the level crossings the energy levels are nearly linear. For deep lattices¹⁴, the formula breaks down because the ground state band structure is so dramatically deformed. The level crossing no longer can be approximated by linearly sloped band structure, and moreover the level-crossings start overlapping in time.

The Landau-Zener tunneling result mapped to Bloch oscillations describes leakage from the ground state of the lattice per Bloch oscillation, as a function of the lattice depth and acceleration rate. The survival probability of an atom in the ground state after a single Bloch oscillation is given by:

$$P_{surv} = 1 - e^{-2\pi\Gamma} \quad (2.36)$$

with $\Gamma = \Omega^2/(8ka)$ for an acceleration a . Putting this in the same form as in reference [15]:

$$P_{surv} = 1 - e^{-\frac{F_c}{F}} \quad (2.37)$$

where the critical force $F_c = \frac{\pi k \hbar^2 \Omega^2}{8E_r} = \frac{\pi k V_0^2}{32E_r}$ when using their nomenclature for $V_0 = 2\Omega$.¹⁵ Putting Eq. 2.36 in dimensionless form:

$$P_{surv} = 1 - e^{-\frac{2\pi\Omega^2}{4r}} \quad (2.38)$$

where r is the dimensionless ramp rate as in Eq. 2.30.

Fig. 2.5 compares this theoretically predicted Landau-Zener equation to direct numerical integration of the Schrodinger equation. We can see that for (dimensionless) ramp rates greater than ≈ 3 the Landau-Zener equation starts to fail. Experimentally, we use

¹²The results were actually published in the same year by four separate people. I cited the paper that was written in English.

¹³i.e. lattice depths $\Omega \ll 8E_r$. See e.g. [15] for similar statements. Multiple factors of two combine to set $8E_r$ as the relevant energy scale for shallow/deep lattices.

¹⁴ $\Omega \gtrsim 8E_r$

¹⁵We agree on factors of 2! Hallelujah this took a while to get correct.




Figure 2.5: Figure comparing simulated Bloch oscillation efficiency (blue) compared to the Landau-Zener tunneling equation (orange).

$rr = 8.58027\omega_r^2$ which corresponds to exactly 10g acceleration - experimentally we therefore operate in a regime where the Landau-Zener formula is not a good prediction of Bloch oscillation efficiency.¹⁶

2.6.2 Role of excited states during Bloch oscillations

Typically Bloch oscillations are understood as an atom adiabatically following the ground state of the optical lattice. However, it is also possible to populate and even adiabatically follow the excited states. Fig. 2.4d) helps illustrate this concept. Suppose an atom has a free-space momentum $2.5k$, and then is adiabatically loaded into an optical lattice (at rest). In Fig. 2.4, this state would lie in the middle of the right branch of the green curve - the atom has been loaded into the second excited state. Now, suppose that we accelerate the lattice in the negative direction (or equivalently apply a positive force to the atom) - the state will increase in quasimomentum until reaching the avoided level crossing between the second and third excited states (the red and green curves). If the acceleration is slow enough/ if the lattice depth is large enough, the atom will adiabatically stay in the second excited state. It will receive a $-6\hbar k$ momentum kick, indicated by the arrow, such that the state now lies on the left branch of the green curve. It will continue increasing quasimomentum until reaching the next avoided crossing, where it will receive a $+4\hbar k$ momentum kick, indicated by the other arrow. This completes one "Bloch oscillation" in an excited Bloch band. The total momentum transfer is still always $2\hbar k$ per Bloch period, but the dynamics of the atom will look much different in the higher state compared to the ground state. Bloch oscillations in excited states of the lattice have been proposed as a way to decrease experimental sensitivity to different parameters in an atom interferometer [53]. Note also that a single one of these

¹⁶But it's not too far off for our experiment and still gives a good rule-of-thumb for general scalings.

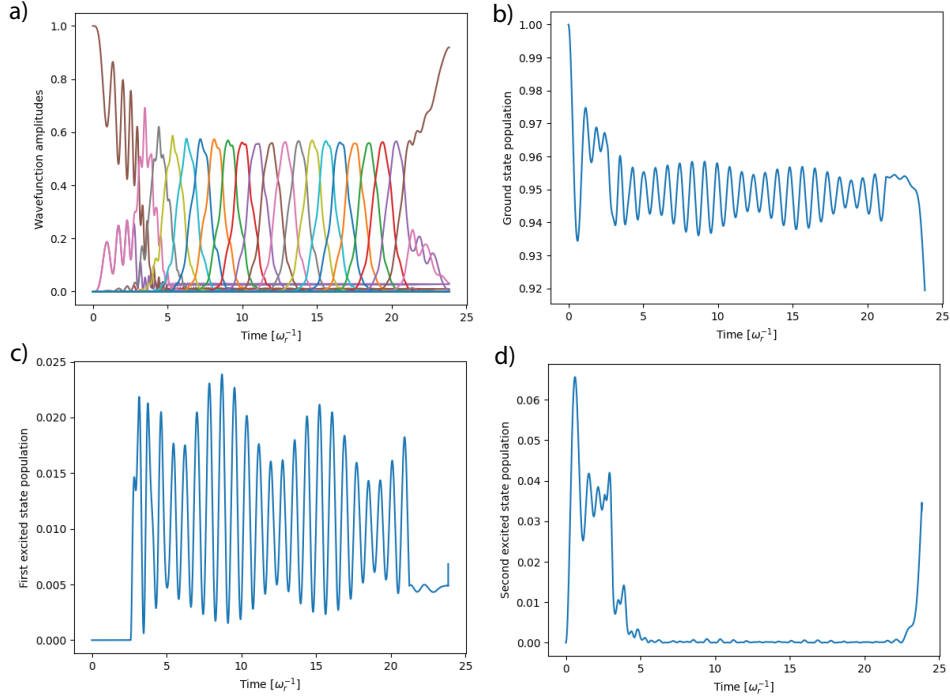


Figure 2.6: Wavefunction evolution over time of a lattice load, 20 Bloch oscillations, and then a lattice unload. All figures are projections of the same evolution with a lattice depth of $10E_r$. a) Amplitudes on the basis states show momentum being transferred each Bloch oscillation. The next figures show the projection squared (to get population) of this state onto the b) ground state, c) first excited state, and d) second excited state of the lattice.

higher-order transitions is equivalent to adiabatic rapid passage (ARP) Bragg diffraction, as explored in the reference [47].

So there's some cool physics possible with higher-order Bloch bands, but why does this matter when we keep atoms in the ground state during our experiment? Although we try to adiabatically load atoms into the ground state by ramping up intensity over several hundred microseconds, this process isn't perfectly adiabatic. Moreover, any phase slips within the DDS output that is generating the modulation frequency can cause discontinuous jumps in the optical lattice. So even though our atom cloud initially has a velocity spread on-axis of much less than one recoil velocity (well within the first Brillouin zone), non-adiabaticity still can excite some population into the excited states.

Figs. 2.6 and 2.7 illustrate an example of what the Schrodinger equation dynamics can look like for a lattice depth of $10E_r$ and $20E_r$ respectively. The plots show the wavefunction population in the ground state and first few excited states as a function of time during the simulation. Parameters used are identical to experimental parameters: the lattice is linearly turned on/off in $200 \mu s$, which is about 2.5 in the plotted dimensionless units. $N=20$ Bloch oscillations occur in about 20 dimensionless time units with an acceleration $a = 10g$. See Section 2.8 for further discussion of the dimensionless unit system.

For $10E_r$ lattice depth, we can see that there is some non-adiabatic loading which transfers population to the second excited state. Then, as the lattice accelerates, population in all

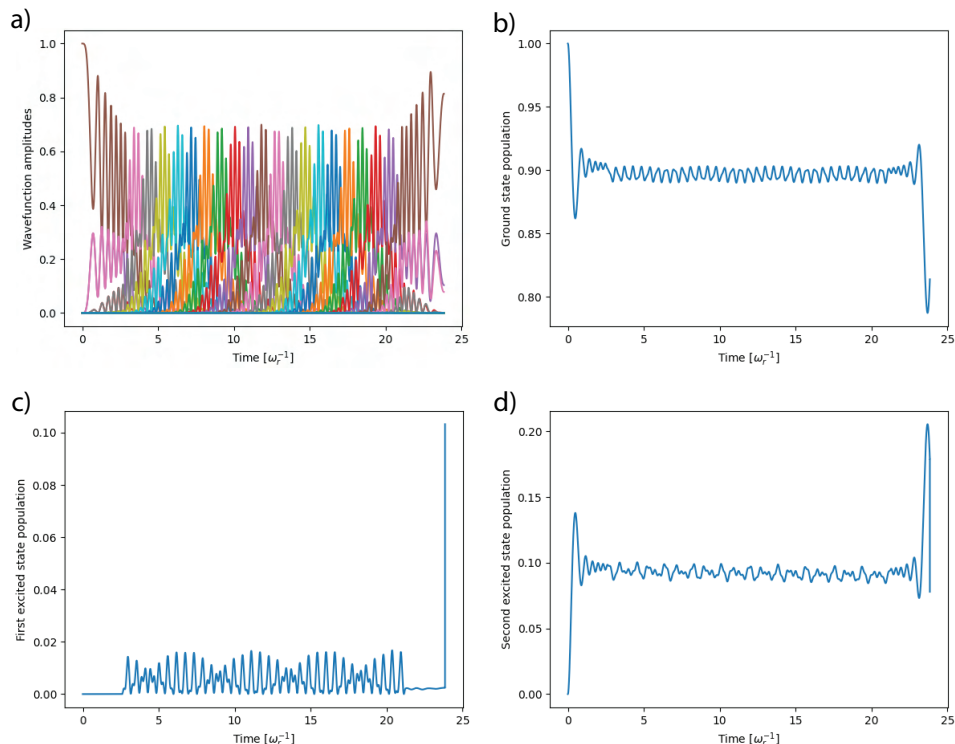


Figure 2.7: The same information as Fig. 2.6, but with a lattice depth of $20E_r$, to illustrate population becoming trapped in an excited state. See text for further discussion.

of the excited states quickly decays away to higher and higher excited states, not shown here - it's the equivalent of 'being left behind' by the lattice as it accelerates away. This effect can also be seen in the raw evolution of the basis states: the fast oscillations in the amplitudes at early times come from population in the ground state and excited states coherently interfering with one another.¹⁷ After the excited state populations have decayed away/been left behind, the evolution appears much 'cleaner' because this interference is no longer happening.

Fig. 2.7 shows the same simulation but with a much deeper optical lattice. Now, more population is transferred to the second excited state during lattice load during the lattice load, and the lattice depth is deep enough that this population doesn't tunnel away but instead remains trapped in the second excited state. The interference effects are evident in Fig. 2.7a) over all 20 Bloch oscillations.

Projecting the simulated state evolution from Eq. 2.32 onto the Bloch bands shown in Fig. 2.4 is not very straightforward because the two are calculated in different frames of reference. The simulation is done in a freely-falling frame where the lattice accelerates away, whereas the Bloch bands are calculated in a frame co-moving with the lattice as it accelerates. Therefore, after calculating the eigenstates of the Hamiltonian 2.28 for each time step, the eigenstates then need an additional unitary transformation to boost them to

¹⁷The ground state and excited states are orthogonal basis states, so why would we see them beating together? Because Fig. 2.6a) shows the free-space momentum states, not the projections onto the Bloch eigenstates. Since we're not in the correct basis, we see interfering amplitudes on the free-space basis states.

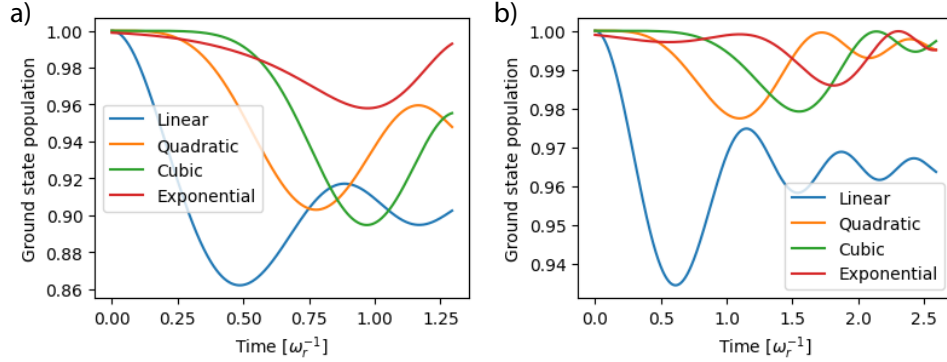


Figure 2.8: Numerical simulations of loading an optical lattice to find the optimal intensity ramp profile in a) $100\mu s$ and b) $200\mu s$. Exponential ramping appears to be the winner.

the frame-of-reference used in the simulations. Only then can you take the projection of the simulated wavefunction onto the calculated eigenstates. There's an example of this in the Monte Carlo code.

Also, it is possible to minimize the non-adiabaticity of loading the lattice for a given duration of ramp time. See e.g. [20] for similar discussion, though countless AMO experiments have surely looked at this at some point. The old alpha experiment ramped on the lattice intensity linearly in $100\mu s$ to a lattice depth of around $10E_r$. Figure 2.8a) shows the projection into the lattice ground state for different ramping profiles with the old alpha parameters. Figure 2.8b) shows same simulation but with $200\mu s$ ramp time, which is what is currently being used in the new experiment. The population in the ground state at the end of the intensity ramp (right side of the plot) is the takeaway from these plots, although note that the oscillations are from non-adiabatic interfering amplitudes and complicate the interpretation somewhat. Even so, it's clear that an exponential ramping profile is optimal. The old experiment was likely losing around 10% population during loading and also during unloading the lattice. The linear, quadratic, and cubic waveforms are straightforward to implement. The exponential waveform used for the simulation is $\exp(4(t - t_{\text{load}})/t_{\text{load}})$, so that the lattice is ramped (fractionally) from e^{-4} up to 1 in the time t_{load} .

There is one other source of non-adiabaticity in our experiment that leads to very similar effects. Experimentally, the modulation frequency¹⁸ is ramped from an initial value to a final value during the Bloch oscillation acceleration. During loading and unloading, this frequency is constant/not ramped. If the direct digital synthesizer (DDS) controlling this ramp has some phase slip in the output waveform between loading the lattice and beginning the Bloch oscillation, there is a non-adiabatic jump in the beam frequency between the two lattices¹⁹. Experimentally this phase jump clearly causes non-adiabatic losses. A different DDS or better configuration of the current DDS would fix this issue in the long-term, but for now the phase needs to be carefully adjusted when setting new Bloch parameters.

¹⁸The frequency that is controlling the velocity splitting between the two lattices

¹⁹The $\cos(\omega_m t)$ part of the coupling term in the Hamiltonian.

2.6.3 Systematic effects from Bloch oscillation excited states

With this new understanding, are there potential systematic effects associated with the higher Bloch bands? Fortunately, our experiment symmetrically accelerates both interferometers at the same time, so the main effects are common mode to both interferometers. We would therefore only be sensitive to asymmetries between the two interferometers. Experimentally, we have seen some evidence of what might be population in excited states - when we look at the time of flight traces, there appears to be an initial rapidly decaying population in the first 10 Bloch oscillations, after which there is a much slower decay in population. This is exactly the behavior to expect from an initial population in an excited state when the lattice starts ramping, similar to the dynamics shown in Fig. 2.6. Since we know there is a phase jump on our DDS when beginning the lattice acceleration, it would not at all be surprising if this is happening. The same non-adiabatic losses will also happen when unloading the lattice, though the effect is harder to see because the population in the excited states will stay near the target state instead of tunneling away. Though I don't think these shifts will ultimately lead to a systematic, there is more work that needs to be done here bounding the effect for the final configuration in the alpha measurement.

2.7 Phase evolution during Bloch oscillations

Another important application of the above theory is to demonstrate an understanding of a wavefunction's phase evolution during Bloch oscillations. In a freely falling frame, a free-space wavefunction's evolution is simply governed by the kinetic energy of the atom. Inside an optical lattice, however, we have now seen that the wavefunction dynamics are much more complicated, so the phase evolution of the wavefunction could be similarly complicated. Let's simulate Bloch oscillations by numerically integrating Eq. 2.32 and work to explain the entire phase evolution of the wavefunction.

The most natural way to analyze the problem is to project the atomic state onto the ground state of the optical lattice, since we're adiabatically following the ground state. This projection is complex - the magnitude squared tells you the population in the ground state, and the phase of the projection tells you the phase of the ground state wavefunction. Once you do this, you find that the phase evolution of the wavefunction comes mostly from kinetic energy and potential energy (as expected), as well as a small correction from non-adiabatic effects.

Fig. 2.9 shows the numerical results of this projection as a function of time.²⁰ Fig. 2.9 shows the full evolution of the wavefunction over five Bloch oscillations, which was found by numerically integrating Eq. 2.32. a) Shows the wavefunction amplitudes in each basis state for reference. b) Shows the phase of the simulated wavefunction projected onto the ground state of the lattice, with no assumptions. c) shows the same projection but where the phase shift from kinetic energy is removed.²¹ d) removes light shift energy in addition to the kinetic

²⁰More details about how to obtain this projection were discussed in Section 2.6.2, it's not straightforward because the simulation frame of reference is different from the frame of reference used to calculate the band structure

²¹Since kinetic energy is diagonal in the momentum basis, you do this by multiplying each momentum state by the integrated kinetic energy phase for that state (with a minus sign so you remove it).

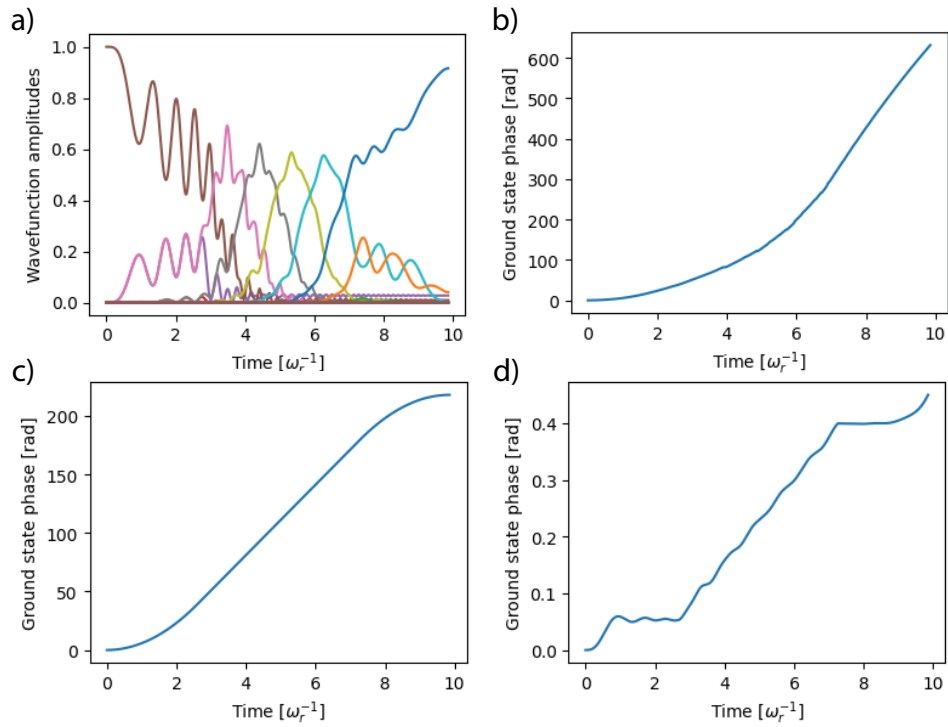


Figure 2.9: Example bookkeeping of the phase evolution during a simulation of Bloch oscillations. a) The wavefunction amplitudes for reference, b) the phase evolution including phase from kinetic and light shift energies, c) phase evolution with kinetic energy phase removed, d) phase evolution with light shift phase removed as well, leaving only the non-adiabatic phases.

energy. We can see that kinetic and potential energy explain most of the phase evolution. The remaining phase is from non-adiabatic effects during lattice loading, acceleration, and unloading. For a good reference on how to calculate these non-adiabatic phase shifts, see section V in Tim Kovachy’s paper [48]. For the example shown in Fig. 2.9, I’ve found that the first-order term in this paper doesn’t accurately replicate the phase seen numerically, and this is likely because the deep lattice/fast ramp rate in our experiment requires one to include more terms in the non-adiabatic calculation. If you are determined to get this calculation correct and the Bloch band picture requires too many higher-order terms, you might find more luck using a Wannier-Stark picture to analytically calculate phase shifts [28].

2.7.1 Potential systematics from Bloch oscillations

So there’s more information than you ever wanted to know about Bloch oscillations. If you read through this entire section, you’re probably burned out trying to get the experiment to not fall apart and you really just want to know ”Zack, does it shift our measurement or not?” As it turns out, the systematic shifts from Bloch oscillations very nearly cancel our in our experiment - this is the reason that the previous α measurement didn’t have to do a very detailed analysis of systematic effects from Bloch oscillations.

Within one of the interferometers in the SCI geometry, both arms are accelerated in the same way by the same optical lattice/ same frequency components of our laser. Therefore, any systematic effects shifting the phase on one arm of the interferometer will (very nearly) phase shift the other arm as well. The systematic effects you need to worry about are from very slight intensity differences between the arms of a single interferometer. For $n = 5$ Bragg order and $T = 100$ ms, the arms within a single interferometer are separated by 3.5 mm so we’re talking about changes in the laser propagation on a very short length scale.

Moreover, the differential measurement of the SCI between the upper and lower interferometers can cancel much of any systematic shift that remains. Refer to the SCI geometry as illustrated in Fig. 1.3 in the first chapter. For example, suppose there was a gradient in lattice depth vertically in the chamber due to e.g. the Rayleigh range of the beam. In this case, the upper arm of each interferometer will see a higher/lower lattice depth compared to the lower arms. Systematic effects from this type of gradient in a potential are cancelled by the SCI geometry - they have the same symmetry as the gravitational acceleration or a uniform B-field, which are uniform gradients in some type of potential energy.

The systematic effects that can remain from Bloch oscillations are due to the difference in the light shift gradient between the upper and lower interferometer at the time of the Bloch oscillations. The upper interferometer is only offset from the lower interferometer by $\approx 175 \mu\text{m}$ at the time of Bloch oscillations due to the ≈ 5 ms propagation time between the second Bragg diffraction pulse and the start of Bloch oscillations. The only Bloch oscillation phase shifts that could then remain are related to higher order gradients in the laser intensity on the length scale of $175 \mu\text{m}$, which will be extremely small.

Let’s walk through an example of the old alpha experiment assuming a Gaussian beam with waist 3.24 mm beam waist.²² Section 2.7 just showed that almost all of the phase from

²²Chenghui’s thesis says their beam waist was 2.29 mm, but I’m almost positive he meant $\sqrt{2} * 2.29$ mm as

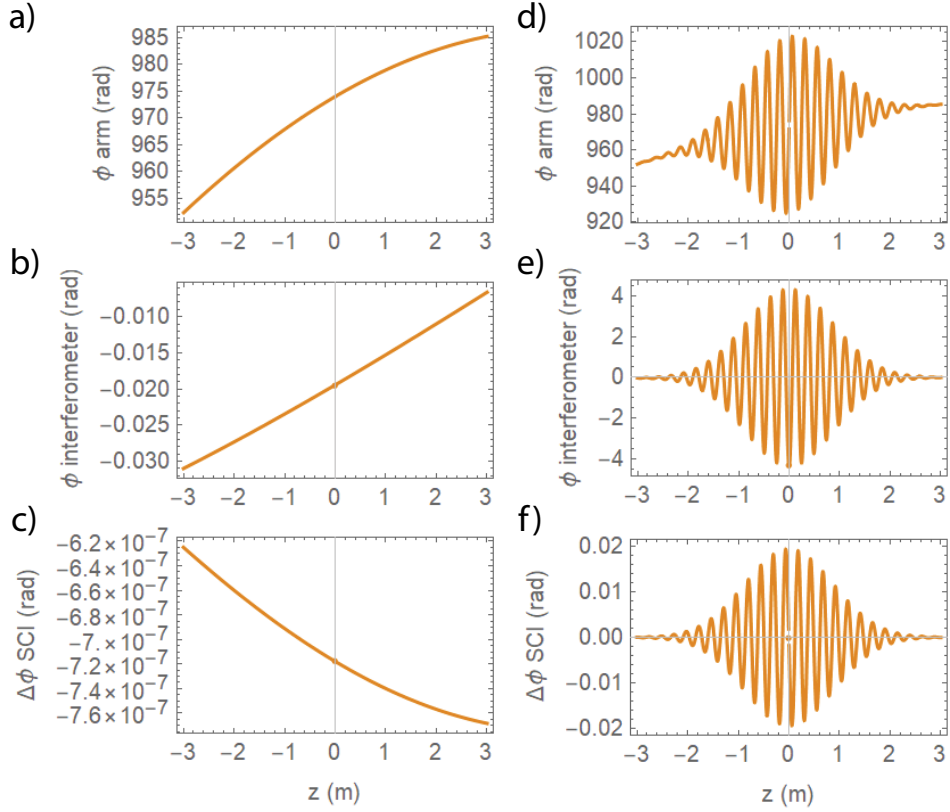


Figure 2.10: An estimate of phase shifts from the integrated light shift during Bloch oscillations as a function of the height in the chamber where Bloch oscillations occur. a-c) show an estimate of the Gaussian beam profile used in the old alpha experiment, d-f) shows the same estimate but with a highly perturbed beam. a) and d) show the total light shift phase on a single interferometer arm. b) and e) show the phase difference of the upper arm minus the lower arm for a single interferometer. c) and f) show the total SCI phase shift from the upper interferometer minus the lower interferometer. See text for more details.

Bloch oscillations (aside from kinetic energy) comes from light shift energy - we can calculate typical systematic effects from gradients in the light shift energy of a beam.

Fig. 2.10a-c) shows the light shift for the configuration of the old alpha experiment to give a rough estimate of the size of the systematic shift. These results are very insensitive to the location of the beam focus because retro-reflection reverses the direction of divergence of the beam. d-f) show the same configuration but with a dramatically perturbed beam. The beam here is assumed to have a sinusoidal perturbation of $\pm 5\%$ of the total amplitude and a spatial wavelength of the perturbation of $q \approx 300 \mu\text{m}$ - it's chosen as an example of a complicated propagation pattern that has substantial intensity gradients along the propagation axis.

the $1/e^2$ radius. The corrected number reproduces a much more accurate Gouy shift as given in their error table, and also $2.29 \text{ mm } 1/e^2$ radius would just be way too small for their $2 \text{ mm } 1\sigma$ atom cloud - if this were the case, the 1σ point of the atom cloud distribution would see 50% laser intensity relative to the maximum of the beam. Since Bragg diffraction is non-linearly sensitive to lattice depth, they would have had much worse contrast experimentally than we see in their old data.

For details about the propagation of this beam, see Chapter 5. The calculation assumes a lattice depth of $20E_r$, a Bloch acceleration of $10g$ where $g \approx 9.8 \text{ m/s}^2$ is the gravitational acceleration on earth's surface, and $N = 200$ Bloch oscillations. The lattice depth and number of Bloch oscillations are slightly exaggerated relative to the old experiment. These parameters give roughly 1000 rad phase from light shift in total. Fig. 2.10c) shows that the phase shift in the old experiment would have been extremely small and not a systematic effect in the measurement.²³ Fig. 2.10f) shows that even a highly perturbed beam might lead to $\approx 10 \text{ mrad}$ shifts in a single location of the beam, but when averaged over the atomic ensemble this effect will average to zero and not contribute to a systematic effect.

The conclusion from this section is that we don't need to worry about systematic effects inherent to the basic Schrodinger equation dynamics during Bloch oscillations, particularly from the light shift. However, there **are** still systematic effects due to Bloch oscillations - primarily from intensity inhomogeneities in the beam causing deviations δk in the local k-vector of the laser beam. This leads to a change in the kinetic energy of the atom after receiving $\hbar(k+\delta k)$ momentum kicks from the laser beam. Chapter 5 explores these systematic effects related to the spatial intensity profile of the laser beam and deviations in the k-vector of the laser.

2.8 Dimensionless units

This section outlines the dimensionless unit system used for deriving the dimensionless Schrodinger equations presented in this chapter, and it is also the units system used for all of the Monte Carlo code. The most useful dimensionless unit system for our Hamiltonians, also called 'recoil units', is one where $\hbar = M_{Cs} = k = 1$, k being the wavevector of the laser and M_{Cs} being the mass of Cesium 133.

In the Monte Carlo discussed in Chapter 5, unit conversion is controlled automatically by a library for adding and removing units so that you don't need to think about numerical prefactors when converting quantities. Once the library is set up correctly, it dramatically reduces possible conversion errors as long as you're using it correctly. If you're spending a lot of time with the Monte Carlo code, I would still recommend you understand how to use these relationships in order to understand the unit engine in the code/ how to bring simulated quantities back to experimental units correctly.

Table 2.8 summarizes many of the core relationships in this dimensionless unit system. The left column defines a dimensionless variable, typically indicated with a tilde over the normal variable. The middle column shows an equivalent dimensionless expression where the standard variable with dimensions is multiplied by a conversion to remove units from the expression. The final column again shows an equivalent expression but where I have set $\hbar = k = M_{Cs} = 1$ from the previous column. This last column is useful when converting equations with dimensions to equations without dimensions - the procedure for converting an equation is to set $\hbar = k = M_{Cs} = 1$ throughout the equation, then substitute the equivalent of any variables with dimensions. For example, if the acceleration a is in the equation, substitute $a = \tilde{a}/4$. When this is done for all variables in an equation in a consistent way, you end up with a valid dimensionless equation with dimensionless variables.

²³The error bars in their final measurement were on the order of 1 mrad, in units of total phase

Dimensionless variable	Dimensional equivalent	Numerical equivalent
Position, velocity, and acceleration relationships		
\tilde{x}	kx	x
$\tilde{v} = \frac{d\tilde{x}}{d\tau}$	$\frac{k}{\omega_r}v$	$2v$
$\tilde{a} = \frac{d^2\tilde{x}}{d\tau^2}$	$\frac{k^2}{\omega_r^2}v$	$4a$
$\tilde{\lambda}$	$\lambda k = \frac{2\pi}{k} * k$	2π
Time, frequency and energy relationships		
$\tilde{\omega}_r$	$\omega_r \times \frac{M_{Cs}}{\hbar k^2} = \frac{\hbar k^2}{2M_{Cs}} \times \frac{M_{Cs}}{\hbar k^2}$	$\frac{1}{2}$
$\tilde{E}_r = \tilde{\omega}_r$	$E_r \times \frac{M_{Cs}}{\hbar^2 k^2} = \frac{\hbar^2 k^2}{2M_{Cs}} \times \frac{M_{Cs}}{\hbar^2 k^2}$	$\frac{1}{2}$
τ	$\omega_r t$	$\frac{t}{2}$
$\tilde{\omega}$	$\omega_r^{-1}\omega$	2ω
$\tilde{r} = \frac{d(\Delta\tilde{\omega})}{d\tau}$	$\omega_r^{-2}r$	$4r$
\tilde{F}	$F \frac{k}{\omega_r^2 M_{Cs}}$	$4F$
$\tilde{\Omega}$	$\frac{\Omega}{E_r}$	2Ω

Table 2.1: A table of relationships between dimensionless and dimension-full variables in a recoil unit system.

I'll also document some other useful relationships between some of these variables. An optical lattice is described by the following cosine argument

$$\cos(2kx - \Delta\omega t) = \cos(2\tilde{x} - \Delta\tilde{\omega}\tau) \quad (2.39)$$

from which we can see the lattice velocity is given by

$$v = \frac{\Delta\omega}{2k} \quad (2.40)$$

or in dimensionless units,

$$\tilde{v} = \frac{\Delta\tilde{\omega}}{2} \quad (2.41)$$

Converting units is very easy in an already dimensionless expression such as within a cosine or exponential argument because the units already cancel out, so there are no extra factors of 2. Also, note that the optical lattice has a spatial period of $\lambda/2 \Rightarrow \tilde{\lambda}/2 = \pi$.

The frequency ramp rate r is related to the force F and acceleration a that the atoms see by

$$r = \frac{d(\Delta\omega)}{dt} = 2k \frac{dv}{dt} = 2ka = \frac{2kF}{M_{Cs}} \quad (2.42)$$

or equivalently

$$F = \frac{rM_{Cs}}{2k} \quad (2.43)$$

or in dimensionless units,

$$\tilde{F} = \frac{\tilde{r}}{2} \quad (2.44)$$

A useful experimental quantity is the frequency ramp rate that corresponds to gravitation acceleration of the optical lattice. This is relevant experimentally because this is how fast we need to ramp our frequencies to compensate for gravity.²⁴

$$r = 2ka = 2(2\pi/(852\text{nm})) (9.8\text{m/s}^2) = 1.45 * 10^8\text{s}^{-2} \quad (2.45)$$

This equation gives us the angular frequency ramp rate. Experimentally, we use units of Hz to program our DDS frequency synthesizers. If you divide by 2π , you get 23 kHz/ms, which you should have memorized if you are working on any Cesium atom interferometer that experiences gravitational acceleration. The mobile gravimeter project in our group measured this quantity at the 10^{-9} level [96]. In our experiment, we need to know this ramp rate much better than the 10^{-5} level.²⁵

²⁴Experimentally, we actually ramp the drive of a double-passed AOM so the frequency ramp rate that we generate is half this rate.

²⁵We need errors in the optical lattice velocity relative to the freely falling atoms to be much less than a recoil velocity during the interferometer time. The Bragg velocity bandwidth is on the order of $0.2v_r$, and even smaller velocity shifts cause diffraction phases. One recoil velocity/recoil frequency in one second is $3.5k\text{Hz/s}$ which is $\approx 10^{-5}$ shift on the $23k\text{Hz/ms}$ ramp rate.

The Bloch period²⁶ is given by [45]:

$$T_b = \frac{2\pi\hbar}{Fd} = \frac{4\hbar k^2}{rM_{Cs}} \quad (2.46)$$

where $d = \lambda/2$ is the optical lattice period. In dimensionless units,

$$\tilde{T}_b = \frac{4}{\omega_r \tilde{r}} = \frac{8}{\tilde{r}} \quad (2.47)$$

A useful number to remember experimentally is that the Bloch period due to gravitational acceleration in an 852 nm lattice is around 700 μs . We ramp at 10g in our experiment so the Bloch period is about 70 μs .

Note that in the above unitless system, $t = \omega_r \tau = \tau/2$, or in words, one unit of dimensionless time is an inverse recoil frequency. I could have alternatively defined $t = 2\omega_r \tau = \tau$. Either way, you'll end up with factors of two in your equations no matter what you choose. I chose to define things this way so that we can think of dimensionless time as x number of recoil times, or dimensionless frequency as x number of recoil frequencies, or dimensionless energy/ lattice depth as x number of recoil energies. The alternate definition would have made converting equations easier, but interpreting the variables harder. I think this way we have to carefully derive equations, but then after this is done correctly we will make far fewer factor-of-2 mistakes converting between experimental and dimensionless quantities.

2.9 Laser power to lattice depth

For lack of a better place to put this section, it's going here. Have you ever wondered what ' Ω ' actually is, and how it relates to the optical power you're sending at the atoms in lab? Here's a conversion for our experiment that was checked experimentally. Note that all frequencies here are assumed angular frequencies.

The light shift²⁷ here will be called the optical dipole potential U_{dip} . The lattice depth Ω , relevant once you have a retroreflected beam, is given by (in terms of single photon Rabi frequency Ω_0 and detuning δ):

$$\Omega = \frac{\Omega_0^2}{2\delta} = \frac{\Gamma^2}{4\delta} \frac{I}{I_{\text{sat}}} \quad (2.48)$$

where the second equality follows from Eq. 49 in Steck D Line Data [77]:

$$\frac{I}{I_{\text{sat}}} = 2 \left(\frac{\Omega}{\Delta} \right)^2 \quad (2.49)$$

After much looking into this equation, it turns out that we need to amend this slightly for our experimental situation on the α experiment - the above applies to beams far-detuned compared to the splitting between the D1 and D2 lines. When we're around 10 GHz from

²⁶The time needed for an atom to traverse the first Brillouin zone and transfer $2\hbar k$ recoil momentum. This will be explained in detail in subsequent sections

²⁷Formerly known as AC Stark shift, but Stark was a Nazi so I try to use the term light shift instead. The two terms are interchangeable as far as the physics in this thesis is concerned

the D2 line but ≈ 15 THz from the D1 line, the D1 line contribution can be ignored and the contribution from the D2 line is only 2/3 what you would expect from the above equations due to the relative line strengths. See Eq. 19 in [36] for a quantitative treatment of this. Ignoring the D1 line, the lattice depth becomes:

$$\Omega = \frac{\Omega_0^2}{2\delta} = \frac{\Gamma^2}{6\delta} \frac{I}{I_{sat}} \quad (2.50)$$

This definition of lattice depth/Rabi frequency Ω corresponds to a Hamiltonian:

$$H = \frac{\hat{p}^2}{2M_{Cs}} + \hbar\Omega \cos(2kx - \Delta\omega t) + \hbar\omega_{Stark} \quad (2.51)$$

In my opinion, saturation intensity is a messy parameter to use in calculations for our experiment. Generally, it depends on laser detuning as shown in Eq. 2.49, however in references the resonant saturation intensity is always quoted [77]. It also depends on polarization, but that is not immediately obvious from Eq. 2.49. It is a useful quantity for order-of-magnitude lab estimates, and was probably very useful for thermal vapor spectroscopy back in the day, but it's a bit outdated for our coherent operations.

Another way of approaching this calculation in addition to Eq. 2.48 above, follows reference [36]: the light shift for a single traveling wave beam is given by (see Eq. 19 in [36]):

$$U_0 = \frac{\pi c^2}{\omega^3} \frac{\Gamma}{\Delta} I(r) \quad (2.52)$$

Note here that this assumes that we are very very far detuned from the D1 line so we only need to use the term from the D2 line.

If you have two counter propagating beams, you get the above light shift from each beam, and then two two beams interfere such that the shift oscillates between zero and twice the two light shifts, i.e. the total amplitude of the lattice depth is four times the traveling wave light shift (for equal intensities in each beam - otherwise it will differ from four). Since we define the lattice depth as the amplitude on a cosine wave, then we have $\Omega_{lattice} = 2\Omega_{traveling}$. When defined this way, the intensity in Eq. 24 refers to the intensity in one of the beams, not the total intensity.

This agrees with the text near equation 42 in [36]. Also, this numerically agrees with the above Eq. 2.50 if you use $I_{sat} = 1.1mW/cm^2$, which is the saturation intensity for σ^\pm polarized light, as we use in our experiment. This agrees with the text below Eq. 52 in [77], that the saturation intensity is typically quoted as the smallest value saturation intensity which corresponds to circularly polarized light.

Now we can write an equation for the total optical power in a Gaussian beam giving a certain lattice depth. A Gaussian beam has a peak intensity of twice the power over the mode area πw_0^2 , such that the intensity distribution is given by

$$I(\vec{r}) = \frac{2P}{\pi w_0^2} e^{-2\frac{|\vec{r}|^2}{w_0^2}} \quad (2.53)$$

Putting these together, we can relate the total power to the lattice depth:²⁸

²⁸Note that when written this way, the power P is the power in a single frequency component. A single

$$U_0 = \frac{2\pi c^2}{\omega^3} \frac{\Gamma}{\Delta} \frac{2P}{\pi w_0^2} e^{-2\frac{|\vec{r}|^2}{w_0^2}} \quad (2.54)$$

or, solving for the optical power,

$$P = U_0 \frac{\omega^3}{2c^2} \frac{\Delta}{\Gamma} \frac{w_0^2}{2} e^{2\frac{|\vec{r}|^2}{w_0^2}} \quad (2.55)$$

If the powers in each beam are not equal, then this relation is slightly modified. The lattice depth to intensity equation becomes

$$U_0 = \frac{2\pi c^2}{\omega^3} \frac{\Gamma}{\Delta} \sqrt{I_1(r)I_2(r)} \quad (2.56)$$

since it is actually the electric field interference that is creating the lattice depth. Let's introduce a power ratio β such that $I_2 = \beta I_1$. Then the power needed in beam 1 is

$$P_1 = \frac{U_0}{\sqrt{\beta}} \frac{\omega^3}{2c^2} \frac{\Delta}{\Gamma} \frac{w_0^2}{2} e^{2\frac{|\vec{r}|^2}{w_0^2}} \quad (2.57)$$

and $P_2 = \beta P_1$. We relate this to the dimensionless lattice depth using the relation $U_0 = \tilde{U}_0 E_r$.

This equation was verified with experimental data from December 14th and 15th, 2022. At the chamber we needed 15.6 mW and 6.0 mW in our two beams used to drive a 3rd order Bragg diffraction $\pi/2$ pulse. Pulse width was 20 μ s, beam waist ($1/e^2$ intensity radius) was 7.75 mm. Simulation predicted we needed a lattice depth of $11.3E_r$ for these conditions. I'll outline the calculation here just to avoid any confusion in the future. Here, P_1 is the smaller of the two powers (6.0 mW):

$$P_1 = \frac{11.3 E_r}{\sqrt{2.6}} \frac{2\pi 351 \text{ THz}}{2(3 * 10^8 \text{ m/s})^2} \frac{2\pi * 1.81 \text{ GHz}}{2\pi * 5.22 \text{ MHz}} \frac{(7.75 \text{ mm})^2}{2} * 1 \quad (2.58)$$

where $E_r = \hbar * 2\pi * 2066 \text{ Hz}$. This gives $P_1 = 5.96 \text{ mW}$, so we have a 1% error. This calculation is definitely not accurate to this level of accuracy - we knew the beam waist to 5% accuracy, and the experimental Rabi flopping likely over-drove the on-axis atoms, which is what we calculate here. In any case, the above is likely accurate at the 5-10% level which should at least be free of 2 or π errors.²⁹

frequency component could be retro-reflected to create a lattice with this lattice depth, or two frequency components each with power P could co-propagate and then be retro-reflected each other. For the latter, this equation describes the lattice depth of one of the moving lattices formed by e.g. the up-going ω_1 beam interfering with the down-going ω_2 beam.

²⁹not to say that there couldn't be errors that are cancelling out...

Chapter 3

Symmetric Bloch oscillations of matter waves

3.1 Pre-introduction

The work in this chapter was published as [65]. When I first joined the fine-structure constant project, I was working on fixing their Bloch oscillation simulations and I was peppering Richard Parker a non-stop barrage of questions. One of the questions was - what happens when you bring the two lattices during multi-frequency Bloch oscillations to zero velocity difference? If you start the lattices at the same velocity and then accelerate the lattices away from one another, could you create a beamsplitter? According to our PI Holger Mueller, some Stanford folk were discussing the possibility of this process back in the day but they decided that it wasn't possible for the process to be adiabatic. I had just worked out how to simulate normal Bloch oscillations correctly and testing this question was a very simple extension of the existing simulation. It worked! Weicheng Zhong, an old labmate, helped me run the old- α experiment in order to realize experimental symmetric Bloch oscillations. In the background of this experimental push, I was working on the theory to show why it's coherent, and found some beautiful symmetrical physics.

Experimentally, we demonstrated a coherent momentum splitting within an interferometer of up to $240\hbar k$ -the second highest to date at the time of publication. This was no small achievement. That said, I do want to be frank with the reader about some pitfalls in using symmetric Bloch oscillations for metrology. I realized late in the publication process that symmetric Bloch oscillations play very poorly with finite velocity width atom clouds (in the axial direction). If an atom is moving slightly upwards, it is much more likely to be swept up by the upward-moving lattice instead of being coherently split in two. This symmetry breaking means that you can't just ramp slower to get a more adiabatic/coherent process - actually you need to ramp quickly in order to address most of the atoms in your finite-temperature atom cloud. Some of the plots in the paper show efficient beam-splitters that efficiently split the population, but didn't split the individual atomic wavefunctions coherently.

Another nasty effect is that the mechanics of how the reflection / recombination pulses work rely on a coherent interference between the even- and odd-parity ground states. The

relative phase of these states depends on the lattice depth. As a result, lattice depth fluctuations can directly cause errors in the reflection/recombination pulses. Since the atoms acquire a lot of light shift phase in the lattice, these errors scale very poorly with lattice depth shot-to-shot noise. This didn't seem to be an issue in the old-alpha experiment - a) they had a highly optimized pulse-shaper servo that was used for the actual measurement of alpha, and b) we used very fast ramp rate parameters which would have minimized the phase from lattice depth noise.

It's good to keep these in mind while reading this chapter. In any case, the math behind the process is still extremely satisfying - it's not every day that you get to find the energy band structure of a new Hamiltonian using a beautiful symmetry of the problem.

3.2 Abstract

Cold atoms in an optical lattice provide an ideal platform for studying Bloch oscillations. Here, we extend Bloch oscillations to two superposed optical lattices that are accelerated away from one another, and for the first time show that these symmetric Bloch oscillations can split, reflect and recombine matter waves coherently. Using the momentum parity-symmetry of the Hamiltonian, we map out the energy band structure of the process and show that superpositions of momentum states are created by adiabatically following the ground state of the Hamiltonian. The relative phase and velocity of the two lattices completely determines the trajectories of different branches of the matter wave. Experimentally, we demonstrate symmetric Bloch oscillations using cold Cesium atoms where we form interferometers with up to $240\hbar k$ momentum splitting, one of the largest coherent momentum splittings achieved to date. This work has applications in macroscopic tests of quantum mechanics, measurements of fundamental constants, and searches for new physics.

3.3 Introduction

Bloch oscillations and the Wannier–Stark ladder of matter waves in a periodic potential were first studied in the context of electrons in crystals in the presence of a homogenous electric field [7, 87]. Their counterintuitive nature—that a constant electric field should lead to an ac current—triggered a debate about their existence [72, 102] and led to the formulation of criteria for their observability [62]. Bloch oscillations were first experimentally observed in semiconductor superlattices [26, 88], and have since been studied in a wide variety of physical systems ranging from Bloch oscillations of light [101, 21] to cold atoms [18, 92]. Bloch oscillations are particularly useful in matter wave interferometers, which have found widespread applications in precision measurements of fundamental constants [67, 8, 74, 30], tests of the weak equivalence principle [107, 75] and dark energy theories [37, 41], as well as precision gravimetry [106, 96] and gradiometry [4].

Matter wave interferometers use optical lattices to coherently transfer momentum, allowing one to split a matter wave between different spatial trajectories, then later recombine them and create interference. The measured phase can be increased by using larger momentum splitting between the trajectories [67, 8]; Bloch oscillations enable such a process [18, 11]

and have recently shown to coherently transfer the momentum of more than 10^4 photons to the atoms [98]. With two superposed lattices that are independently accelerated, it might even be possible to realize large-momentum-transfer beam splitters for matter waves, by performing Bloch oscillations of two different velocity classes of atoms simultaneously [52]. However, this process has never been demonstrated. Near velocity degeneracy of the two accelerated lattices, it was expected that non-adiabatic effects would prevent coherent ground state dynamics. Instead, Bloch oscillations have only been used to accelerate atoms after an initial momentum splitting was already made with Bragg diffraction [61, 67], resulting in up to $408\hbar k$ momentum splittings [1, 32], where k is the wavevector of the laser.

Here, we show that Bloch oscillations of atoms in two symmetrically accelerated lattices can remain adiabatic and coherent even as the two lattices pass through velocity degeneracy. Theoretically, we show that it is possible to split, reflect, and recombine atoms simply by allowing them to adiabatically follow the ground state of the Hamiltonian while accelerating the two lattices. The dynamics result in symmetric Bloch oscillations where the matter wave is in a coherent superposition of interacting with each of the two lattices, and the relative phase and velocity of the two lattices completely determines the trajectories of different branches of the matter wave. Experimentally, we demonstrate symmetric Bloch oscillations and realize $240\hbar k$ coherent momentum splitting of a superposition state as well as interferometry with nearly fully-guided matter waves.

Using only accelerated lattices for momentum transfer is desirable for a number of reasons. In comparison with resonant processes such as Bragg diffraction, 1) the dynamics are adiabatic, and can therefore be much more efficient per $\hbar k$ momentum transfer, 2) the processes require less laser power, 3) the velocity class of atoms addressed can be larger, relaxing temperature requirements on atom clouds, and 4) the optical lattices prevent thermal expansion of the atom cloud, further relaxing temperature requirements. As a result, symmetric Bloch oscillations can find applications in next-generation precision measurements of fundamental constants, searches for gravitational waves, and searches for new physics [100, 75, 31, 35, 41].

Section 3.4 presents a theoretical treatment of the Hamiltonian and the resulting dynamics. The Hamiltonian is symmetric under momentum inversion, allowing one to simultaneously diagonalize the Hamiltonian in momentum parity and energy. For the beamsplitter process described above, we show that an atom adiabatically follows the even-parity ground state of the Hamiltonian. The momentum-parity basis is then used to study effects such as non-adiabatic losses, dynamics while ramping the lattices through velocity degeneracy, and effects from different experimental imperfections.

In section 3.5, we describe how we implement symmetric Bloch oscillations experimentally. We use the relative phase between the two lattices to control the populations in the two lattices after ramping through velocity degeneracy; in effect, this creates a fully tunable matter-wave switch each time the lattices cross through velocity degeneracy. We demonstrate the first interferometers created only using accelerated lattices, including a Mach–Zehnder (MZ) interferometer with a momentum splitting of up to $240\hbar k$. Prior to this work, the largest momentum transfer from a single beamsplitter operation was $24\hbar k$ [58]. In order to confirm that symmetric Bloch oscillations are first-order phase coherent, we implement a differential measurement between two simultaneous MZ interferometers and see a stable phase between the interferometer outputs.

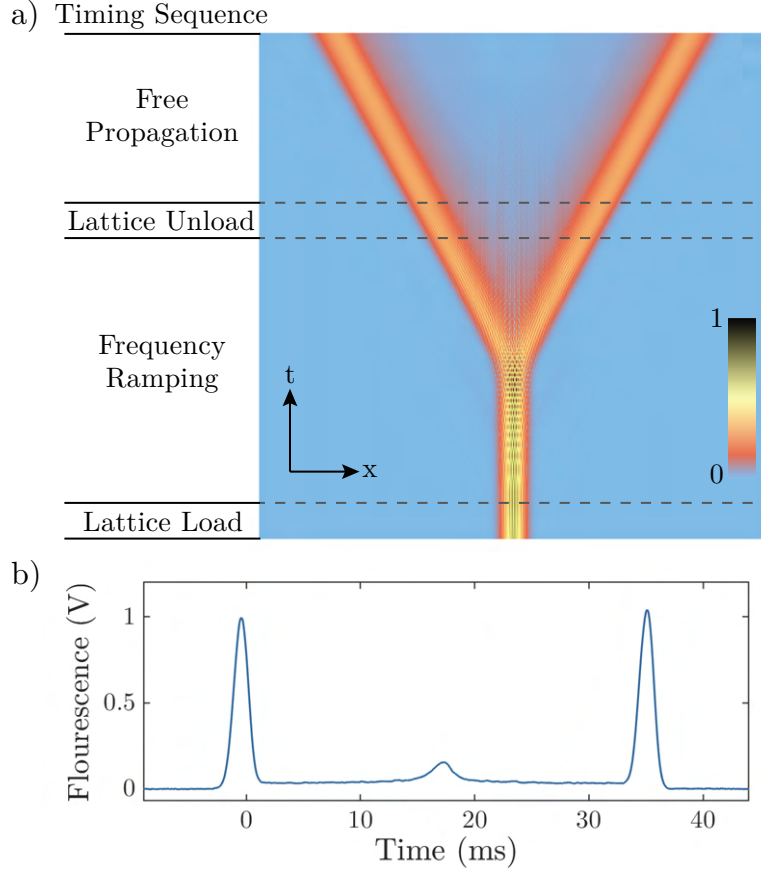


Figure 3.1: a) Density plot of $|\psi(x, t)|$ from numerical simulation of a symmetric beamsplitter with a lattice depth $U_0 = 1.3 E_r$ and ramp rate $r = 0.2 \omega_r^2$. Timing sequence is indicated on the left. The initial wavefunction is a Heisenberg-limited Gaussian wavepacket with velocity spread $0.05 v_r$, corresponding to our experimental temperature. Frequencies are ramped for one Bloch period, corresponding to a $4\hbar k$ splitting between arms. b) Experimental time of flight fluorescence trace showing an efficient $60\hbar k$ beamsplitter with a ramp rate $r = 0.26 \omega_r^2$, and a lattice depth of around $1.5 E_r$.

3.4 Theory

When two superposed optical lattices are far apart in velocity, it is well known that atoms can undergo efficient Bloch oscillations in either of the lattices [61, 67, 1, 32]. Near velocity degeneracy, however, it was previously expected that near-resonant effects from the second lattice would cause too large of a perturbation to the standard Bloch oscillation dynamics to permit an efficient beamsplitter. We first derive a unitary transformation that isolates the relevant dynamics (Sec. 3.4.1), and then show that the effects of the perturbation terms can remain small within the rotating wave approximation under certain conditions (Sec. 3.4.2). Throughout the analysis, it is useful to stress the parallels between Bloch oscillations in a single lattice (SLBO) and Bloch oscillations in two lattices which we call dual-lattice Bloch oscillations (DLBO). The simplified DLBO Hamiltonian is nearly identical to the SLBO Hamiltonian, differing only in being invariant under momentum inversion. As a result, the eigenstates of DLBO are symmetric and anti-symmetric in momentum space.

We then study non-adiabatic loss mechanisms, which include standard Landau–Zener tunneling due to avoided level crossings as well as higher-order transitions which are possible due to perturbation terms dropped in the rotating wave approximation (Sec. 3.4.3). These conditions are combined to place limits on the permissible lattice accelerations and lattice depths, and in total they allow for the DLBO to approach 100% efficiency in the limit of slowly accelerated lattices (Sec. 3.4.4). The dynamics are also discussed for lattices that are ramped through velocity degeneracy, showing that an offset laser phase can be used to coherently control the output population in the two lattices (Sec. 3.4.5). Lastly, we discuss some important experimental requirements in order to realize these methods in the laboratory (Sec. 3.4.6), and supporting material is left for Sec. 3.7.

3.4.1 Hamiltonian and unitary transformation

SLBO are most easily studied using a coordinate system that is comoving with the accelerating lattice [14, 45, 68], and a unitary transformation can be used to boost the Hamiltonian between the atom’s inertial frame and the accelerating lattice frame [48, 68]. For DLBO, it is not possible to transform to a coordinate system that is simultaneously comoving with both lattices. Instead, using a basis of momentum states it is possible to independently transform each momentum state so that positive (negative) momentum states are boosted to a coordinate system comoving with the positively (negatively) accelerating lattice. This unitary transformation is shown to capture the core coherent dynamics of DLBO. The analysis that follows is relevant for zero temperature atoms comoving with the initially degenerate lattices: a similar analysis can be explored for atoms with a small initial velocity, and the band structure of the Hamiltonian can still be studied. One finds that any initial velocity breaks the parity symmetry discussed in the following sections and leads to asymmetric dynamics. A full analysis is beyond the scope of this paper.

We begin with a Hamiltonian containing the AC Stark shift of two superposed optical lattices that are far detuned from single-photon transitions (see Fig. 3.2). Experimentally, the lattices are realized with one upward-propagating laser frequency ω_1 , and two downward-propagating frequencies $\omega_2 \pm \omega_m(t)$. We work in the frame of reference where $\omega_1 = \omega_2 = \omega$, and denote $\omega_{\pm} = \omega \pm \omega_m(t)$. The relative speed of the two lattices is given by $\omega_m(t)/k$,

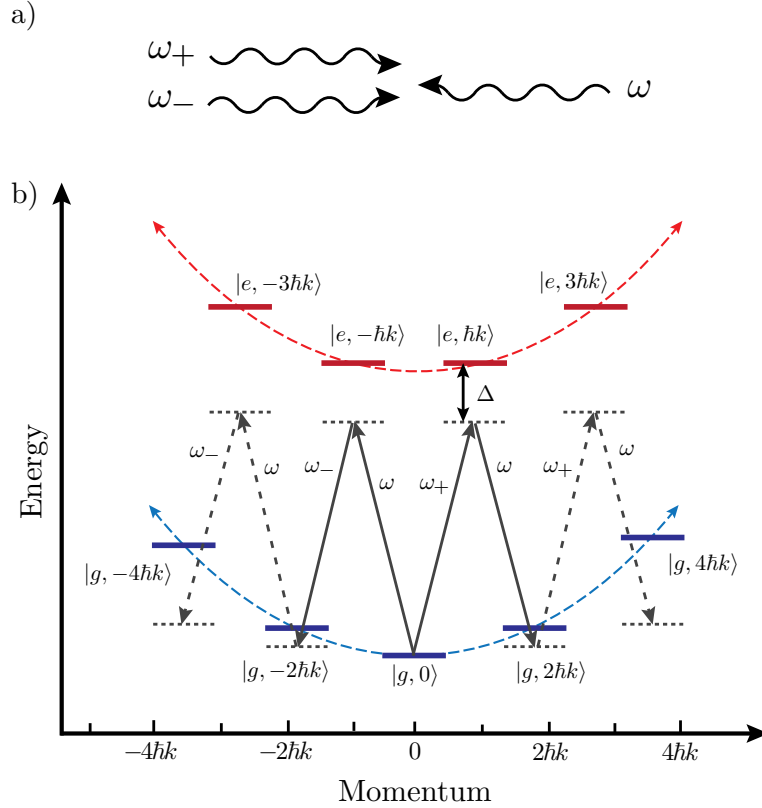


Figure 3.2: a) Counter-propagating lasers form two superposed optical lattices. The frequency differences are $\omega_+ - \omega = \omega - \omega_- = \omega_m$. b) Energy-momentum level diagram showing relevant atomic states. The lasers drive two-photon transitions between neighboring momentum states such that the atom remains in the same internal ground state. The detuning from the excited states Δ (many GHz) is much larger than the separation between adjacent ground states (few kHz). As the modulation frequency ω_m is swept away from zero, the lasers sweep past a succession of two-photon transitions between adjacent ground states. Off-resonant transitions driven by the extra oscillating terms in the Hamiltonian (Eq. 3.4) are omitted for clarity.

where k is the wave number of the laser defined as $k = \omega/c$. Two-photon transitions leave atoms in the same internal state but different external momentum states. After adiabatic elimination of the excited state, the Hamiltonian for an atom in these two optical lattices can be written as:

$$\begin{aligned}
H_{\text{BBS}}(t) &= \frac{\hat{p}^2}{2m} + \frac{U_0}{2} \left(\cos \left[2k_+ \hat{x} + \int_0^t \omega_m(t') dt' + \phi_1 \right] \right. \\
&\quad \left. + \cos \left[2k_- \hat{x} - \int_0^t \omega_m(t') dt' + \phi_2 \right] \right) \\
&= \frac{\hat{p}^2}{2m} + U_0 \cos[2k\hat{x}] \cos \left[\int_0^t \omega_m(t') dt' + \phi_0 \right]. \quad (3.1)
\end{aligned}$$

Constant terms are dropped in the second form, which will be used for analytics and simulation. The wave numbers $k_+ = \omega_+/c$ and $k_- = \omega_-/c$ are nearly identical to k , so we approximate $k_+ \approx k_- \approx k$ in the second form as well. For Cs atoms separated by $n = 1000$ photon momenta, k_+ , k_- , and k differ by less than one part in 10^8 . The phases ϕ_0 , ϕ_1 , and ϕ_2 are offsets between counter-propagating lasers at time $t = 0$. The lattice depth $U_0 = \hbar\Omega_R^2/(2\Delta)$ is the AC Stark shift for a single, far-detuned lattice [14], where Δ is the detuning from the excited state and Ω_R is the on-resonance Rabi frequency between the ground and excited states. The integral $\int_0^t \omega_m(t') dt'$ keeps track of the phase evolution of the lattice for time dependent frequencies. Specializing to linear frequency ramps with rate r , the modulation frequency can be written as $\omega_m(t) = rt$ so that the lattices are velocity degenerate at time $t = 0$ and $\int_0^t \omega_m(t') dt' = rt^2/2$. This ramp rate corresponds to an acceleration $a = r/2k$.

We now write the Hamiltonian in a momentum-state basis $|l\rangle$, where l is an integer that labels the basis states such that the state $|l\rangle$ has $2l\hbar k$ momentum. Plane-wave basis states are a good approximation to initial atomic states when the velocity spread is much smaller than the recoil velocity $v_r = \hbar k/m$. Projected into this basis, the Hamiltonian is:

$$\begin{aligned}
H &= \sum_{l=-\infty}^{\infty} \left(\frac{(2l\hbar k)^2}{2m} |l\rangle\langle l| \right. \\
&\quad \left. + U_0 \cos \left(\frac{rt^2}{2} + \phi_0 \right) (|l\rangle\langle l+1| + |l\rangle\langle l-1|) \right) \quad (3.2)
\end{aligned}$$

The unitary transformation used to boost the different momentum states in this Hamiltonian is given by:

$$U = \sum_{l=-\infty}^{\infty} e^{i\frac{d(t)|\hat{p}|}{\hbar}} e^{i\frac{\theta(t)}{\hbar}} |l\rangle\langle l| \quad (3.3)$$

where $d(t) \equiv at^2/2 + \phi_0/k$ and $\theta(t) \equiv ma^2t^3/6$. The first term corresponds to the position translation operator, and the absolute value sign ensures that positive momentum states are

translated with the positive-moving lattice while negative momentum states are translated with the negative-moving lattice. The $d(t)$ term in Eq. (3.3) also absorbs the offset phase ϕ_0 into the definition of the basis states. The $\theta(t)$ in Eq. (3.3) corresponds to a global energy shift to each state such that the energy of the ground states comoving with either of the lattices stays near zero at all times [48]. See Section 3.7.1 for the analogous treatment of the SLBO Hamiltonian.

The transformed Hamiltonian $H' = UHU^\dagger + i\hbar\frac{dU}{dt}U^\dagger$ is:

$$\begin{aligned}
H' = \sum_{l \neq 0} \left[\frac{(2|l|\hbar k - Ft)^2}{2m} |l\rangle\langle l| \right. \\
+ \frac{U_0}{2} \left(1 + e^{is_l(rt^2+2\phi_0)} \right) |l\rangle\langle l+1| \\
+ \left. \frac{U_0}{2} \left(1 + e^{-is_l(rt^2+2\phi_0)} \right) |l\rangle\langle l-1| \right] \\
+ \frac{(Ft)^2}{2m} |0\rangle\langle 0| + \frac{U_0}{2} \left(1 + e^{-i(rt^2+2\phi_0)} \right) (|0\rangle\langle 1| + |0\rangle\langle -1|) \quad (3.4)
\end{aligned}$$

where $s_l \equiv l/|l|$ is the sign of the momentum state, and the force $F = rm/2k$ is adapted from the standard treatment of SLBO [14].

The nearest-neighbor coupling terms proportional to $|l\rangle\langle l \pm 1|$ include both a stationary term and an oscillating term. In a two-level system, oscillating coupling terms of this type can be dropped under a rotating wave approximation (RWA) provided the terms time-average to zero on the relevant timescale of the dynamics. Here, the couplings between neighboring momentum states can be treated with an analogous RWA to arrive at the reduced DLBO Hamiltonian:

$$\begin{aligned}
H_{\text{DLBO}} = \sum_{l=-\infty}^{l=\infty} \frac{(2|l|\hbar k - Ft)^2}{2m} |l\rangle\langle l| \\
+ \frac{U_0}{2} (|l\rangle\langle l+1| + |l\rangle\langle l-1|) \quad (3.5)
\end{aligned}$$

The validity of this RWA is discussed in Sect. 3.4.2, where we derive bounds on the ramp rate for which the Hamiltonian in Eq. (3.5) is valid.

The DLBO Hamiltonian in Eq. (3.5) and the SLBO Hamiltonian derived in Section 3.7.1 are nearly identical; the only difference is the absolute value $|l|$ in the kinetic energy term for H_{DLBO} , which makes H_{DLBO} symmetric under momentum inversion. This symmetry is already apparent in the original Hamiltonian (3.1), which commutes with a momentum inversion operator. Using a basis of momentum eigenstates that are also eigenstates of momentum-parity, the even- and odd-parity states are decoupled.

Figure 3.3 (a,b) shows the energy band structure over time of the Hamiltonian (3.5) for even- and odd-parity states, respectively, where the two lattices are ramped away from velocity degeneracy beginning at time $t = 0$. The energy bands are calculated by finding

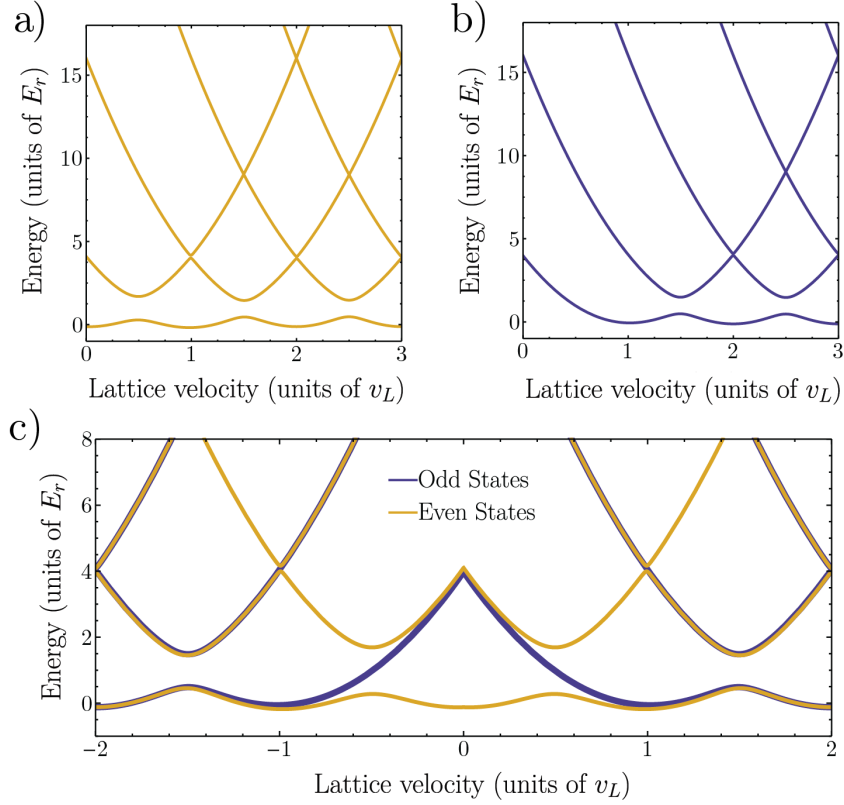


Figure 3.3: Energy band structures of the reduced Hamiltonian (3.5) as a function of the lattice velocity, using a lattice depth $U_0 = 1E_r$. The lattice velocity is defined as $v_L = rt$, such that the two lattices at time t have velocities $\pm v_L$. a) Even-parity and b) Odd-parity energy eigenvalues starting from velocity degeneracy. c) Combined band structure as lattices are ramped through velocity degeneracy.

eigenvalues of a truncated version of the Hamiltonian in Eq. (3.5) as a function of time. Note that in plotting the energy bands in Fig. 3.3c, for negative times we use the substitution $d(t) \rightarrow -d(t)$ in Eq. (3.3) in order to use the coordinate frame comoving with the lattices driving amplitude towards zero momentum instead of driving amplitude away from zero momentum.

A beamsplitter can be understood as an atom adiabatically following the even-parity ground state of the Hamiltonian (3.5), and higher efficiency beamsplitters can be achieved by making the process more adiabatic. At every time $t = (a + 1/2)T_B$ for integers $a \geq 0$, where the Bloch period $T_B = 8\omega_r/r$ and the recoil frequency $\omega_r = \hbar k^2/(2m)$, there is a level crossing such that the even-parity state receives an additional $4\hbar k$ momentum splitting; the positive momentum component of the even state acquires an additional $+2\hbar k$ momentum and the negative momentum component acquires an additional $-2\hbar k$ momentum. This is the momentum-symmetric analogue of SLBO in the ground Bloch band, where atoms receive $2\hbar k$ momentum at the edge of the first Brillouin zone at each avoided level crossing between the ground band and first excited band.

3.4.2 Limits on ramp rate from the rotating wave approximation

A RWA can be used to drop the oscillating coupling terms in Eq. (3.4) provided that the time-average of the oscillating term e^{irt^2} is $\ll 1$ on the relevant timescale of the dynamics, namely the duration of first level crossing between the ground even band and the first excited even band. This crossing occurs at time $t = T_B/2$, and the time interval during which the level crossing happens is given by $\Delta t = 2\sqrt{2}U_0/(\hbar r)$. A simplified form of the resulting inequality gives an upper limit on the ramp rate for which the RWA is valid:

$$r \ll 4U_0(2\sqrt{2}E_r - U_0)/\hbar^2 \quad (3.6)$$

where we define the recoil energy $E_r = \hbar\omega_r$. The RWA is therefore valid in the limit as $r \rightarrow 0$. See Section 3.7.3 for a full derivation of this condition.

The validity of the RWA can be further studied with numerical simulation. By solving for the evolution of $|\psi(t)\rangle$ from the Hamiltonian in Eq. (3.4), the full state evolution is captured without using the RWA. We numerically integrate the Schrödinger equation with the Hamiltonian Eq. (3.2). The initial condition is a free particle (plane-wave) momentum state which is adiabatically loaded into the lattice; the modulation frequency is then ramped to its final value, and finally the lattice is adiabatically unloaded. This state evolution can then be compared with the eigenstates of the Hamiltonian in Eq. (3.5) after the RWA. Fig. 3.4a) shows the probability amplitude in the ground state of Eq. (3.5) during the frequency ramping, defined as $P_0(t) = |\langle +_{gs}(t)|\psi(t)\rangle|^2$. The state $|+_{gs}(t)\rangle$ denotes the even-parity ground state of Hamiltonian (3.5) as a function of time. Fig. 3.4a) shows that the true state evolution is nearly identical to that of the ground state of the Hamiltonian in Eq. (3.5), which generally holds true when Eq. (3.6) is satisfied.

To stress the parallel between SLBO and DLBO, we also plot the probability amplitude in the ground state for SLBO using eigenstates calculated from the Hamiltonian in Eq. (3.10) in Section 3.7.1. In both SLBO and DLBO, the states pass avoided level crossings at times $t = (a + 1/2)T_B$ for integer a , where there is mixing with the second band as well as Landau–

Zener tunneling losses, which are discussed in Sec. 3.4.3. The dual-lattice simulation doesn't project perfectly onto the ground eigenstate around time $t = 0$ due to the perturbation terms dropped in the RWA.

3.4.3 Limits on ramp rate from Landau–Zener tunneling and higher-order transitions

Non-adiabatic Landau–Zener losses arise from the level crossings in Fig. 3.3 between the first and second even-parity energy bands. For SLBO with weak lattices and slow ramp rates, the survival probability per Bloch oscillation is given by $P_{LZ} = 1 - e^{-2\pi\Gamma_1}$ where $\Gamma_1 = U_0^2/(4\hbar^2 r)$ is the Landau–Zener parameter [63, 83]. For ramp rates $r < \omega_r^2$, this formula also describes losses from all level crossings of the DLBO Hamiltonian, Eq. (3.5), except for the two level-crossings at $t = \pm T_B/2$. These two crossings between even-parity eigenstates have an energy gap that is increased by a factor of $\sqrt{2}$, as derived in Section 3.7.2. The Landau–Zener parameter Γ_2 for these two crossings is therefore given by $\Gamma_2 = U_0^2/2\hbar^2 r$. All subsequent crossings in DLBO have the same energy gap as SLBO and are described by the same tunneling parameter Γ_1 . The dual-lattice beamsplitter is therefore more robust to Landau–Zener losses at the first level crossing than SLBO at a fixed lattice depth U_0 , as shown in Fig. 3.4a).

Fig. 3.4b) shows the simulated efficiency of a single Bloch oscillation at a constant Landau–Zener parameter for both the SLBO and DLBO Hamiltonians in Eq. (3.2) and (3.7) respectively. The efficiency is defined as the total population in the desired final momentum states relative to the initial population. In order to have the same expected Landau–Zener losses for both simulations, the SLBO lattice depth is increased by a factor of $\sqrt{2}$ for each ramp rate compared to the DLBO simulation such that $\Gamma_1 = \Gamma_2 = 0.3$. There is asymptotic agreement with the Landau–Zener formula for ramp rates $r \ll \omega_r^2$ for both single-lattice and dual-lattice level crossings, as well as additional oscillatory behavior of the DLBO efficiency compared to the SLBO efficiency owing to the oscillatory terms dropped in the RWA.

The rotating terms being dropped in the RWA can also contribute to higher-order processes that couple amplitude from the ground band to higher energy bands, and are further discussed in Section 3.7.4. The dominant loss channel is a third-order transition that couples the first and second energy levels around time $t = T_B/6$. These higher-order losses place a lower limit on the ramp rate for a fixed lattice depth, below which losses from the ground band begin to be appreciable.

3.4.4 Comparison of limits on the ramp rate

The RWA condition in Eq. (3.6) and Landau–Zener tunneling losses both place an upper limit on the ramp rate. For Landau–Zener losses, efficient dynamics require $r \ll (\pi/2)U_0^2\hbar^2$; when $U_0 \lesssim \sqrt{2}E_r$, the RWA condition in Eq. (3.15) is automatically satisfied if the lattice depth is large enough to sufficiently suppress Landau–Zener tunneling. The RWA that leads to the Hamiltonian (3.5) is therefore asymptotically correct in the limit $r \rightarrow 0$ provided that $\hbar\sqrt{r} \ll U_0 \lesssim \sqrt{2}E_r$. On the other hand, when $U_0 \gtrsim \sqrt{2}E_r$, both the RWA condition and the standard Landau–Zener criterion begin to fail because the time windows for successive transitions begin to overlap non-negligibly.

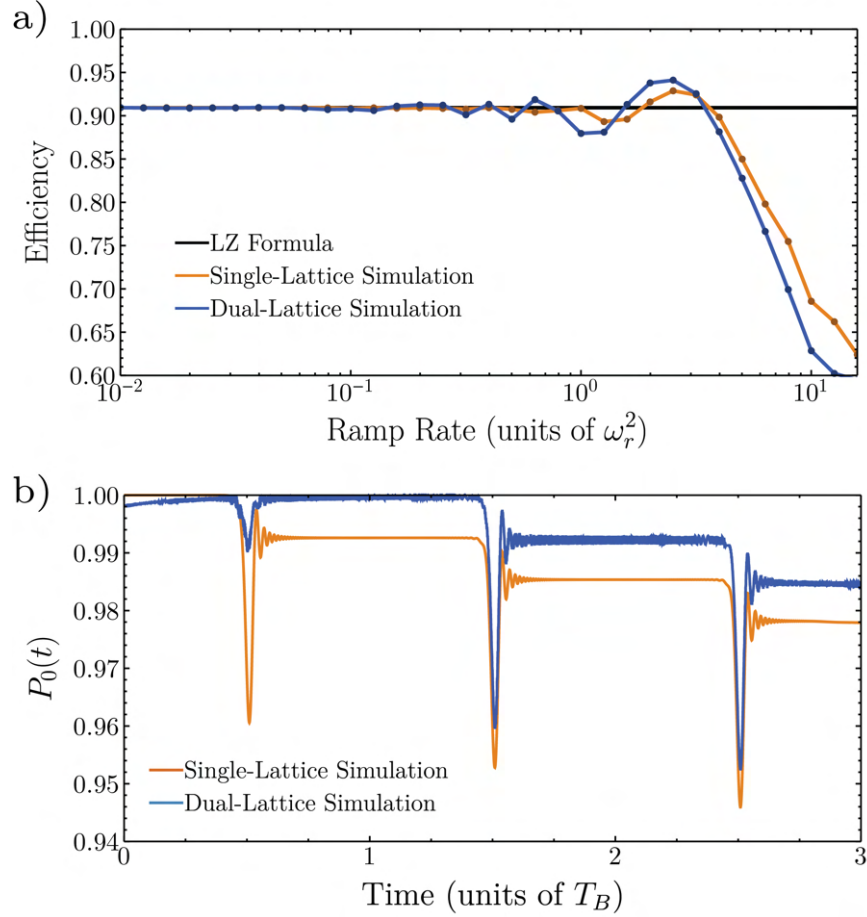


Figure 3.4: Comparison of single-lattice and dual-lattice Bloch oscillations. a) Probability amplitude in the ground state over three Bloch periods. A lattice depth of $U_0 = 0.5E_r$ and ramp rate $r = 0.02\omega_r^2$ are used for both simulations. The lattice depth is intentionally chosen to be low in order to illustrate loss mechanisms for SLBO in comparison with DLBO. See text for discussion. b) Simulation of efficiencies after one level crossing. For each ramp rate, the lattice depth is chosen to keep the Landau–Zener (LZ) parameters constant at $\Gamma_1 = \Gamma_2 = 0.3$ such that the expected losses from the LZ formula are constant. The atom begins in the ground state at time $t = 0$ with $\omega_m(t = 0) = 0$ and $\phi_0 = 0$.




Figure 3.5: Numerical simulation of beamsplitter losses showing the dependence on the two most important parameters in the Hamiltonian: frequency ramp rate and lattice depth. Darker color corresponds to lower losses, or higher efficiency. The simulation includes adiabatic loading of lattice, frequency ramping for four Bloch periods, and adiabatic unload, such that the final momentum splitting is $16\hbar k$. Efficiency is defined as the probability amplitude on the desired momentum states after unloading the lattice. See text for discussion of loss mechanisms.

Higher-order losses place a lower limit on the ramp rate, and for $r \leq \omega_r^2$, this limit and the upper limits on the ramp rate from Landau–Zener losses and the RWA condition can all easily be satisfied. Because of the non-linear scaling of these different limits on the ramp rate, the maximum possible efficiency of the processes quickly approaches 1 as $r \rightarrow 0$; for $r = 0.5\omega_r^2$, the maximum efficiency of the initial $4\hbar k$ momentum splitting in a Bloch beamsplitter is already $> 99\%$.

Fig. 3.5 illustrates beamsplitter losses as a function of the ramp rate r and the lattice depth U_0 . Losses towards the top-left of the plot correspond to Landau–Zener tunneling losses, and losses towards the bottom-right correspond to higher-order transitions. Moving towards higher lattice depths and ramp rates, the maximum efficiency of the beamsplitter decreases because of the competing loss mechanisms.

The two loss channels result in non-zero wavefunction amplitude in momentum states different from the target states, and these additional momentum states could contribute to parasitic interferometers. This analysis is beyond the scope of the paper, however we note that there exist methods to reduce the effects of parasitic interferometers [78].

3.4.5 Crossing through velocity degeneracy

In addition to a beamsplitter, one can also ramp the two lattices through velocity degeneracy to create atom mirrors and combiners. This process has previously been attempted experimentally [2], but the dynamics were seen to be inefficient and uncontrolled because the ramp rate, lattice depth, and relative phase between lattices were not optimized. The intuition for the dynamics through a level crossing are described below, and for a more mathematical treatment see Section 3.7.5.

Consider two optical lattices with velocities that are initially far apart. One arm of an interferometer that is initially comoving with one of the two lattices can be understood as a superposition of an even-parity and an odd-parity ground state. Relative phase shifts between the even and odd states causes amplitudes to add constructively or destructively for positive or negative momentum states, which means that a controlled relative phase shift between the even- and odd-parity states can be used to control the momentum distribution of the atomic state after crossing through velocity degeneracy.

Figure 3.3c shows the band structure as the lattices are ramped through velocity degeneracy at time $t = 0$. Far from velocity degeneracy, the even and odd ground state energy bands overlap and have the same level crossing structure. Near time $t = 0$, however, these energy bands deviate because, by definition, an odd-parity state in momentum space cannot have amplitude on the zero-momentum basis state $|0\rangle$. As a result, when crossing through velocity degeneracy the odd-parity ground state has no level crossing coupling momentum into or out of the zero momentum state, so the even parity ground state passes through two additional level crossings at times $t = \pm T_B/2$ compared to the odd parity ground state.

Through the coherent interactions with photons from each of the lattices, the relative phase ϕ_0 of the two optical lattices is ultimately added to amplitude in the even-parity state, but not the odd-parity state. As a result, the offset phase ϕ_0 can coherently control the population in the two lattices after a degeneracy crossing. This allows one to create reflection or recombination pulses in an interferometer, and together with the beamsplitter process described previously, this comprises a full set of atom-optics tools for atom interferometry

(see Fig. 3.6 for experimental implementation).

3.4.6 Experimental considerations

The dynamics of symmetric Bloch oscillations are sensitive to the initial velocity distribution of an atom. Efficient beamsplitter dynamics are observed for atoms with velocity spreads of more than $\sigma_v = 0.5v_r$, where σ_v is the standard deviation in velocity of a Heisenberg-limited Gaussian wavepacket. However, this spatial separation does not necessarily result in a superposition state in momentum space. For matter wave sources where different velocity classes are uncorrelated, only amplitude within a certain momentum window Δp results in a superposition state, and amplitude to the left (right) of this window in momentum space will preferentially follow the right-moving (left-moving) lattice [52]. Intuitively, this can be understood by considering the dynamics in the Brillouin zone. When an atom begins at zero velocity, symmetric Bloch oscillations apply a force in both directions, and the quasimomentum can be thought of as being ramped in both directions simultaneously such that the state reaches both edges of the Brillouin zone at the same time, splitting the atom symmetrically in a superposition state. If the atom has some initial velocity, however, it will reach one edge of the Brillouin zone before the other, and as a result amplitude will preferentially be driven by this first transition.

Numerical integration of the Hamiltonian (3.1) can be used to solve for evolution of a wavefunction $\psi(x, t)$ with arbitrary initial conditions (see Fig. 3.1) using the Crank–Nicolson method to discretize the Schrödinger Equation [86, 29]. These simulations confirm that faster ramp rates result in higher fidelity superposition states in momentum states, which in turn results in higher contrast interferometers.

Diffraction phases are fundamental to any asymmetric Bragg diffraction beamsplitter [10, 24], and must be accounted for in precision measurements [67]. For symmetric Bloch oscillations, if the center of the initial atomic velocity distribution is non-zero, the initial state has some projection onto the odd-parity eigenstates which leads to asymmetry and diffraction phases. The symmetry of the Bloch beamsplitter (see Fig. 3.1) ensures that there is no diffraction phase that is fundamental to the technique. An initial velocity of the atoms, however, breaks the symmetry and creates a diffraction phase between interferometer arms. The numerical study discussed in Section 3.7.6 shows that there are “magic” lattice depths where the diffraction phase vanishes. For realistic experimental control over the stability of the lattice depth, the diffraction phase can be limited to ± 10 mRad, independent of the momentum splitting. Increasing the momentum splitting will therefore fractionally suppress the diffraction phase, and diffraction phases can also be measured directly by varying the time between pulses in an interferometer. Note also that an ensemble of atoms with different center velocities will result in phase spreading in an interferometer.

The analytic results derived for Landau–Zener tunneling and the rotating wave approximation only apply to slow ramp rates that satisfy the condition in Eq. (3.6). Experimentally, we use larger ramp rates of up to $r = 10\omega_r^2$ and lattice depths around $8E_r$ in order to maximize interferometer contrast, which is a region of parameter space that breaks the assumptions used to derive this inequality. Although the analytical efficiency predictions break down in this regime, we still observe reasonably efficient dynamics both numerically and experimentally. In fact, the velocity bandwidth of the beamsplitter is larger at faster

ramp rates which results in higher contrast interferometers. See Fig. 3.5 for an illustration of a beamsplitter for different values of lattice depth and ramp rate. Notably, even in regions of parameter space outside where the RWA is valid, one can still achieve relatively low loss beamsplitters.

3.5 Experiment

Our experimental apparatus has been described previously in [67]. A magneto-optical trap of Cesium atoms is launched vertically in an atomic fountain. The cloud is further cooled to a few hundred nK using polarization gradient cooling and Raman sideband cooling. Three successive Raman transitions prepare the atom in the internal state $|F = 3, m_F = 0\rangle$ with a vertical velocity spread around 0.05 recoil velocities v_r .

The frequencies ω_1 and ω_2 in the Hamiltonian (3.1) are ramped in the lab to compensate for Doppler shifts from gravitational acceleration of the atoms such that in the atom’s inertial frame, $\omega_1 = \omega_2 = \omega$. The optical lattice is detuned by +80 GHz (blue) from the Cs D2 line, and is formed from a roughly Gaussian beam with $1/e$ waist of about 3 mm that is retroreflected. The frequency components ω_1 and $\omega_{2\pm}$ are cross-polarized and a quarter waveplate is placed in front of the retroreflecting mirror such that the desired lattices are formed upon retroreflection. The laser intensity is actively stabilized by feeding back to the drive power of an acousto-optic-modulator (AOM) [23].

The modulation frequency $\omega_m(t)$ from the Hamiltonian in Eq. 3.1 determines the velocity splitting between the two lattices. It is generated experimentally by mixing the output of an AD9959 digital frequency synthesizer with a 10-MHz clock and low-pass filtering the output, after which $\omega_m(t)$ is mixed into the drive frequency for an AOM to generate frequency sidebands that are written onto the laser. The offset phase ϕ_0 is a tuneable parameter on the digital frequency synthesizer.

3.5.1 Atom optics with symmetric Bloch oscillations

To create a beamsplitter, atoms are adiabatically loaded into two velocity-degenerate lattices that initially add constructively to form a single lattice, which corresponds to $\omega_m = 0$ and $\phi_0 = 0$ at time $t = 0$ in Eq. (3.1). The modulation frequency is then ramped linearly at a rate r such that $\omega_m = rt$ and the two lattices accelerate away from one another. The resulting momentum distribution is then measured using time-of-flight detection, as shown in Fig. 3.1b. For a $\pm 2n$ -photon beamsplitter, the final atomic state after the beamsplitter is mostly in the $|\pm n\rangle$ states, with a small number of atoms left in the $|0\rangle$ state.

In addition to an initial beamsplitter, a full interferometer sequence requires reflection pulses to reverse momentum of the interferometer arms and a recombination pulse to interfere the two arms together. We find that varying the offset phase ϕ_0 between the two lattices from 0 to π controls the population in the two lattices after the degeneracy crossing, varying from reflection to transmission with a beamsplitter/recombination behavior at an intermediate ϕ_0 . This phase also dictates the interference (“beat”) between the two optical lattices at the time of the modulation frequency zero crossing, as shown in Fig. 3.6b. The optimal phase offsets ϕ_0 are found which maximize population in the desired output channels, as shown in

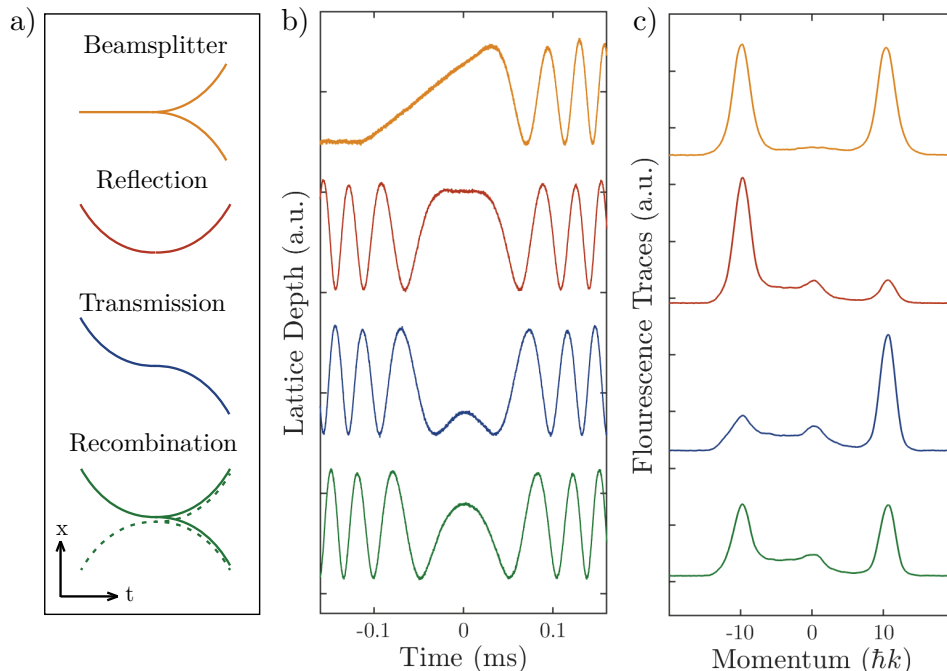


Figure 3.6: Experimental realization of a Bloch beamsplitter (yellow), reflection (red), transmission (blue), and recombination (green) as lattices are ramped through velocity degeneracy. a) Space-time trajectories. b) Intensity profiles of the $\omega_2 \pm \omega_m(t)$ interferometry beams, which are measured by imaging the laser beams on a photodiode just before entering the vacuum chamber. The profiles show beats between the two frequencies, which is the temporal part of the potential in the Hamiltonian (3.1). Time $t = 0$ indicates when $\omega_m = 0$. Different phase offsets ϕ_0 result in different beat profiles on the beam. c) Fluorescence traces of atoms from time-of-flight imaging showing the resulting distribution after various operations.

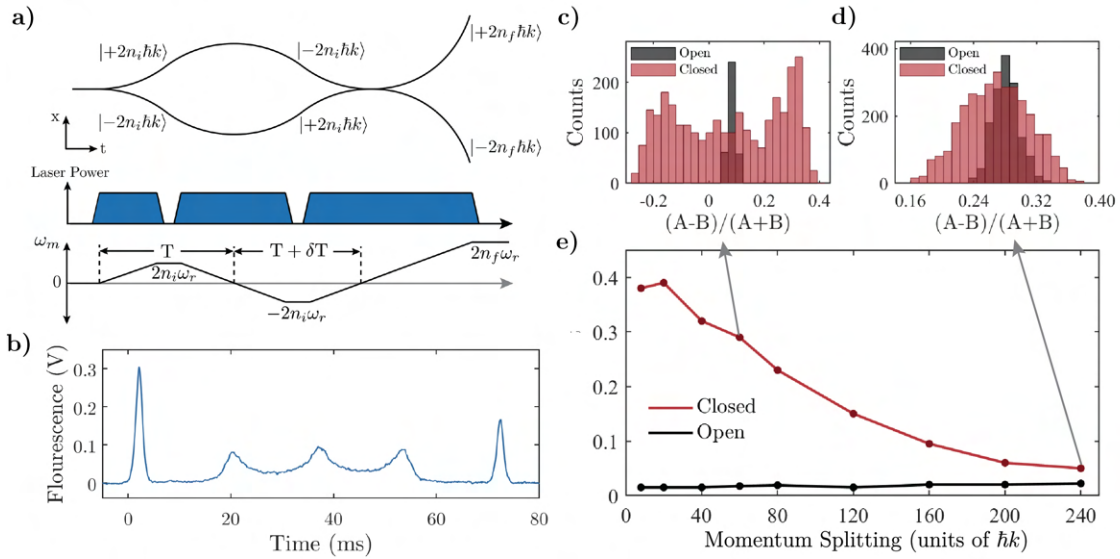


Figure 3.7: Experimental realization of a Mach-Zehnder interferometer. a) Interferometer geometry, laser intensity profile, and profile of the modulation frequency ω_m vs. time. The time interval T is defined based on ω_m -zero crossings. A time offset can be used to open the interferometer to eliminate interference, while $\delta T = 0$ leads to maximum contrast. b) Sample fluorescence trace of a $T = 8.5$ ms, $60\hbar k$ MZ interferometer. c,d) Histogram of population fractions for $60\hbar k$ and $240\hbar k$ momentum splittings, respectively, in $T = 8.5$ ms interferometers. Population fraction is defined as $(A - B)/(A + B)$ where A and B are populations in the two output ports. e) Contrast versus momentum splitting for closed and open interferometers. For all data points, ω_m is ramped at a rate of $r/(2\pi) = 249$ MHz/s.

Fig. 3.6. Note that the optimal phase offset ϕ_0 is dependent on both the ramp rate and the lattice depth U_0 ; the beat profiles shown in Fig. 3.6b are specific to the lattice depth and ramp rate used experimentally, and will need to be optimized anew if either parameter is changed. This is due to the fact that the dynamical phase ϕ_d in Eq. (3.18) is a function of both the ramp rate and the lattice depth. For the parameters used in our experiment, the simulated efficiencies are similar to those realized experimentally, but experimentally we see more atoms lost to the zero momentum state.

3.5.2 Mach–Zehnder interferometer

Combining these techniques, we implement a MZ interferometer, see Fig. 3.7. The sequence starts with a Bloch beamsplitter that is ramped to some final momentum splitting n_i . After this, a reflection sequence is performed and the phase ϕ_0 in the Hamiltonian Eq. (3.1) is arranged as shown in Fig. 3.6b. The two halves of the wavefunction are then interfered using a recombination sequence and the outputs are separated to some final momentum state n_f .

To optimize the contrast of the detected interferences, we need to separate signal atoms from background atoms that arise from loading and unloading the lattices. Using $n_f > n_i$ separates “signal” atoms from those backgrounds in time-of-flight imaging, see Fig. 3.7b.

A ramp rate for ω_m of $r = 2\pi * 249 \text{ MHz/s} = 9.3\omega_r^2$ and a lattice depth of around 8 recoil energies (for each lattice individually) are used, as these parameters resulted in the largest interferometer contrast. The phases ϕ_0 for the two degeneracy crossings were also optimized experimentally to maximize contrast. In between different interferometer operations, we switch the direction of the modulation frequency ramp by switching RF frequency sources for the modulation frequency $\omega_m(t)$, and we adiabatically unload the lattice during this time to avoid losses from the ground state.

We observe up to 40% contrast in a $T = 8.5 \text{ ms}$, $20\hbar k$ interferometer where atoms are guided in the lattices during 16.7 ms of the 17 ms interferometer duration (Fig. 3.7e). Because of vibrational noise in the experiment, it was not possible to observe a stable fringe, so contrast was determined by measuring the fluctuations in the output populations on a histogram. Without changing the laser intensity profile, momentum transfer is increased by changing the profile of ω_m as shown in Fig. 3.7a, and contrast is observed up to $240\hbar k$ momentum splitting.

3.5.3 Gradiometer

Observing contrast in an interferometer does not show that the interferometer is phase-stable. In order to show phase-stability and first-order coherence [76], we also perform a differential measurement between two MZ interferometers in a gradiometer configuration, see Fig. 3.8. In this configuration, phase noise from vibrations is common to both MZ interferometers, so the differential measurement can reveal a stable relative phase. The two MZ interferometers are separated vertically by roughly 11 cm by using a $500\hbar k$ Bloch beamsplitter with a ramp rate of $r = 2\pi * 249 \text{ MHz/s} = 1.9\omega_r^2$. Within each MZ, a momentum splitting of $20\hbar k$, an interferometer time $T_i = 10 \text{ ms}$, and a ramp rate of $r = 2\pi * 249 \text{ MHz/s} = 9.3\omega_r^2$ are used. The slower ramp rate for the first beamsplitter minimizes background atoms in the time-of-flight traces, and the faster ramp rate during the interferometer maximizes contrast.

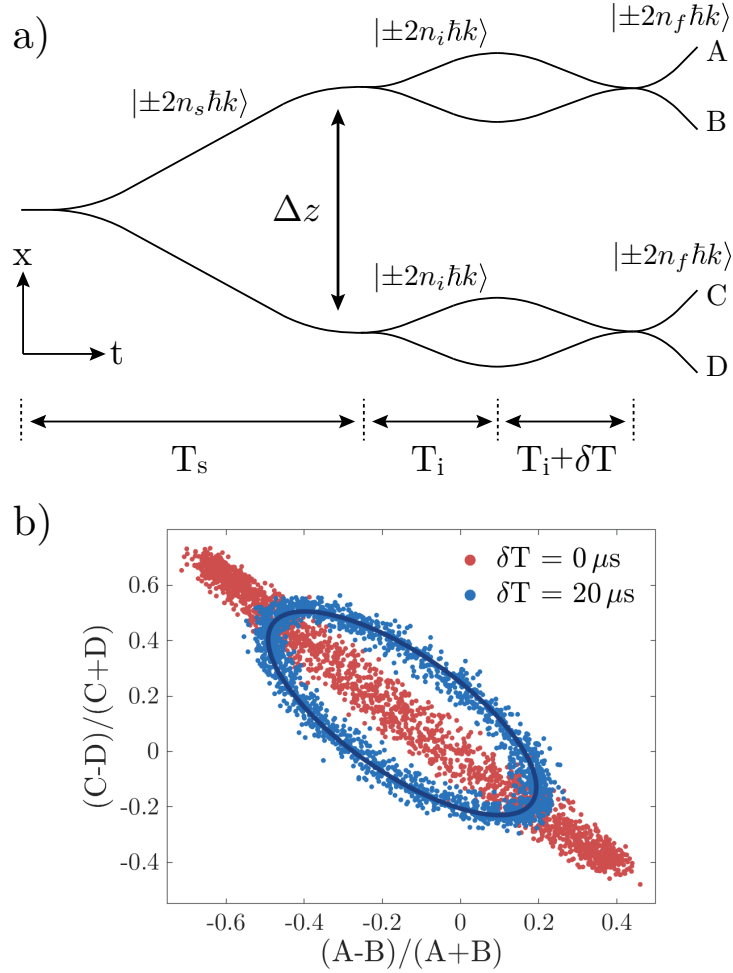


Figure 3.8: a) Schematic of a dual-lattice gradiometer. A and B are populations in output ports of one MZ, and C and D are populations in output ports of the second MZ. b) Parametric plot of data taken using $n_s = 125$ and $T_s = 150$ ms for $\Delta z \approx 11$ cm vertical splitting between the two MZ interferometers. Within the interferometers, $n_i = 5$ and $T = 10$ ms, and $n_f = 10$ to resolve the outputs. The dark line is an ellipse fit to the $\delta T = 20 \mu s$ data.

Phase-stability is observed between the interferometers by plotting the relative populations parametrically (see Fig. 3.8b).

If there is no differential phase acquired between the interferometers, one sees perfect correlation in the outputs, and common-mode vibration noise causes data to fall at different points on this line. Instead, we see that the outputs are anti-correlated, owing to a $\pm\pi/2$ phase shift imprinted on the upper and lower MZ interferometers, respectively, during the opening pulse of the interferometer. Similar phases are well known in higher-order Bragg transitions, and come directly from Schrodinger equation dynamics [23]. Note that this phase is not permitted from symmetry arguments: the opening of the upper and lower interferometers around time T_s are asymmetric, since upper (lower) arm has positive (negative) velocity prior to the splitting.

Differential phase shifts between the two MZs results in an elliptical distribution in the parametric plot. We find that a timing delay δT of the final recombination, as defined in Fig. 3.8a), introduces a controlled phase difference into the interferometer that scales linearly with the timing delay, $\Delta\phi = (40 \text{ rad/ms})\delta t$. Differences in gravity between the two MZs also creates a differential phase shift which is proportional to the gravity gradient. However, this phase is around 5 mrad for the parameters used experimentally and is too small to be observed. The phase coherence between the two MZs demonstrates that the technique is first-order coherent and phase-stable, and can therefore be used for measurements in atom interferometry. We achieve as large as 50% contrast in the differential measurement, which is similar to the largest contrast we ever observed with Bragg diffraction in the same instrument. The contrast is higher than the contrast in the Mach-Zehnder interferometers in Fig. 3.7 because the lattice in the gradiometer configuration is turned off when the lattices are not being accelerated. The timing delay causes loss of contrast because of not fully closing the interferometer.

3.6 Conclusions and Outlook

We have developed new techniques for coherently controlling superpositions of momentum states by generalizing Bloch oscillations to two independently accelerated optical lattices. First, the Hamiltonian was treated analytically, and it was shown that the dynamics can produce efficient and coherent atom optics elements, even when the lattices pass through velocity degeneracy. For slow ramp rates, the process is adiabatic and atoms can adiabatically follow the even-parity ground state of Hamiltonian (3.5). When ramping lattices through velocity degeneracy, the populations in the two lattices can be controlled by changing the relative phase of the two optical lattices, allowing for all atom-optics elements required to form an interferometer. Using only accelerated lattices, we create LMT interferometers with high contrast, and we showed that the resulting dynamics were first-order coherent using a differential measurement.

Compared to existing atom optics techniques [13, 61, 32], DLBO offer a number of advantages. Applications with constraints on laser power and free-fall distance, such as space-based interferometry [22, 49] or portable gravimeters [96], can use these techniques to maximize momentum transfer and thus sensitivity. Being based on adiabatic processes, these methods are robust to fluctuations in experimental parameters like lattice depth or laser frequency

[14]. Symmetric Bloch oscillations are more robust to small laser intensity variations than Bragg diffraction beam splitters, and can eliminate systematic phase shifts known as diffraction phases [24, 66, 42]. Moreover, large momentum transfer can be obtained with modest laser power, whereas in multi-photon Bragg diffraction the required laser intensity scales proportional to n^2 or even n^4 , if scattering losses are to be kept constant [57]. In contrast, the laser power required for DLBO is independent of the momentum splitting, relaxing the laser power requirements in an experiment. Compared to combinations of Bragg diffraction and Bloch oscillations [61, 1], DLBO requires less laser power and can achieve higher efficiencies. For example, two sequential $4\hbar k$ double-Bragg beamsplitters used in reference [32] use a peak lattice depth of $3 - 4E_r$ and achieve a total efficiency around 90%, and higher-order double Bragg pulses require considerably more laser power. In contrast, the $60\hbar k$ beamsplitter in Fig. 3.1b uses a lattice depth of $1.5E_r$ while achieving an efficiency greater than 90%.

A generalization of these dual-lattice techniques shows promise for new measurements of the fine-structure constant α . A set of realistic experimental parameters are outlined in Section 3.7.8, where we show that 10^8 radians of phase are attainable. This paves the way for a measurement of alpha at the 10^{-11} level, an order of magnitude improvement on existing measurements. Another generalization of the Bloch beamsplitter uses a multi-photon, $4n\hbar k$ transition to open the interferometer where $n > 1$. Our numerical simulations show that this multi-photon process also leads to an efficient beamsplitter for appropriate ramp rates and lattice depths, see Section 3.7.4 for further discussion.

The authors thank Matt Jaffe, Victoria Xu, Sven Abend, Ernst Rasel, Justin Khoury, and Tanner Trickle for useful discussions and feedback, as well as the past group members who helped develop the experimental apparatus including Brian Estey, Joyce Kwan, Chenghui Yu, Pei-Chen Kuan, and Shau-Yu Lan. This work was supported by the National Science Foundation Grant No. 1806583, the National Institute of Science and Technology Grant No. 60NANB17D311, and the W.M. Keck Foundation Grant No. 042982. Z.P. acknowledges funding from the National Science Foundation GRFP.

3.7 Supporting material

3.7.1 Unitary transformation for single-lattice Bloch Hamiltonian

In an inertial frame initially co-moving with the atoms, the SLBO Hamiltonian can be written as:

$$H = \sum_{l=-\infty}^{\infty} \left(\frac{(2l\hbar k)^2}{2m} |l\rangle\langle l| + U_0 e^{i\left(\frac{vt^2}{2} + \phi_0\right)} (|l\rangle\langle l+1| + |l\rangle\langle l-1|) \right) \quad (3.7)$$

The Hamiltonian, Eq. (3.10), is derived by transforming this Hamiltonian, Eq. (3.7), into a rotating frame that puts the time dependence of the rotating terms into the diagonal. This

is achieved with the following unitary:

$$U = \sum_{l=-\infty}^{\infty} e^{i\frac{d(t)\hat{p}}{\hbar}} e^{i\frac{\theta(t)}{\hbar}} |l\rangle\langle l| \quad (3.8)$$

$$= \sum_{l=-\infty}^{\infty} e^{il\left(\frac{rt^2}{2} + \phi_0\right)} e^{i\frac{ma^2t^3}{6\hbar}} |l\rangle\langle l| \quad (3.9)$$

with $d(t) \equiv at^2/2 + \phi_0/k$ and $\theta(t) \equiv ma^2t^3/6$. This same transformation is used in reference [16], and it is almost identical to the transformation used in Eq. (3.3), except there is no longer a absolute value sign on the momentum operator. Acting on the Hamiltonian, Eq. (3.7), with the unitary transformation in Eq. (3.8) results in H_{SLBO} :

$$H_{\text{SLBO}} = \sum_{l=-\infty}^{\infty} \frac{(2l\hbar k - Ft)^2}{2m} |l\rangle\langle l| + \frac{U_0}{2} (|l\rangle\langle l+1| + |l\rangle\langle l-1|) \quad (3.10)$$

The Ft term that appears in the kinetic energy is related to the quasimomentum k_q through the relation $\hbar k_q = Ft$.

3.7.2 Symmetrized Hamiltonian

The Hamiltonian in Eq. (3.5) can be explicitly symmetrized by applying a rotation to the basis states. This is achieved by rotating to new basis states that are symmetric and antisymmetric combinations of the free-space momentum basis states, namely we will have (unnormalized) even parity basis states $|+_l\rangle = |l\rangle + |-l\rangle$ and odd parity states $|-_l\rangle = |l\rangle - |-l\rangle$. The zero momentum state remains unchanged under this rotation, as it is already an even-parity state. The following rotation matrix achieves this transformation:

$$R = |0\rangle\langle 0| + \sum_{l>0} \frac{1}{\sqrt{2}} (|l\rangle\langle l| + |-l\rangle\langle -l|) + \frac{1}{\sqrt{2}} (|l\rangle\langle -l| - |-l\rangle\langle l|) \quad (3.11)$$

The Hamiltonian (3.5) can then be rotated to the symmetric Hamiltonian $H_{\text{sym}} = R H_{\text{DLBO}} R^T$ to arrive at the following:

$$\begin{aligned}
H_{\text{sym}} = & \frac{(Ft)^2}{2m} |0\rangle\langle 0| \\
& + \sum_{l>1} \left(\frac{(2l|\hbar k - Ft)^2}{2m} (|+l\rangle\langle +l| + |-l\rangle\langle -l|) \right. \\
& + \frac{U_0}{2} \left(|+l\rangle(\langle +l+1| + \langle +l-1|) + |-l\rangle(\langle -l+1| + \langle -l-1|) \right) \\
& \left. + \frac{U_0}{2} (|+1\rangle\langle +2| + |-1\rangle\langle -2|) + \frac{U_0}{\sqrt{2}} (|0\rangle\langle +1| + |+1\rangle\langle 0|) \right) \quad (3.12)
\end{aligned}$$

In this rotated basis, there is no coupling between $|0\rangle$ and $|-1\rangle$, so we can explicitly see why the odd-parity states have no level crossing at times $t = \pm T_B/2$ in Fig. 3.3c). Moreover, the coupling between $|0\rangle$ and $|+1\rangle$ is $\sqrt{2}$ larger than any of the other couplings, resulting in suppressed Landau-Zener tunneling from the level-crossings of the even-parity ground state at times $t = \pm T_B/2$ in Fig. 3.3c).

3.7.3 Rotating wave approximation condition

To make the rotating wave approximation (RWA) in Eq. (3.4), we average the oscillating term e^{irt^2} over the duration of the transition between momentum states. This term is oscillating most slowly around the first level crossing between the first and second even bands at time $t = T_B/2$. In the limit of small lattice depths $U_0 \ll 4E_r$, the energy gap $E_g(t)$ near this level crossing is given by

$$E_g(t) = \sqrt{\hbar^2 r^2 (t - T_B/2)^2 + 2U_0^2} \quad , \quad (3.13)$$

such that the center of the level crossing occurs at time $t = T_B/2$, and the duration of the level crossing is $\Delta t = 2\sqrt{2}U_0/\hbar r$.

Taking the time average of the rotating term e^{irt^2} over the duration of the level crossing gives the following:

$$\langle e^{irt^2} \rangle \approx -\frac{i\hbar^2 r}{4U_0} e^{i\alpha} \frac{U_0 \cos \beta - 2\sqrt{2}iE_r \sin \beta}{8E_r^2 - U_0^2} \quad (3.14)$$

where we define $\alpha = 2(8E_r^2 + U_0^2)/(\hbar^2 r)$ and $\beta = 8\sqrt{2}E_r U_0/(\hbar^2 r)$, and we have assumed that $r \ll 2(2\sqrt{2}E_r - U_0)^2/\hbar^2$. The rotating term can be dropped so long as this average is small compared to 1, i.e., when

$$|\langle e^{irt^2} \rangle| < \frac{\hbar^2 r}{4U_0(2\sqrt{2}E_r - U_0)} \ll 1 \quad (3.15)$$

or equivalently, $r \ll 4U_0(2\sqrt{2}E_r - U_0)/\hbar^2$. We note that varying the time window of integration in Eq. (3.14) changes the numerical factors in Eq. (3.15), but not the limiting behavior as $r \rightarrow 0$.

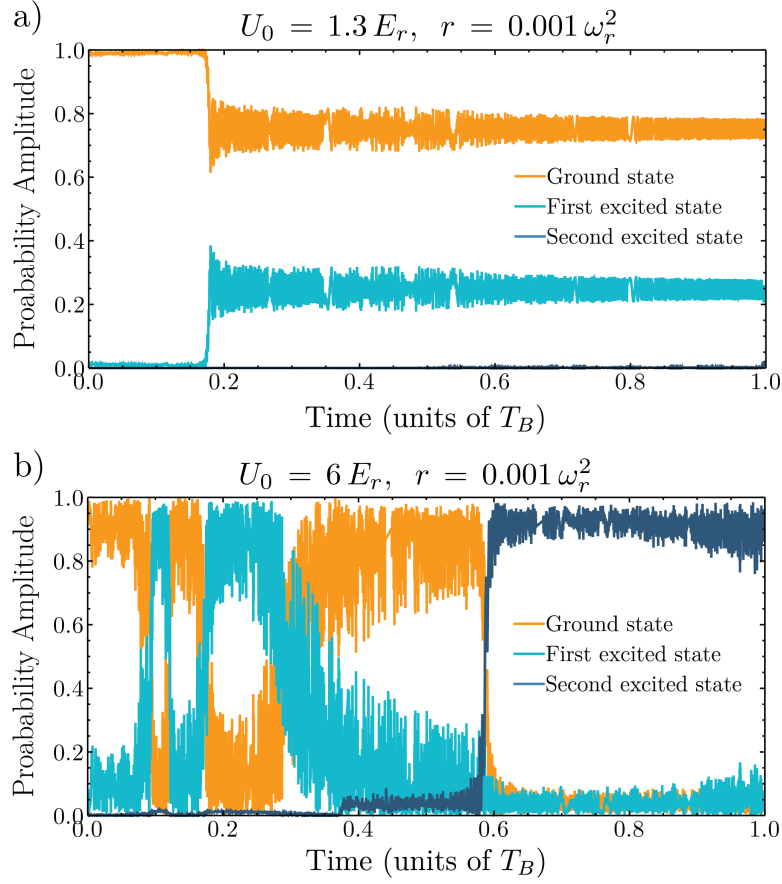


Figure 3.9: Simulations of one Bloch period of a Bloch beamsplitter illustrating losses from the ground band due to higher-order transitions. The states used for determining the probability amplitude are even-parity eigenstates of Hamiltonian (3.5). A slow ramp rate is used so that the various transitions are resolved from one another. a) The first losses to occur are due to a third-order transition coupling the ground state and first excited state. b) A much larger lattice depth shows a number of different higher-order transitions. Before time $t = T_B/2$ there are four separate higher-order resonances between the ground state and first excited state that transfer population between the levels. Around time $t = 0.6 T_B$ there is a transition between the ground state and the second excited state.

3.7.4 Higher-order loss mechanisms

When the lattice depth is too large, the oscillating terms dropped in the rotating wave approximation from the Hamiltonian in Eq. (3.4) can contribute to higher-order parasitic transitions. The dominant loss mechanism at ramp rates $r \ll \omega_r^2$ is a third-order (six-photon) process coupling the states $|0\rangle$ and $|+1\rangle$ around time $t = T_B/6$, where $|+l\rangle$ refers to the symmetrized basis states derived in Section 3.7.2. There are two possible energy and momentum conserving pathways for the transition to occur; $|0\rangle \rightarrow |+1\rangle \rightarrow |0\rangle \rightarrow |+1\rangle$ and $|0\rangle \rightarrow |+1\rangle \rightarrow |+2\rangle \rightarrow |+1\rangle$. For lattice depths much less than the spacing between energy levels, $U_0/2 \ll 4E_r$, the effective coupling between these states scales like $(U_0/2)^3/(4E_r)^2$, which is the same scaling as the Rabi frequency in higher-order Bragg diffraction [57, 99].

During a Bloch beamsplitter, the laser frequencies are swept across this parasitic resonance, as seen in Fig. 3.9a, which can be thought of as a parasitic level-crossing between $|0\rangle$ and $|+1\rangle$; for an efficient Bloch beamsplitter, amplitude should remain in $|0\rangle$ by tunneling through this level-crossing adiabatically. To first order, the adiabatic population transfer to the state $|+1\rangle$ during this level crossing is given by $P_{LZ} = 1 - e^{-2\pi\Gamma} \approx 2\pi\Gamma$ when the Landau-Zener parameter Γ is close to zero. For $U_0 \ll 8E_r$ and $r \ll \omega_r^2$, we therefore expect losses from the Bloch beamsplitter $P_{loss} = 2\pi\Gamma_3 \propto (\omega_r^2/r)(U_0/8E_r)^6$ where $\Gamma_3 \propto (\omega_r^2/r)(U_0/8E_r)^6$. This scaling of the higher-order losses in the limit of $r \rightarrow 0$ agrees with our numerical simulations.

In addition to the third-order process discussed above, there are an infinite number of these higher-order processes that conserve energy and momentum, but the transition rates are highly suppressed at lower lattice depths. Fig. 3.9b show the result of a simulation with an increased lattice depth, to a regime in which many of these higher-order transitions can couple amplitude to higher-excited states. The parameters chosen for this simulation happen to drive five of these higher-order transitions within the first Bloch period. A ramp rate $r \ll \omega_r^2$ is chosen for the simulation so that the transitions are well-resolved. In contrast, Fig. 3.4a) illustrates negligible higher-order losses because all higher-order transitions are highly suppressed at lower lattice depths.

3.7.5 Crossing through velocity degeneracy

The dynamics while crossing through velocity degeneracy are determined by studying the eigenstates of the DLBO Hamiltonian, Eq. (3.5). An initial momentum state $|n\rangle$, where $n > 0$, can be decomposed as

$$|n\rangle = \frac{1}{\sqrt{2}}(|+n\rangle + |-n\rangle) \quad (3.16)$$

where are the symmetric and antisymmetric combinations of the free-space momentum basis states $|\pm n\rangle$ as derived in Section 3.7.2. Similarly, $|-n\rangle$ can be decomposed as

$$|-n\rangle = \frac{1}{\sqrt{2}}(|+n\rangle - |-n\rangle) \quad (3.17)$$

Without loss of generality, we restrict our attention to one arm of an interferometer with momentum $|n\rangle$. Then when one of the two lattices is initially comoving with the state $|n\rangle$,

this state will be loaded into the ground state of the DLBO Hamiltonian in Eq. (3.5) as a superposition of odd-parity and even-parity ground states according to Eq. (3.16).

Crucially, relative phase shifts between the even- and odd-parity eigenstates cause amplitude to add constructively or destructively for the positive momentum or negative momentum states; for example, if the state $| -_n \rangle$ acquires a π phase shift relative to the state $| +_n \rangle$, then the state $| n \rangle$ in Eq. (3.16) will transform the the state $| -_n \rangle$ in Eq. (3.17). There are two sources of relative phase shifts between the even- and odd-parity states as the lattices are swept through velocity degeneracy. First, since these states are energy eigenstates of the Hamiltonian, there is a dynamical phase difference ϕ_d between the two states given by $\phi_d = (1/\hbar) \int dt' (E_-(t') - E_+(t'))$, where E_{\pm} denotes the energy of the even- and odd-parity ground states over time, as shown in Fig. 3.3c. Since the even- and odd- parity states have different level structure near the degeneracy crossing, this gives a non-trivial phase shift. In addition, there are two additional level crossings for the even state near velocity degeneracy compared to the odd-parity state, as discussed in Sec. 3.7.2. These level crossings correspond to transferring photons to and from the laser field, so the phase of the laser field is imparted to the atomic state during these crossings.

Laser phase is a well known source of phase in atom interferometers, and is the primary phase contribution for certain interferometer configurations such as Mach-Zehnder interferometers [69]. In a single optical lattice, laser phase arises when the position of the laser standing wave shifts position with respect to the atom, resulting in a phase shift $\Delta\phi = 2k\Delta x$. In the case of two optical lattices, there is an additional degree of freedom, namely the relative position of the two lattices. This changes the offset phase ϕ_0 in Eq. (3.1), and it is reasonable to expect this phase term to play a coherent role in the dynamics.

There are two ways to understand the laser phase effects, mathematically and physically. Mathematically, one can see that the even-parity state is shifted relative to the odd-parity state from the definition of the unitary transformation in Eq. (3.3). As mentioned previously, the sign on $d(t)$ in Eq. (3.3) is changed at time $t = 0$, which changes the phase offset on every basis state except for the zero momentum state $|0\rangle$. Just before time $t = 0$, the odd-parity ground state is approximately given by $| -_{gs} \rangle = (e^{-i\phi_0} |1\rangle - e^{-i\phi_0} | -1 \rangle) / \sqrt{2} = e^{-i\phi_0} | -1 \rangle$, whereas after time $t = 0$ the state becomes $| -_{gs} \rangle = (e^{i\phi_0} |1\rangle - e^{i\phi_0} | -1 \rangle) / \sqrt{2} = e^{i\phi_0} | -1 \rangle$. The odd state is therefore phase shifted by $2\phi_0$. Since the state $|0\rangle$ is unchanged, the even- and odd-parity states see a relative phase shift of $2\phi_0$. Physically, the nature of the degeneracy crossing is a result of constructive or destructive interference between amplitudes. Since there are two additional level crossings of the even state compared to the odd state, the even state receives a laser phase shift $\phi_l = 2\phi_0$. At time $t = T_B$, after the two additional crossings, both the even- and odd-parity states are mostly superpositions of the states $| \pm l \rangle$, but the extra phase shift of the even state results in coherent interference and changes the resulting output state. This phase shift can also be observed in our numerical simulations, where the even-parity state is phase shifted by ϕ_0 at each of the two level crossings near velocity degeneracy. Up to a global phase, the new state after the degeneracy crossing can be written as:

$$|\psi_f\rangle = \frac{1}{\sqrt{2}} (e^{i(\phi_d + \phi_l)} |+_n\rangle + |-_n\rangle) \quad (3.18)$$

By controlling the phase shifts ϕ_d and ϕ_l in an experiment, one has control over the output nature of the degeneracy crossing. For example, arranging for $\phi_d + \phi_l = 2m\pi$ for some

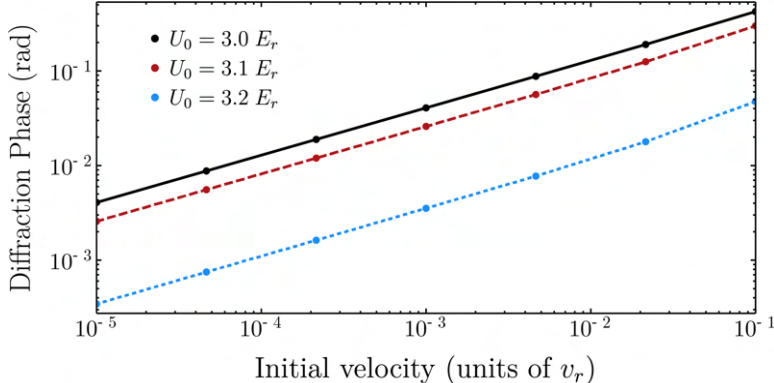


Figure 3.10: Numerical simulation of diffraction phase from a Bloch beamsplitter as a function of velocity with respect to the initial optical lattice. Simulation includes adiabatically loading the lattice, frequency ramping at a rate $r = 1.0\omega_r^2$ for four Bloch periods, then adiabatic unloading of the lattice. See text for further discussion.

integer m ensures that the state after the crossing will be identical to the state before the crossing, which corresponds to transmission through the crossing. For $\phi_d + \phi_l = (2m + 1)\pi$ for some integer m , the output state becomes $-|+n\rangle + |-n\rangle = |-n\rangle$, which has opposite momentum compared to the input state $|n\rangle$ and corresponds to a reflection. Intermediate values of the phase can be used to split amplitude between the two momentum states $|\pm n\rangle$. In practice, it is easiest to change ϕ_0 , and therefore ϕ_l , since this phase is directly controllable experimentally. Our simulations show that ϕ_d also depends on ϕ_0 at the moment that the lattices are velocity degenerate, but this dependence does not prevent one from continuously transforming between different output behaviours by changing only ϕ_0 .

The phase ϕ_d is dependent on the lattice depth, and therefore the lattice depth needs to be well controlled in order to see coherent dynamics after the zero-crossing. In the limit $U_0 = 0$, the dynamical phase ϕ_d is given by $\phi_d = 16\omega_r^2/r$, such that $\phi_d \gg 2\pi$ when $r \ll \omega_r^2$. When $U_0 > 0$, this phase term is also a function of the lattice depth; as a result, fluctuations in U_0 lead to fluctuations in ϕ_d . Similarly, variable U_0 across a finite laser beam leads to a variable ϕ_d across an atom cloud. Both of these effects result in unreliable zero-crossing behaviour at slow ramp rates, and both effects likely explain why we see the largest interferometer contrast for fast ramp rates around $r = 10\omega_r^2$.

3.7.6 Diffraction phase

Here, we consider the diffraction phase acquired from a beamsplitter, which is the phase difference between the positive and negative momentum components of the resulting wavefunction. If the atomic state initially has some free-space velocity with respect to the lattice, the momentum-parity symmetry of the problem is broken and the resulting dynamics will be asymmetric, leading to a diffraction phase.

Figures 3.10 and 3.11 show numerical simulations of the diffraction phase for a $16\hbar k$ Bloch beamsplitter. Almost all of the diffraction phase from the beamsplitter comes from the first $8\hbar k$ momentum splitting near velocity degeneracy; further increasing the momentum




Figure 3.11: Numerical simulation of diffraction phase from a Bloch beamsplitter as a function of lattice depth for two different ramp rates, using an initial velocity with respect to the lattice of $0.001v_r$. Points are from simulation, lines are an interpolation between points to guide the eye. The zero crossings in the diffraction phase allow for one to operate an interferometer at a “magic” lattice depth to suppress sensitivity to diffraction phase. See text for further discussion.

transfer beyond this does not further increase the dynamical phase ϕ_d . The diffraction phase for a beamsplitter scales like the square root of the initial velocity, but the prefactor in front of this scaling can be controlled by varying the lattice depth and the details of loading or unloading the lattice. The simulations in Figures 3.10 and 3.11 use a linear intensity ramp for loading an unloading over a time $t_{load} = 6\pi\omega_r^{-1}$.

Figure 3.11 shows the diffraction phase as a function of the lattice depth, and oscillations in the diffraction phase allow one to operate at a “magic” lattice depth with suppressed sensitivity to diffraction phases from missing the center velocity of the atom cloud. For precision measurement, such magic lattice depths could be used to significantly reduce the diffraction phases caused by fluctuations in experimental parameters. For example, a ramp rate of $r = 4\omega_r^2$ and a lattice depth around $U_0 = 5.9E_r$ gives 80% efficient beamsplitters with minimized diffraction phase sensitivity (see Figures 3.5 and 3.11). We can reasonably operate within $0.001v_r$ of the center velocity of the atom cloud, and by intensity stabilizing the lattice to 1% fluctuations, the diffraction phase can be limited to ± 10 mRad. This diffraction phase can then be measured directly by varying the duration of the interferometer, as done in reference [67].

3.7.7 Higher-order generalization of the dual-lattice methods

The transitions driven in DLBO are two-photon processes that transfer $2\hbar k$ momentum. By sweeping past multiple of these transitions in successions, LMT can be easily achieved. In contrast, higher-order transitions are also possible that transfer $2n\hbar k$ momentum in a single, multi-photon process.

It is instructive to first understand single-lattice higher-order processes before understanding the dual-lattice analogues. SLBO can be thought of as adiabatically sweeping past a successions of $2\hbar k$ Bragg transitions [68]. The higher-order, multi-photon analogue has been implemented experimentally in reference [47]. The laser is adiabatically swept across

a $2n\hbar k$ Bragg resonance, which adiabatically drives a $2n$ -photon process. Though not discussed directly in [47], this process can be interpreted using a Bloch band picture where atoms have an initial quasimomentum outside of the first Brillouin zone such that they are loaded into higher Bloch bands. As the lattice is accelerated, the state sweeps past a level crossing between higher Bloch bands, and successful momentum transfer requires the state to adiabatically traverse the crossing and stay in the same Bloch band.

DLBO can be thought of as adiabatically sweeping past a succession of “double Bragg” transitions [34]. A first-order double Bragg transition symmetrically drives $\pm 2\hbar k$ Bragg resonances such that the two arms are split by $4\hbar k$ momentum. One can also symmetrically drive two higher-order Bragg resonances that transfer $\pm 2n\hbar k$ momentum to obtain a $4n\hbar k$ beamsplitter, as are implemented in references [1, 32].

It is also possible to adiabatically sweep past a higher-order double Bragg transition. In terms of the modulation frequency ω_m in Eq. (3.1), these resonances occur at $\omega_m = (2m+1)\omega_r$ for integers m . A $4n\hbar k$ adiabatic dual-lattice beamsplitter can be achieved by sweeping past one of these resonances adiabatically. An experimental sequence would consist of the following: 1) atoms are adiabatically loaded into a lattice with a modulation frequency slightly below the desired resonance, 2) the modulation frequency is swept across the resonance, and 3) the atoms are adiabatically unloaded from the lattice. It is important that the modulation frequency does not become close to other resonances during this sequence. Unlike a Bloch beamsplitter, continued ramping of ω_m after a high-order beamsplitter process will not transfer more momentum, but rather alternate between increasing and decreasing the momentum splitting between arms. The average momentum transfer per Bloch period will still be $4\hbar k$, as in the ground band.

Our simulations of this process show that it can be more efficient than a Bloch beamsplitter at a given ramp rate. However, there are two major downsides to these higher-order dual-lattice techniques. First, much more laser power is required to drive the transition; the power required to drive an n^{th} -order Bragg transition scales sharply with the order n , namely as n^2 to maintain the same Rabi frequency, and n^4 to also maintain the same single-photon scattering rate [57]. Second, continued ramping of the lattices does not continue to increase momentum splitting in any advantageous way compared to using the ground band. As a result, the first-order dual-lattice methods discussed in the main text are easier to use if the goal is to achieve very large momentum splitting without the need for significantly more laser power.

3.7.8 Application to recoil measurements

A generalization of DLBO shows promise for atom recoil measurements, and therefore in measurements of the fine-structure constant α [67]. This section is included as an example of the potential applications of DLBO, however we note that before such a measurement, many new systematic effects would likely need to be studied.

By removing the assumption that $\omega_1 = \omega_2$ and are independent of time in the Hamiltonian in Eq. (3.1), asymmetric lattice guided geometries can be created [48]. Additional light frequencies can also be added to the laser in order to address more than two velocity classes of atoms at the same time. Figure 3.12 shows an example interferometer configuration that would be sensitive to an atom recoil phase. The phase in the interferometer can be calculated




Figure 3.12: Interferometer geometry sensitive to an atomic recoil phase. The asymmetry between the two upper (lower) trajectories leads to a kinetic recoil phase acquired by the upper (lower) interferometer. The simultaneous conjugate interferometer configuration is used for a differential measurement that cancels gravitational phase to first order, and adds the recoil phases in the upper and lower interferometers. Addressing the four velocity classes of light requires one left-moving frequency and four right-moving frequencies, similar to Fig. 3.2a).

by integrating the energy of the atoms over the various trajectories [99]. Assuming that the time to accelerate atoms is much less than the time between beamsplitter or reflection pulses, the phase of the interferometer is given by

$$\phi = 16\omega_r n_s^2 T \quad (3.19)$$

where ω_r is the recoil frequency of the matter wave, n_s is defined in Fig. 3.12, and T is the time between beamsplitter and reflection pulse in the upper (or lower) interferometer.

The following outlines a set of realistic experimental parameters that could lead to 10^8 radians of recoil phase, an order of magnitude improvement in sensitivity over the leading recoil measurement [67]. Based on the results discussed in Section 3.5, atoms in our apparatus can interact with up to 1000 photons inside an interferometer where contrast can still be observed. Choosing $n_i = 100$, $n_s = 80$ and $n_f = 100$ (defined in Fig. 3.12) requires atoms to interact with 840 photons before closing the interferometers. For the calculation, we use a time of 80 ms between opening the interferometers and slowing the arms back to having the same velocity, the same timing used in reference [67]. Using a frequency ramp rate of $r = 250$ MHz/s, Cesium atoms can be accelerated from $|2n_i\hbar k\rangle$ to $|2(n_s + n_i)\hbar k\rangle$ in roughly 6 ms, which is much less than the time between different pulses. This ramp rate was shown to give good interferometer contrast in the main text for atoms with a vertical velocity spread of $0.05v_r$.

Chapter 4

Experimental Overview

4.1 Introduction

This chapter begins by describing the design decisions made when designing the next-generation α experiment. This is followed by details about the vacuum design as a reference so that future experiments hopefully don't need to go through all of this process again. As a whole, the new experiment is very similar to the previous generation's - the main changes we made were in designing a vacuum chamber to minimize distortions to the laser intensity profile and to design a very large clear aperture in the vacuum chamber. Chenghui Yu and Brian estey's theses [23, 99] already cover of many of the details of the experimental system, and since there are already good references I won't describe the entire system here. Brian's thesis [23] has a more in-depth overview of the apparatus, and Chenghui's thesis [99] has more details about all of the systematic effects as understood at the time of their 2018 publication [67]. In addition to the new large-aperture vacuum system, some new aspects of the experiment covered in this thesis include a compact 2D MOT (Section 4.5.2), a 3D MOT with larger magnetic coils and a new control system. Last I'll discuss some interesting new ideas that either aren't quite publication worthy or haven't been implemented yet.

4.2 Motivation and design decisions

The previous generation of this experiment published their results in 2018 [67] with an error bar of 0.2 ppb in α . Of critical importance to this section is their error table: see Fig. 4.1 for the error table from the 2018 measurement. In order to build an experiment with 10x improved accuracy, we need to make sure the new design addresses every item in the error table including statistical sensitivity and systematic uncertainties. A 10x better measurement requires 10x smaller statistical error bars and 10x improvements on the largest systematic uncertainties. Looking forward, many of the terms in the error table can be addressed with more careful characterization in the future, and don't need to govern the design of the chamber. Systematic effects critical to the new experiment's design are discussed below.

Table 1. Error Budget. For each systematic effect, more discussion can be found in the listed section of the Supplemental Materials.

Effect	Sect.	$\delta\alpha/\alpha$ (ppb)
<i>Laser Frequency</i>	1	-0.24 ± 0.03
<i>Acceleration Gradient</i>	4A	-1.79 ± 0.02
<i>Gouy phase</i>	3	-2.60 ± 0.03
<i>Beam Alignment</i>	5	0.05 ± 0.03
<i>BO Light Shift</i>	6	0 ± 0.002
<i>Density Shift</i>	7	0 ± 0.003
<i>Index of Refraction</i>	8	0 ± 0.03
<i>Speckle Phase Shift</i>	4B	0 ± 0.04
<i>Sagnac Effect</i>	9	0 ± 0.001
<i>Mod. Frequency Wavenumber</i>	10	0 ± 0.001
<i>Thermal Motion of Atoms</i>	11	0 ± 0.08
<i>Non-Gaussian Waveform</i>	13	0 ± 0.03
<i>Parasitic Interferometers</i>	14	0 ± 0.03
<i>Systematic Error</i>		-4.58 ± 0.12
<i>Statistical Error</i>		± 0.16
<i>Electron Mass (19)</i>		± 0.02
<i>Cesium Mass (10,18)</i>		± 0.03
<i>Rydberg Constant (10)</i>		± 0.003
<i>Total Uncertainty in α</i>		± 0.20

Table 4.1: The error table from the 2018 alpha measurement [67].

4.2.1 Beam related effects

The largest error bar of any item in the error table was the 'Thermal Motion of Atoms'. See Chenghui Yu's thesis, section 5.5 for a summary of this analysis [99]. 'Thermal Motion of Atoms' is a lumped term that included uncertainties in initial conditions of the atom ensemble positions and temperatures as well as errors in the Bragg pulse intensities. The value published in the error table came straight from Monte Carlo estimates.

The dominant effect being captured is diffraction phase, which is a phase shift from Bragg diffraction due to the lattice depth and velocity of an atom relative to the lattice. Experimentally, we take data at multiple values of interrogation time T in order to extract the diffraction phase and remove it from the measurement - see Brian Estey's thesis, Fig. 5.8 [23]. This method assumes the diffraction phase is not a function of T , however various effects in our experiment can give the diffraction phase a slight T dependence. The diffraction phase removal method will therefore not work properly and will introduce a systematic shift. The 'Thermal Motion of Atoms' term in the error bar bounds this effect and its dependence on the various parameters mentioned above.

A main takeaway is that this effect is entirely the result of a finite-sized atom cloud with finite temperature inside of a finite-sized laser beam, such that the atoms across the ensemble are seeing different optical intensities and therefore different diffraction phases. The ensemble average across the atom cloud becomes a function of the size and position of the atomic cloud within the beam.

To lower this effect in the future generation of the experiment, the size of the laser beam relative to the size of the atom cloud needs to be increased. Making the atom cloud smaller would reduce signal size, and making the cloud denser could make the atom-atom interactions become a troublesome systematic effect. The best course of action is then to make the laser beam larger. In terms of the beam waist w_0 , the thermal motion of atoms term will scale as $1/w_0^2$ or even $1/w_0^4$, since the beam intensity locally gets flat near the center of a Gaussian beam.

Increasing the waist w_0 of the laser beam will also lower other terms in the error table. 'Gouy Phase' (the largest correction to the measurement) decreases at $1/w_0^2$, and Beam Alignment decreases as $1/w_0$ because the condition for back-fiber coupling becomes more sensitive.

Speckle phase shift is also related to the beam intensity profile, but it describes effects from small-scale intensity inhomogeneities on the beam profile instead of the overall Gaussian beam profile. The speckle phase shift can be addressed by increasing the clear aperture of the experiment to prevent an clipping of the beam, and by having very clean optics/ clean-up optics before the beam enters the chamber so that the beam profile is very nearly Gaussian.

To address all of these issues, the new vacuum system was designed with an 8" clear aperture, compared to the old experiment's 2" clear aperture. This allows for a 2x larger waist laser beam while still having a 2x larger ratio of the beam waist to the chamber clear aperture.¹ We are also investigating using a pinhole filter to mode-clean the laser beam before sending into the chamber, though in a first implementation we found that the

¹I'm not sure it's documented anywhere other than word of mouth. In the old experiment it was believed that reflections of small amounts of light off of the vacuum chamber walls was creating interference patterns on the laser beam inside the chamber.

returning beam reflects off of the pinhole and causes large stray reflections back towards the chamber. Figuring out how to use a pinhole filter before the chamber is a potentially useful addition to the experiment moving forward, although it would likely introduce new systematic effects.

4.2.2 In-vacuum optics and beam delivery

Another important design decision was whether or not to use a single fiber and retro-reflection mirror for generation of the optical lattice, or whether to try using two separate fibers for the upgoing- and downgoing-beams respectively, with no retroreflection mirror. With a single fiber, the analysis of beam-related systematic effects is simpler because you only have a single laser beam propagating and back-fiber coupling ensures good beam alignment. For two beams, the benefit is that you don't have a bunch of parasitic lattices formed inside the chamber due to the retroreflection of three optical frequencies against three other optical frequencies. These extra lattices can interfere with Bragg diffraction and Bloch oscillations when the atoms near the lattice at zero velocity in the lab-frame, or moving lattices formed by additional unwanted interferences. See Section 4.6.3 for further discussion. We chose to use a single beam that is retro-reflected in order to maximize the beam quality and the ability to analyze beam-related systematics.

Having decided on retroreflected beam, we next needed to decide whether or not to include $\lambda/4$ waveplates near the retro-reflection mirror and before the the beam enters the chamber. With these $\lambda/4$ waveplates, the up-going and down-going Bragg frequencies need to be orthogonally polarized before going to the chamber such that the circular polarizations inside the chamber can drive Bragg diffraction. This means you can combine frequencies on a PBS without losing optical power. On the other hand, these waveplate can cause additional etaloning and beam distortions - beam distortions caused inside the chamber are especially hard to characterize because we can't image the beam until it leaves the chamber. Without the quarter waveplates, the beams must have the same linear polarization entering the chamber. This means they need to be combined on a non-polarizing beamsplitter and half of the laser power is sacrificed. See Chenghui's thesis Section 3.4 for a good discussion of the relevant Clebsch-Gordon coefficients involved in driving Bragg/Bloch and Raman transitions [99]. In the end, we stuck to the philosophy of maximizing beam quality, so we left out the quarter waveplates. Not to mention the fact that I don't think it would be possible to procure an 8" zero-order quarter waveplate with the optical quality specs we would need.

There are two main consequences of omitting quarter waveplates. First, as mentioned above, half of the optical power needs to be sacrificed when combining beams on a non-polarizing beamsplitter. Second, having all Bragg frequencies the same linear polarization creates extra parasitic moving lattices compared to the old experiment, where orthogonal polarizations ensured that there were only a few parasitic lattices that were formed. This effect is discussed in detail in Section 4.6.3.

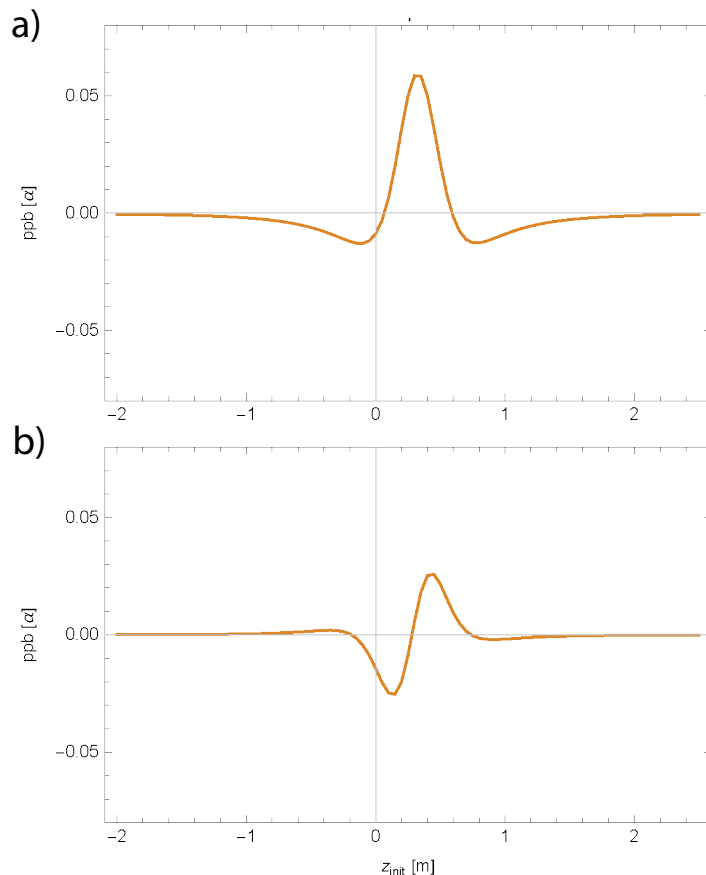


Figure 4.1: The estimated phase shift for an atom interferometer occurring near a ring mass approximately the shape and weight of our vacuum flanges. See text for further discussion.

4.2.3 Acceleration Gradient

The Acceleration Gradient term is not an easy one to deal with looking forward to 10x improvement - it was already a 10x larger correction than the final uncertainty in the previous measurement, so the next-generation measurement is looking at a 100x larger correction than the total error bars on the α measurement. The gravity gradient phase contribution scales with T^3 [99], so we could decrease the interrogation time T to dramatically lower the effect, but then we take a hit in sensitivity of the experiment.

When you start looking at the gravity gradient at the 1% level, you find that it is no longer constant across the interferometry region. Fig. 4.1 shows a theoretical estimate of the effect in α from a ring mass around the chamber, as a function of the initial height of the interferometer inside the chamber. I use an analytical expression for the gravitational potential from a ring mass and numerically integrate the potential energy over the interferometer trajectories to get the phase shift plotted here. A 50 kg ring mass with a radius of 35 cm is used to model the effect from the upper and lower ends of the actual support structure around the chamber. Fig. 4.1a) shows the effect in a normal SCI interferometer,

and b) shows the effect in the OSCI 'BC' port.² I tried to keep this in mind while designing the experiment, but in the end we had to (safely) mount the chamber inside an optics table that had a 1 m hole in it, so there was necessarily going to be support structure mass near the chamber. I probably should have made the support structure with thinner walls or with aluminum, but this would have decreased the stiffness of the structure and then hurt the vibration isolation effectiveness.

In the future of the experiment, the gravity gradient characterization will likely require measuring the gravity gradient as a function of height in the chamber, which will be difficult to do with high sensitivity. The sensitivity of a typical dual-Mach Zehnder measurement comes from long interrogation times T , however this means atoms will travel over a large region of space and average over the small-scale changes that you want to measure.

4.2.4 Statistical sensitivity

In order to get better statistical sensitivity, we need to either 1) get more total phase in the interferometer, or 2) improve our signal-to-noise ratio (SNR) so that we average down to a smaller phase uncertainty. In the end, we will likely need to use some of each strategy.

The old experiment was limited to $n = 5$ Bragg order based on laser power, and $N = 125$ Bloch oscillations based on loss of atoms from single photon scattering. The interrogation time T was limited to $80ms$ because of loss of contrast at long times, as well as limiting the size of the gravity gradient term in the error budget. In this experiment, we have taken a factor of 2 hit in laser power due to not having a quarter waveplate in-vacuum, meaning it will be difficult to get $n = 4$ Bragg order. We have already seen that our contrast decay in T is better than the old experiment, so we can likely get away with more total phase. Since increasing T would dramatically increase the gravity gradient systematic, we will likely want to focus on increasing the Bloch order. $N = 200$ Bloch oscillations will hopefully be reasonable to achieve. In total, $n = 4$ and $N = 200$, and $T = 80ms$ would be a 30% increase in total phase relative to the old experiment. A higher power laser in the future might enable $n = 5$ or $n = 6$ Bragg orders, which would give another boost in total phase. See Section 4.6.3 for further discussion about the limits to the number of Bloch oscillations.

Statistical uncertainty can also be improved by improving the SNR of the experiment. This can be accomplished by a) decreasing the shot-to-shot ellipse noise, such as phase noise, contrast noise, offset noise, etc., or b) improving the ellipse contrast so that ellipse fitting and shot-noise errors are both reduced. If SNR is limited by noise in the detection readout electronics or by atom number shot-noise, there could be easy gains in SNR. For example, we know the old experiment had a large baseline on their detection signal from stray detection light bouncing around the chamber, and we eliminated that in this design. The old experiment believed their ellipse noise was dominated by polarization noise in the final optical fiber to the chamber, so improving the polarization alignment/stability in this fiber could also lead to a reduction in ellipse noise. More time needs to be spent understanding

²See reference [105] for details about the 'OSCI' geometry and how it cancels a constant gravity gradient. In summary, one creates two overlapping SCI interferometers such one output from each of the interferometers is used to create ellipses. The vertical spacing between the SCI interferometers is chosen such that the a constant gravity gradient cancels from the ellipse data, leaving only higher-order gravity gradients

what is limiting the SNR of the experiment, but there are likely some easy-ish wins here moving forward.

Without a dramatic increase in optical power for the interferometer beam, most of the improvements in statistical sensitivity in this experiment will likely need to come from improving the SNR of the experiment, or from increased integration time T .

4.3 The rest of the vacuum chamber design

Having decided on an enormous clear aperture to an already tall atomic fountain, the next step was to flush out the rest of the vacuum design. There were still many other design aspects that need to be sorted out before having a functional atomic fountain. Figures 4.2 and 4.3 show the outside and inside of the vacuum chamber CAD model, respectively. Many of the features discussed in this chapter can be seen in these figures.

4.3.1 Magnetic shielding

An atomic fountain for measuring α requires a uniform bias field in the interferometer region as well as shielding from external fields. This is accomplished using a solenoid inside of a multi-layer magnetic shield - this section concerns the design of the solenoid and the magnetic shield to avoid systematic effects in the α measurement.

The bias field is used to ensure atoms can adiabatically remain in the $m_F = 0$ magnetic sublevel after state preparation is complete. Non-adiabatic losses to other m_F states can cause loss of contrast in the experiment, or even worse could cause systematic phase shifts in the measured ellipse if the other m_F states create parasitic interferometers. The bias field also needs to be uniform over the interferometry region because atoms acquire a second-order Zeeman shift from the bias magnetic field [99]. Effects from uniform Zeeman shifts cancel out in the SCI geometry, but gradients in the magnetic field act asymmetrically on the interferometer arms and cause systematic phase shifts. For similar reasons, the experiment also requires magnetic shielding to ensure the stray magnetic fields from the environment don't shift our measurement, especially stray fields that are changing in time.

The first large design decision was to decide whether to put the magnetic shield and solenoid inside or outside of vacuum. If we had put the shield outside, we likely would have needed a 12" diameter vacuum chamber after including the clear aperture plus in-vacuum baffles (to-be discussed soon). Such a chamber would have 16" diameter flanges, so the solenoid would have needed to be around 17" diameter, and the magnetic shield inner diameter around 18". By comparison, the magnetic shield currently installed inside the vacuum chamber has an inner diameter of 12". When the diameter of a magnetic shield is increased, fringing effects from external fields penetrate deeper into the shielded region. This decreases the shielding factor and also allows magnetic field gradients to impinge on the interferometry region.

More quantitatively, the magnetic field decays exponentially near the ends of a cylindrical magnetic shield at a length scale proportional to the radius of the shield [51]. The field generally peaks at the entrance to the shield at a value larger than the background field, since magnetic field lines are being sucked to that region. The field then decays roughly




Figure 4.2: An overview of the entire vacuum chamber, not including the support structure. The lower chamber is used for trapping atoms in a 3D MOT, where they are launched vertically using a molasses launch. Interferometry is performed in the interferometry region. On the way down atoms pass through a light sheet in the detection chamber where fluorescence is detected. The retro-reflection mirror is mounted to bellows in the top chamber so that closed loop piezo micrometers can control the angle of the mirror.




Figure 4.3: CAD model cross-section of the vacuum system highlighting some features relevant to this section. 1) Bottom viewport attached to the MOT chamber where the interferometry beam enters. The viewport is angled at 3-4 degrees to prevent etaloning inside the chamber. 2) Black copper-oxide tubes in the detection chamber to prevent stray light from reaching the detection photodiode. 3) The main interferometry region - black copper oxide baffles inside of an aluminum tube with the solenoid wrapped around it, inside of a three-layer magnetic shield. 4) The support structure (yellow) that can sit on Minus-K vibration isolation stages (black). The structure is attached to the chamber by six hexapod struts (light blue). 5) The gold retro-reflection mirror with black copper oxide ring around it. The mirror is attached to a bellows so that the angle can be controlled by closed-loop piezo micrometers outside of vacuum.




Figure 4.4: Effect of Bloch oscillation order on the shifts in alpha from magnetic field fringing at the ends of the solenoid/magnetic shield. Larger Bloch orders leads to a smaller relative shift in alpha from Zeeman energy, since the total phase is getting larger but the Zeeman shift stays similar. However, the atoms are also accelerated towards the edge of the solenoids which causes the region without fringing effects to narrow.

according to $e^{-k_L z/R}$ inside the tube, where R is the radius of the tube and k_L 2.405 is a constant derived from theory. For small aspect ratio magnetic shields like ours, the shielding factor is determined by when the exponential decay in field from each end of the tube meet. A larger aspect ratio tube would therefore have decreased the shielding factor on-axis.

Similarly, increasing the diameter of the solenoid also increases fringing effects on the bias field and limits the usable interferometry region. We need a solenoid inside of the magnetic shield for creating a bias field, so fringing effects from the solenoid can also cause magnetic field gradients that penetrate into the interferometry region.

The inner shield layer in our current experimental shield is 92" in length. In order to get the external fields shielded to 1 mGauss or better, for a shield radius of 6" you would end up with a usable interferometry region of 1.5m. This length would limit you to around $N = 300$ Bloch oscillations and $T = 150$ ms, as a rough estimate. If you moved the magnetic shield outside of vacuum and had a shield radius of 9", you would limit the usable region to 1.1m. This would limit the total phase ultimately achievable in the experiment, since the space needed for interferometry scales with the interrogation time T and the number of Bloch oscillations N .

The above outlines the rough line of thinking that went into the design. Fig. 4.4 shows a quantitative treatment of the estimated shift in α as a function of initial height of the interferometer in the chamber for the field from just the bias solenoid. Even in a solenoid with strong fringing effects and $N = 800$ Bloch oscillations, there is plenty of vertical space in which interferometry can be performed. This differs from the previous estimate of 1.5m total usable interferometry region because even though the ends of the tube have strong fringing effects, the atoms spend the majority of the interferometer time in the well-shielded regions so these fringing effects end up being smaller contribution.

The conclusion from Fig. 4.4 is that the shielded interferometer region is likely longer than we will ever need in this generation of the experiment. The benefit of this is that the

interferometer can be operated well-inside the shielded region to be sure we never have to worry about fringing effects (assuming you can launch the atoms high enough...). Note that this systematic analysis does not include non-adiabatic losses to other m_F states due to magnetic field gradients - these would be an additional consideration that could further limit the usable interferometry region.

In total, by having a magnetic shield and bias solenoid outside of the chamber, we would have reduced the usable interferometer region. This seemed like an unacceptable limit to future experiments when we didn't know what the future experiments would look like. So, we put the magnetic shield and solenoid inside vacuum, as was done in the previous generation experiment. The vacuum chamber diameter was bumped out to 16" and all of a sudden the vacuum chamber needed to be enormous. Everything needed to be supersized, viewports are far from the atoms, we need 200A, 16" diameter MOT coils, etc.

Also, note that there are ways to minimize fringing effects, if it ever becomes an issue. Even if you have a poor shielding factor on-axis, you can use coils near the entrance to both ends of the shielding tube to smooth out the field at the ends of the tube [93]. In this reference, they use the trim coils to null the external field, but I'm sure a similar idea will work in our system to smooth the transition between background Earth's field and the bias field inside the magnetic shield.

Note also that magnetic-field systematic effects are relatively easy to characterize because the phase shifts scale with B^2 . If you typically operate with a 300 mGauss background field, you can bump the background field up to 1 Gauss to make the effects 11 times larger, for example. The old experiment went up to 3 Gauss to make the effects 100 times larger. This allows you to quickly bound and/or measure and systematic shifts.

The mumetal magnetic shield installed in the experiment was designed by Magnetic Shield Corporation (I highly recommend this company, they did a great job). The shield is three-layers diameters 14", 13", and 12", and lengths 104", 98", and 92" respectively. This leaves the layers offset by 3" on each end, for each shield, to improve the shielding factor. The shield layers are spaced by permanently welded mumetal, and assembly of the shield required aligning holes in the three layers at the location of the spacers so that 1/4-20 bolts can hold everything together. There are small endcaps used on the middle and outer layers of the shield to help with the shielding factor. I realized later in the design stages that the endcaps would have caused a large region of trapped air with how the final assembly came together, so I had holes laser cut into the larger endcaps to allow for outgassing. In total, Magnetic Shield Corporation estimated a shielding factor of 1000 for the assembly.

It's possible to use atoms to measure the magnetic field inside the chamber, the easiest way being tracking the frequency of magnetically sensitive microwave transitions. We haven't yet characterized the field homogeneity in the interferometry region. One worry in the design process was whether the magnetic field from the 3D MOT coils would saturate the mumetal in the magnetic shield. If this does end up being a problem, we wrapped the magnetic shield in a couple winds of magnet wire so that the shield could be degaussed someday. These wires are connected to the electrical feedthrough at the top of the chamber. WARNING - the degaussing wire is shorted to the chamber somewhere, but only in one spot. You should be OK degaussing as long as you use a floating ground power supply. But don't just plug a variac into the chamber and go nuts, you could really hurt yourself.



Figure 4.5: The 'rotisserie' that was built to wind the solenoid on a 10ft aluminum tube. The tube was mounted on each end to bearings, and the right side was attached to a drive belt and a car window motor. The UHV kapton wire was mounted to a track that could slide along as the coil was wound. Imaged here was our test winding before winding the real-deal solenoid, so we didn't need to maintain UHV cleanliness.

4.3.2 Solenoid winding

Typically, bias solenoids like ours are wound wire-to-wire so that the winding is guaranteed to be uniform. Since we had to hand-wind a 10ft long, 11" diameter solenoid, we would have needed over 2km over UHV-clean kapton-coated wire for this which was unacceptable expensive. Additionally, a wire-to-wire winding of this scale would have the potential to create very large regions of trapped air that could have looked like virtual leaks and spoiled our vacuum pressure.

We therefore decided to wind a solenoid with a spacing between turns of 1cm. We then used UHV safe epoxy (EPO-TEK 353 ND) to fix the wires in place, since the position of the winding turns wasn't fixed otherwise. The aluminum tube was sourced from Kurt Lesker (version one they sent us was out of spec in its diameter tolerance, version two was still out of spec but I milled down the copper baffles to make up for it). A delrin winding spacer was manufactured to guide the wires and ensure a uniform spacing as we wound the wires. Delrin was chosen because it is a relatively low outgassing plastic, in case the procedure was leaving plastic contamination, but soft enough to not scrape off kapton coating from the wire.

Fig. 4.5 shows the 'rotisserie' as we called it - the tube was fixed to bearings on each end, and on one of the ends we rigged a car window motor to a drive belt so that we could get a very geared-down rotation. We practiced the procedure on the first out-of-spec tube from Kurt Lesker, then did the final run on the second tube that was still out of spec but good enough. After winding, we epoxied each turn of the wire in four spots, for a total of around 1000 epoxy spots. We then epoxied some guard rails over the wires to help protect the winding during the installation step of lowering this assembly into the magnetic shield,

since the magnetic shield had screws sticking towards the coil turns. This procedure ended up working very well.

Figure 4.6 shows some estimates of the effects of an imperfect coil winding procedure. To model imperfections, I calculated the resulting magnetic field if each turn had some Gaussian distributed error with a standard deviation of e.g. $500\ \mu\text{m}$ or equivalently 5% error in spacing each turn. Based on the tolerances of our winding guide and observations while winding, I think the average error was likely much smaller than this, and that the only source of error we would need to worry about would be from having accidentally bumped wires during the installation procedure.

I'll make a note here for future generations of the experiment. Instead of putting the magnetic shield and solenoid in-vacuum, it would way easier to leave them outside of vacuum, but then make the tube much longer in order to maintain a good aspect ratio of the shield and solenoid. Atoms would need to be launched higher into the tube to avoid any fringing effects.³ The current design already pushed the limits of what we could fit inside of the $5m$ tall pit room. If you find a very tall lab space in the future, this could be an option for a more scalable design, e.g. see [93] for what you can do with a huge room, a lot of money and a lot of grad students.

4.3.3 In-vacuum baffles for stray light

Another cool feature of the new vacuum chamber is the in-vacuum baffles - let me first motivate them. Essentially everything gets better in an atom interferometer when you make the laser beam waist larger, assuming the same lattice depth. The atoms will see a more uniform intensity, with smaller gradients of the light shift etc. This is why it was so surprising that the old experiment saw a less contrast decay in their experiment when they added an apodizing filter that roughly cut the beam waist in half. The leading hypothesis for why this improved the contrast is that the apodizing filter was cutting the beam intensity outside of a beam waist $\pm w_0$. There is actually more optical power than one would expect on the wings of a laser beam from an optical fiber because roughly 1% of the beam's optical power lives outside of the step-index of the fiber (see Chengui's thesis Section 4.2 [99]). They expected that this power was reflecting off of vacuum chamber walls, then re-interfering with the main beam to cause intensity ripples and loss of contrast. Recall that interference effects between beams scales with the electric field amplitude, not the intensity, so $O(0.01\%)$ optical power can still give you $O(1\%)$ interference ripples on the main laser beam.

This scattering issue was directly addressed by adding in-vacuum baffles to capture any stray light, in addition to the already much larger clear aperture. The baffles are designed so that any stray light coming from the bottom viewport that will hit a baffle before having the opportunity to re-interfere with the main laser beam. After much research on how to make a quality black surface that's UHV safe for relatively cheap, we chose to follow a neat paper and oxidize copper ourselves in order to make a 98% absorptive black coating [64]. A common alternative that is a great option for small parts/big budgets is to use Magic Black

³We likely can't launch too much faster with a molasses launch as is currently used in the experiment, but Section 4.6 discusses what is needed for a lattice launch of the atoms. A lattice launch is much more scalable to higher launch speeds.




Figure 4.6: Simulations of the effects of imperfect solenoid winding on the measured value of α . In each plot, I use a model for our solenoid (one winding per centimeter), and randomly vary the location of each winding. The default unless otherwise stated is 5% standard error, which is $500\ \mu\text{m}$. The x-axis is sweeping the starting location of the interferometer, so that the total interferometer integrates over a different region of the B-field. a) The effect scales roughly linearly with the average position error of the winding. b) The effect scales quadratically with the B-field since it's a second-order Zeeman shift proportional to the B-field squared. c) The OSCI interferometer geometry [105] only cancels out systematics from uniform gradients in the magnetic field, but effects from coil winding imperfections are far from uniform over the interferometer region.



Figure 4.7: Images of the copper oxide procedure and final parts. Upper left shows the final assembly of the baffles structure. The right image shows this assembly being slid into the solenoid tube. The lower left shows the struts of the baffle structure mid-oxidation - the light colored ones were just added to the solution and hadn't yet formed an oxide layer.

coating by Acktar. It has even better absorption and is UHV safe, but the price scales with surface area and was therefore not a reasonable solution for us.

We followed the exact procedure outlined in [64], except we had to use much stronger concentrations of the Citranox (10%) and much stronger concentrations of the blackening solutions (20%) in order to get consistently good results. The Citranox step is very important to remove all oxidized copper from the surface, and the parts must be immediately placed in the blackening solution after rinsing off with water from a squirt bottle. We tried using sandpaper to roughen the surface, but the resulting oxide coatings looked terrible and very inhomogenous. In the end, we sand-blasted all copper parts in the physics machine shop, which gave reliable uniform finishes when sand-blasted well.

You can see some pictures of the process and resulting parts in Fig. 4.7. The lower left image shows double-boiling of the oxidizing solution in a larger water bath. A stir rod helped ensure uniform temperature. The image shows the struts, part of the baffles, in the early stages of oxidation just after putting the parts into the solution. The struts that are still copper colored were put in the solution last, and the blackest ones were put in first. The upper left image shows the final assembly of the baffles. The right image shows the baffles as they were being lowered into the solenoid tube with our home-built crane visible up above.

In addition to baffles, we also added oxidized copper tubes in the detection region to prevent stray light from reaching the detection photodiode. The old experiment had a voltage baseline in the detection photodiode on the order of 0.5V from the detection light scattering and reaching the detection photodiode.⁴ This was most likely from non AR-coated

⁴If you looked inside of their detection chamber with an IR viewer during the experimental sequence, it




Figure 4.8: An image taken looking up into the chamber from the bottom viewport. The tubes used to block stray light in the detection region are visible, as well as the beam shutter which is moved off to one side. The baffles make the main tube look black, and the gold mirror at the top of the chamber reflects the camera flash. You can also see the bottom end-cap on the magnetic shield, which maintains the 8" clear aperture, and the degaussing wires on the left and right side of the magnetic shield which are UHV-kapton-taped in place.

windows causing extra reflections, or from dust on the viewports causing diffusive scatter.

In an effort to reduce detection noise, I added these tubes and designed them to be as effective as possible without affecting the clear aperture of the interferometer beam or optical access of the detection optics. An image of these tubes in the final assembly is shown in Fig 4.8, which is taken looking up from the bottom viewport. You can see a seam from the rolled copper sheet that was used to make to the tubes, as well as some imperfections on the oxide surface. The inner surface of the tube had a good oxide layer so I said this was 'good enough', since it was difficult to uniformly oxidize parts this large. There is also a ring of black oxide mounted around the gold retroreflection mirror at the top of the chamber in order to avoid any stray reflections off of the mirror mount.

We did UHV tests of oxidized copper samples and saw no outgassing at all from our test chamber. Before doing the UHV tests, the test chamber was hydrogen baked at 300-350C to get extremely low outgassing in the test chamber walls. The turbo pump we were using for the UHV tests was extremely bad⁵, so we were only able to get to 10^{-9} torr after a day or two of water baking in the test chamber, but it still gave a very clear indication of

lit up like a light-bulb.

⁵I estimated less than 1 L/s pumping speed

any outgassing that was much larger than the baseline hydrogen outgassing from the test chamber walls. Also, before beginning with the copper oxide project I discussed outgassing issues with researchers at NIST in Boulder, and the same procedure had been used to make copper oxide parts in experiments with vacuum pressure into the 10^{-12} Torr range so I didn't anticipate any outgassing issues.

In retrospect, the copper oxidizing project ended up taking quite a bit of time to get the procedures ironed out, but it worked and was a neat home-made addition to the experiment.

4.3.4 Hexapod, support structure and vibration isolation

Another neat upgrade of this generation of the experiment is that the chamber is mounted to a hexapod - this allows a bold experimentalist (or a PI after having a few beers) to move the orientation of the chamber relative to gravity. Since all optics are mounted directly to the chamber, this alignment procedure decouples the gravity alignment from the rest of the optical alignments. The hexapod is the brainchild of Joseph Silber, a mechanical engineer at LBL who pitched this in his first 30 minute meeting with us about the α project in Holger's office. The hexapod struts themselves were Joe's design from another project DESI⁶ where they were used to mount their enormous ≈ 1 m lenses. The only change we made to the strut design was changing the threads to ACME threads and making the motion differential between the two threads for much finer alignment control.

We only used the hexapod struts for gravity alignment a few times before my graduation, since a lot of the time we had a car-jack supporting an extension to the interferometry breadboard, and moving the hexapod would push the breadboard against the carjack and add stresses to the bottom viewport.⁷ In the future, I expect the hexapod will be very useful for fine alignment of the launch during data taking, but you'll need to work out a readout scheme that is more sensitive to the alignment since we never found a great procedure during my time.

Joe Silber also helped with the design of the support structure for the chamber. The support structure was designed to allow the entire experiment to be vibration isolated. For our measurement of α , ambient vibrations don't affect the experiment other than giving a random phase that walks data points around the ellipse - see Brian's thesis [23] for details about this. However, the initial idea for this new experiment was that we might potentially want to make a juggling interferometer as a search for oscillating forces from dark matter candidates. A single juggling interferometer would need vibration isolation so that the juggling interferometer phase isn't washed out. Although I think this particular search for dark matter isn't very compelling or sensitive, having an experiment with the versatility of being able to measure gravity and gravity-related phases without a differential measurement opens a lot of possible doors in the future.

Joe and I designed the support structure to be 1) extremely rigid, so that we could vibration isolate the chamber at the 0.5Hz level and keep all mechanical modes much much higher than this, and 2) very safe, since I didn't trust myself to design something safe enough to hold large weight of the chamber. To ensure rigidity, Joe made a finite element analysis

⁶Dark Energy Spectroscopic Instrument

⁷We actually did do this before realizing what a bad idea it was, but didn't see the vacuum pressure move...

(FEA) simulation of the support structure with a simplified chamber mass, and found the vibrational eigenmodes of the chamber - see Fig. 4.9.

The structure is designed to sit on top of four CM-1 vibration isolation stages from Minus K Technology. We used a Minus K stage in the portable gravimeter project in Holger's group [96] and the stages have as good of performance as you can hope for from a simple passive commercial system. The CM-1 stages can have a resonance frequency of 0.5 Hz or even lower. Stiffening the support structure and pushing all mechanical modes well above 0.5 Hz ensures minimal coupling between modes so that the vibration isolation can be as effective as possible.

In the future, if anyone decides to try floating the chamber on these stages and wants to optimize vibration isolation performance, you're going to need to eliminate any lower frequency mechanical modes on the chamber. Primarily these will be from cables - cables can 1) couple to the minus k resonance, and 2) bring vibrations from outside the chamber to the chamber. All cables should be strapped to the table (ideally multiple times) with sorbathane in order to prevent vibrations from reaching the chamber. Then on the chamber side, cables will need to be fixed tightly to the chamber for the entire routing of the cable, eliminating any low-frequency motional modes of the cables. This would probably work well for optical fibers and BNCs, but the cooling lines for the MOT coils will be much hard to isolate - I'm not sure how you would be able to vibration isolate this, if you could at all, since the hoses are rigid and the water carries vibrations. When I realized I had been very overambitious in my chamber design, I quickly gave up on my prospects of floating the chamber and making any sort of gravity-sensitive measurement during my time as a graduate student. As a result, there was no effort put into vibration isolating cables. But if you end up trying it, good luck! For best performance, close the door to the pit room and block the air flow from the AC vent. The portable gravimeter experiment from our lab found that stray air currents were limiting the ultimate sensitivity of the gravity measurements [96]. You'll also somehow have to deal with the fact the switching on and off the MOT coils creates audible vibrations that would absolutely destroy any vibration-sensitive measurement...

4.3.5 In-vacuum shutter for beam imaging

Wouldn't it be nice to stick a CCD camera inside the vacuum chamber and take images of the laser beam? Well of course that's not possible in a UHV chamber. Instead I worked out a scheme to have a ceramic disk shutter that can be lowered in front of the beam. The shutter can be seen in Fig. 4.8 in the 'un-shuttered' position. The idea is that the shutter can be lowered in front of the beam, and the laser beam can be imaged by its diffusive scatter off of the shutter from underneath the chamber. We tested a few options for white UHV-safe materials and a CNC'd macor disk ended up giving the most uniform surface finish/ best image for a laser beam profile. From our images, we could see something like $50\ \mu\text{m}$ grain sizes on the surface, so I wouldn't trust this method at that length scale, but for larger-scale effects on the beam you might find this to be quite useful. To decouple macor surface roughness from actual beam noise, you can hopefully just move the disc around and take images on different parts of the macor disk. The rotation stage has a magnetic feedthrough that will slip if there is too much torque, so I had to take care to minimize the torque on the stage. **Be very careful if you ever move this stage around because you can bang**

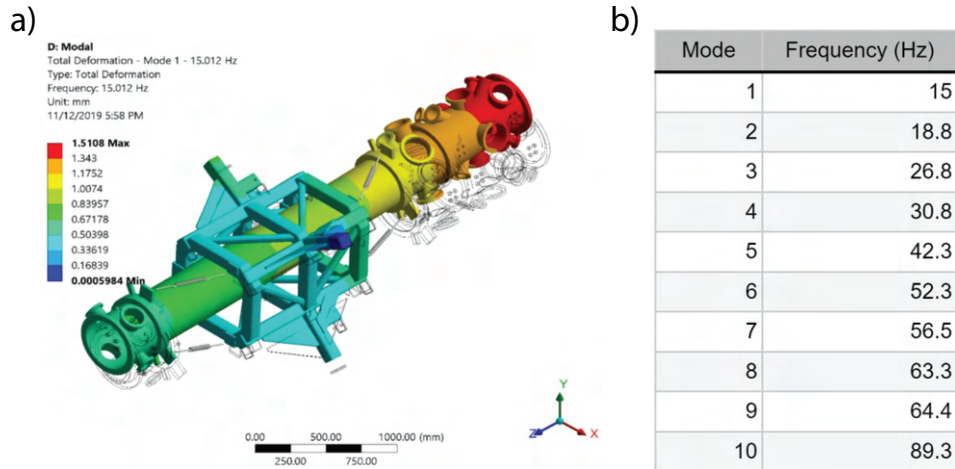


Figure 4.9: FEA simulation of vibration modes of our vacuum system and support structure. a) An illustration of the displacement in the lowest vibrational mode at 15 Hz. The motion is related to bending of the arms of the support structure allowing the vacuum chamber to swing sideways. b) A list of the first 10 vibration modes. Although the first mode is primarily horizontal motion, the second mode (not shown) is primarily vertical motion, which will be most relevant if trying to vibration isolation the system.

the shutter against the detection copper oxide tubes, and probably you’ll create dust in the chamber. There is a small set screw that you release to allow the rotation stage to freely move.

Note that another possible way of imaging the beam in-vacuum is by using the atoms themselves - this idea is discussed in detail in Section 4.6.2.

4.3.6 Designing for UHV

Our vacuum pressure needed to be low enough to ensure a μs vacuum lifetime - this would ensure that our MOT size isn’t limited by vacuum pressure, and that background gas collisions isn’t limiting the contrast of the interferometer. For some numbers here, the cavity interferometer in our group has recently measured the atom lifetime in their optical lattice as a function of vacuum pressure, see Fig. 7.6 in Vicky Xu’s thesis [97]. In our current experiment, we can already see that MOT size goes down when the temperature of our 2D MOT region goes up, which indicates that our MOT loading is still limited by vacuum pressure, even though the pressure is around 1×10^{-10} torr.

In the end, we aimed for a vacuum pressure around 1×10^{-10} torr. There were many things that had to be considered in order to get this large of a vacuum system down to this low of vacuum. For most AMO experiments, there is typically some amount of room for error in a UHV system because it is relatively easy to break vacuum, replace parts, and pump down again. For this chamber, we needed it installed by a professional rigging team and it took \$70k and most of a week, not to mention all of the planning that went into the assembly steps. There were so many parts in-vacuum that vacuum pressure issues likely would have been extremely hard to isolate. Failure to get vacuum would have required a week for the

riggers to disassemble (another \$70k), an unknown but potentially very large amount of time to find the issue, another week to re-assemble (and another \$70k), then another month to set up the bake and bake the chamber. We therefore spent a lot of time making sure we got it right the first time.

How does one get this much hardware into vacuum without extra out-gassing? The first step is to choose UHV-safe materials for everything going in-vacuum. I have described some of this in the sections above, but in general I did a UHV test in our UHV test chamber for any material or procedure that I was unsure about. Beyond material choices, I was mostly worried about trapped gas. We had roughly 100 screws inside of the vacuum chamber, and it would only take one screw with a slow leak to wreck our pressure. You can get vented screws from e.g. McMaster-Carr, but I was still paranoid about trapped air in the threads. In the end I ordered some thread-cut screws which had a milled cut along the entire thread. This allows each thread to individually outgas quickly. We also used split washers for a similar reason. Evidently this plan worked because we got good vacuum pressure.

Another consideration was how to properly UHV clean parts. For small parts we used a standard sonication procedure adapted from Dan Stamper-Kurn's group - 15 minutes of sonication in soapy water (Simple Green and DI water), 15 minutes in acetone, then 15 minutes in methanol. For the 3x magnetic shield layers and the aluminum tube that we wrapped the solenoid around, we needed a procedure for cleaning the tubes without using sonication because the parts were 10 feet long. The strategy that we used was to 1) try to get the large parts delivered to us as cleanly as possible to begin with, then 2) clean the parts by hand with acetone and then ethanol. Kurt Lesker manufactured the aluminum tube that we wound the solenoid around, and shipped it UHV clean. I don't know exactly how Lesker did the UHV cleaning of the aluminum tube on their end. For the magnetic shield layers, I had them only handle the shield with latex gloves after annealing at very high temperature (at 1100C or something similar). This ensured the shield was relatively clean when we got it, since any oils or contaminants would have been burned away during annealing. They then shipped the shield layers disassembled so that we could clean them individually before re-assembling.

For all of these large tubes, we used clean-room wipes to hand-scrub every surface with acetone and then ethanol. For the insides of the tube, we used a clean-room mop to reach the center in order to minimize dust or oil contaminants. Ethanol was used instead of methanol because we had to be in a room with these fumes for days.⁸ The cleaning was done in a second floor Birge lab.⁹ The parts were supported by 8020 struts covered in UHV foil. We tested the cleaning procedure on smaller parts in the UHV test chamber and didn't see any outgassing - and it ended up working in the end. The only thing I would have changed is that the cleanroom wipes we used very slightly dissolved in acetone, which was leaving streaks on the parts. Even if we did leave plastic residue on the parts, I guess the layer was so thin that it baked out very quickly in the final bake of the chamber, as it didn't end up being an issue.

Regarding how to design a pumping scheme for the chamber, AMO experiments have

⁸Whoops, we had to buy 5 gallons of 100% pure ethanol...

⁹I spent a week cleaning the room many times to get rid of as much dust as possible, then we left a window open and got ashes from wildfires all over the room so I had to clean the room all over again. I'm sure some of these ashes made their way into the actual vacuum chamber...




Figure 4.10: If your reaction to this is ‘No way Zack, this vacuum chamber could ruin me’, then you have the correct reaction. Also, I have heard that the rapid change in air pressure from sudden loss in vacuum pressure can cause hearing loss if you’re too close, so in addition to emotional damage you could also potentially get physically hurt.

historically used large ion pumps, but in recent years non-evaporable getter (NEG) pumps have gotten very popular. NEG pumps are made from a specially engineered ceramic that functions like a fancy piece of charcoal, absorbing background gas until the ceramic saturates. Ion pumps are more effective with heavier atomic species, but very low pressure UHV experiments are typically limited by hydrogen outgassing from the stainless steel chamber walls. NEG pumps are most effective with lower mass atoms, and therefore have a very large hydrogen pumping speed. A combination of a large NEG pump and a smaller ion pump works very well to keep hydrogen pressure down while still pumping out the smaller fraction of heavier atoms. Titanium sublimation pumps are also an excellent option for pumping hydrogen if you can afford a large surface area for depositing titanium and if you aren’t too worried about titanium contaminating other parts of your chamber. I opted to use all three of these types of pumps to give us the best chance of getting good vacuum.

Two large ion pumps were installed on the chamber to make sure we had good pumping for heavier atomic species (e.g. Cesium). In the end, we’re only using the 75 L/s ion pump, we never turned on the 55 L/s one because I never got a controller for it. Don’t turn on the 55 L/s one now unless you have a turbo attached after breaking vacuum, since there will be a ton of outgassing initially and it could saturate some of the other pumps.

Our main hydrogen pumping was designed to come from titanium sublimation pumps (TSP) deposited onto chamber walls, and from Capacitorr NEG pumps. In reality, I think almost all of the pumping currently in the experiment is currently from the TSPs: when I activated the NEG pump after already activating the TSPs, there was essentially no change in pressure in the chamber. The surface area numbers in Table 4.11 are a lower bound of the surface area of titanium that would be coated. To activate the TSPs, I ran the current through the filaments up to 50A or even higher, even though the manual said never to exceed 47A. This was based on watching the RGA reading in real-time, and I kept increasing the current until I could see the pressure start dropping after each deposition. Since I exceeded

the manual current limits by so much, it's very likely that I deposited a ton of titanium and that's why we have such good vacuum pressure. I only activated filament 3 on both TSPs - you should use these filaments until they run out, since activating a new filament will initially cause a ton of outgassing which could saturate the other TSP region/ the NEG pumps.

The Capacitorr HV 200 was included in case we needed larger volume, higher vacuum pumping while activating the various pumps. In the end, we didn't really use it at all because it only raised the chamber pressure. The pump needs to be heated in operation which causes extra outgassing, hence why it's only an HV pump and not UHV. It ended up being sufficient to do all of the activation of pumps while the turbo pump was still connected, then I think I fired the TSPs once more after closing the gate valve to the turbo. As I mentioned, the Capacitorr D2000 didn't seem to lower the pressure at all after I activated it (multiple times). I'm not sure exactly why that's the case, but just letting you know in case you're baking down again. I don't think this should be interpreted that Capacitorr pumps aren't effective, but rather that our TSP pumping was both extremely effective.

Vacuum conductance also needs to be considered when designing a chamber like this. I calculated the conductance of the constrictions before every pump to ensure the conductance was larger than the pumping rate of the pump, otherwise the pump would be choked and not operate effectively. Note that estimated conductance from the center of the shielded region to the pumps is on the order of several hundred L/s. Since our pumping rate is much higher than this, it's possible that the vacuum pressure inside the shielded region is much worse than the readings on the ion gauges. This couldn't really be avoided though, since there's no way we would have made the chamber any larger diameter. In any case we now have a much larger vacuum conductance in the interferometry region than the old alpha experiment did, so it's still a large improvement in vacuum conductance compared to the previous generation.

Table 4.11 shows an estimate of the final vacuum pressure we might achieve with this chamber design. I never found good literature values for outgassing of mu-metal, but MuMetal manufacturers do say that the outgassing is similar to stainless steel. Looking at outgassing numbers for different materials, it was clear that the large majority of our outgassing would be from hydrogen from the enormous surface areas of the stainless steel chamber walls and the magnetic shield layers. This is typical for well-designed UHV chambers that haven't done a hydrogen bake of the chamber walls. Even if some of the smaller parts in vacuum had 100x larger outgassing rates per unit surface area, the surface area of stainless steel and mumetal was so large as to remain the dominant outgassing source. So, assuming that we used UHV safe materials for all components in-vacuum, didn't have any virtual leaks, UHV cleaned everything adequately, and that we did a long water bake around 150-170C, Table 4.11 should be a reasonable estimate of what's dominating the background pressure inside the chamber. I made this estimate during design phase, and it ended up being almost exactly the final pressure that we achieved. The pressure after baking was initially 6×10^{-11} Torr, and it has now drifted up to 1.7×10^{-10} Torr. The estimated pressure value was also consistent with a naive scaling between the old experiment and ours. The ion gauge sensitivity to hydrogen is only 0.5 though, so the actual pressure in the chamber is likely 2x higher than this reading assuming the pressure is dominated by hydrogen.

Regarding Cesium background pressure, the old experiment used a differential pumping

np.linalg.norm	1.6	3.2%
np.ones	1.5	3.0%
np.outer	1.6	3.2%

Outgassing (Torr * L/s)	
Chamber walls	1.20E-07
Magnetic shield	2.44E-07
Total	3.64E-07
Hydrogen pumping speed (L/s)	
Ion pump 55 L/s	20
Ion pump 75 L/s	20
TSP	1205
Capacitorr HV 200	200
Capacitorr D 2000	2000
Total	3450
Estimated pressure (Torr)	1.06E-10

Figure 4.11: An estimate of the final vacuum pressure achievable with our experiment. The surface area of chamber walls and the magnetic shield were assumed to out-gas at the same rate of $3 \times 10^{-12} \text{Ls}^{-1} \text{cm}^{-2}$, which was a value for non-hydrogen-baked stainless steel from literature. See e.g. [12] but you will find similar numbers in other sources as well. The hydrogen pumping speed is listed for the various pumps installed on the experiment. This pressure estimate met our design goal for pressure, and also ended up matching the actual pressure achieved in the experiment.

tube between the 2D MOT chamber and the main vacuum chamber which cut the 10^{-6}Torr ambient background pressure of Cesium by a factor of 1000 (see Brian's thesis [23]). Since we're aiming for a vacuum pressure of 10^{-10}Torr , this seemed insufficient. I therefore used the same design¹⁰ as the old experiment but then added two apertures after the differential pumping tube to further limit the conductance between the 2D MOT chamber and the main chamber. These were 1/4" holes in OFHC copper pieces, which were attached to a copper vacuum feed-through. This could optionally be operated as a cold-finger in the future, since cold surfaces act as a pump for alkali atoms.¹¹ I never calculated exactly how much the additional apertures would limit conductance by, but it's likely more than a factor of 10. The 1/4" diameter holes were made to be large enough that the solid angle of atoms leaving the differential pumping tube would still pass through both holes and our 2D MOT atom flux would be unaffected. The solid angle is easily calculated from the thermal velocities of atoms 300K temperature on-axis and Doppler-temperature off-axis, giving about 1 mrad angle.

There are three ion gauges in various parts of the experiment, in case some break or in case you want to know the pressure in different regions of the chamber. The ion gauges need to be degassed every once in a while if you want an accurate measurement, since they are hot filaments that attract crap inside the chamber. During initial activation of the pumps

¹⁰Actually I took the old 2D MOT off of their chamber and put it on the new chamber

¹¹I had tolerancing issues in the final assembly of this cold-finger, and in a rush I ended up adding a stainless steel washer between the cold-finger and the copper assembly with the differential pumping holes. As a result, the thermal conductance of the cold finger will be much worse than you would hope for...

where you're getting huge vacuum spikes, this was important.

One of the most useful pieces of vacuum equipment is the residual gas analyzer (RGA) from SRS - for a complicated chamber like this, I can't imagine pumping down without it. We used it for

- Helium leak testing to find our cracked viewport
- Watching water pressure drop during water bake
- Seeing pumping speed of hydrogen relative to other species
- Seeing the virtual air leak when closing the big gate valve - you can see it's ambient air from ratio of N₂ to O₂
- Seeing organic compounds baking off during the bake
- Having convenient software for logging and plotting long-term pressure
- Obsessively checking the pressures during bakeout
- Trying to see Cesium after breaking the Cesium ampule
- Measuring out-gassing rates of each atomic species individually by turning off pumps and measuring the rate of rise in pressures

Since we used the RGA for probably 2 months of integrated time, the filament is starting to burn out. I already had to replace the filament once during our UHV test chamber runs, and it appears it needs to be replaced again very soon. We haven't operated the RGA since noticing this in the hopes that you might be able to get a little more information from it when it dies, and maybe you only need a little more information during the final alpha run... The RGA is on a gate valve, but if you remove it to replace the filament you'll need to add in a vacuum-T in order to add a pumping port to pump the RGA arm back down. This will cause some interference issues where the RGA is mounted, and will further limit the conductance between the RGA and the main chamber. I don't know the best way to go about this but I'm sure you'll figure something out.

We baked the chamber around 170C for several weeks. On either side of this, we slowly ramped the temperature up/down over the course of several days to ensure there were no thermal shocks, particularly to viewports. There were around 20 variacs hooked up to silicone heating strips, and around 35 thermocouples in different locations. This allowed us to carefully control the temperature across the entire chamber. We wrapped the chamber with many layers of UHV foil to help minimize temperature gradients. We began the bake with a large fan inside the pit room to bring in cold air, but it was blowing in a way that created a temperature gradient across one of the viewports in the detection region. The viewport cracked, causing sudden loss in vacuum pressure and the turbo pump shut off. After replacing the viewport, we baked without the fan and the room got warm but not unreasonably hot, so all was good. The RGA was on during the bake-out process so we could monitor the contaminant gasses dropping. Once the initial contaminants had neared a steady-state, we baked for an extra week or two, then slowly cooled down.

After baking, I activated each pump and gauge one-by-one. Activating a pump or gauge causes a bunch of crap to suddenly leave the surface, and this pressure spike can saturate other pumps. You therefore need to go in a circle and keep doing this to each pump/gauge until the pressure stabilizes - you want to see the pressure dropping well below what the turbo pump was doing before you started. The turbo pump had 500 L/s or so pumping speed, so seeing the pressure drop from other pumps will indicate that the pumping speed inside the chamber is closer to a couple thousand L/s. You can then close the gate valve to the turbo pump to see how the pressure behaves without the turbo pump - you should see the pressure drop even further because the turbo pump can only hold about 1000:1 ratio of hydrogen pressure between the output and input ports. After sealing the gate valve for good, I added an extra valve behind the gate valve and pumped this down as well, then sealed this valve. This ensures that failure of the gate valve wouldn't cause an immediate change to atmospheric pressure. Also, note that when we closed the large gate valve it initially looked like a virtual leak for a day or two, but this eventually went away and is apparently a common issue with closing large gate valves.

4.3.7 Summary of design decisions

The following is a summary of pro's (+) and con's (-) of some of the major design decisions discussed in the proceeding section.

- Magnetic shield and solenoid inside vacuum
 - + Higher aspect ratio tubes → less B-field fringing effects, longer interferometry region
 - Vacuum chamber diameter effectively doubles
 - Need to UHV clean enormous parts
- MOT and detection chambers very large diameter
 - + Made vacuum design simpler and interferometry region longer
 - Makes beam alignment much more difficult, especially for 3D MOT
 - Light collection is less efficient
 - Making these chambers smaller should have been more seriously considered during design phase
- Retroreflection mirror instead of two independent fiber ports
 - + Simpler and smaller magnitude beam-related systematics
 - + Easier alignment
 - More single photon scattering
 - Parasitic lattices limit total phase achievable
- Omitting quarter waveplate from vacuum

- + Better beam quality in-vacuum
- Lose 50% of optical power from combining beams on non-polarizing beamsplitter
- Even more parasitic lattices - see Section 4.6.3
- Mount chamber on hexapod struts
 - + Allows for much easier gravity alignment
 - Slightly more work to mount all optics directly to chamber
 - Added lots of mass near interferometry region - \therefore high order gravity gradients
- Add a permanent RGA, and use TSPs for with large surface areas for deposition
 - ++ No cons, these were super helpful for a large vacuum chamber
 - Just don't deposit titanium on viewports, especially the main interferometry one(s)

4.4 Installing chamber and getting vacuum

Another down-side to building an enormous chamber is that you can't move the parts by hand. The individual chambers weighed up to 300 lbs, and the magnetic shield + solenoid + baffles + central vacuum tube assembly weighed around 500 lbs. We talked to a couple moving companies, but the only company that would consider a project like this was Sheedy Drayage¹². I won't bore you with all the details, but needless to say it was very difficult to get this chamber installed. You can find the assembly steps created by Sheedy in the projects' Google Drive folder. I put one page of the assembly in Fig. 4.13 for an example of what was happening. I'll just give some highlights of the installation procedure here.

First, in order to save money, we decided to do some of the assembly steps ourselves - in particular the assembly of the magnetic shield layers, the solenoid tube, and the copper baffles. A horizontal assembly would have created extra dust from parts scraping together and would have likely damaged the solenoid winding, so we needed to work out a procedure to do a vertical assembly. Since the parts were around 10 feet long, we needed around 20 vertical feet in order to accomplish this. Like any reasonable AMO grad student would do in this situation, I found a single 1/4" hole in an I-beam on the bridge between Birge Hall and Physics North and I built a 'crane' that fixed onto this single hole. We came in on a Sunday since of course Anthony the building manager would have absolutely lost his mind if he ever found out we did this. See Fig. 4.12 for an image of some of the excitement.

Some other highlights from the installing include the assembly procedure I had to work out with Sheedy engineers to get the chamber installed (Fig. 4.13) and some images of the sweaty construction workers touching our precious UHV chamber (Fig. 4.14).

¹²This is the company responsible for lowering the BART tunnel into the bay. Yeah, the one that's several miles long and several hundred feet underwater.



Figure 4.12: In case you ever doubted the strength of a single 1/4-20 bolt... The crane is attached to the I-beam with a single extreme strength 1/4-20 bolt from McMaster-Carr, though a thorlabs 1/4-20 bolt should have still had a large factor of safety.




Figure 4.13: Page 4/6 in the vacuum chamber assembly steps from Sheedy. 18) the optics table is lifted up from the side of the room and 19) moved in place over the pit. 20) The table legs are re-installed. 21) The support structure is lowered around the main vacuum tube assembly. 22) The main vacuum tube is lifted up to install the detection chamber under the main tube. This is where I ran into a tolerance issue - the magnetic shield rests on pads inside the detection chamber, but the tolerancing of the braces mounted inside the vacuum chamber was screwed up by Kurt Lesker. I had to remove some washers from my assembly, then use all of the strength in my fingers to move the magnetic shield assembly over a few mm. Then finally everything slid together... 23) The top chamber is brought in and prepared to be mounted.



Figure 4.14: Images of the Sheedy workers during installation.

4.5 The rest of the atomic fountain

While building up the new system, we decided early on to try to replicate the old experimental details as exactly as possible in the new chamber - this would give us a baseline starting point to compare the new experiment's performance to the old one. From there, we could start changing things one by one to avoid changing too many things at the same time. Many of these sections directly compare parameters in the new experiment to numbers in the old experiment - this has allowed us to directly understand what we should expect in the new apparatus compared to the old apparatus.

By the time I'm graduating, we're still working to get the same sensitivity as the old experiment. As a result, there isn't too much different between the two systems. I'll highlight some large upgrades/changes that had to be made, but unless otherwise stated you should see the experimental details given in Brian and Chenghui's theses [23, 99].

4.5.1 Control system

One large upgrade to the experiment was getting a new control system. The old experiment used NI cards and LabView. It would have been very poor design to build on top of an already messy and not very scalable control system, so we decided to start from scratch. We purchased an Artiq control system which is FPGA controlled at its core, but all coding is done in python. We wanted the precision timing offered by the FPGA clocking along with all open-source code to enable thorough debugging. The downside of Artiq is that it was very new when we purchased our system - we were actually the first neutral atom experiment to use Artiq to control our experiment.

We originally tried to design a scalable, organized code structure based on 'blocks' that were assembled into a higher-level experimental sequence. For example, the MOT block would initialize all MOT-related channels, set MOT-related frequencies, and generally house any MOT-related functions. Then there would be other blocks for RSC, launch, interferometry, etc. We ended up running into issues with multiple blocks trying to initial the same channels. We also learned that in order to harness the precision-timing of the FPGA, you need to explicitly schedule events in absolute machine time units. This specifies the exact clock cycle in the FPGA when events will happen. Since many events are entangled during the core of the interferometer sequence, most of the code ends up living in a single file. The current state of the code is some messy mixture of the two of these ideas since the code moved in a direction different from the original design. Probably it should be overhauled, but the overhaul would break everything and require a lot of time to clean-up.

Another nice thing about this control system over the previous is that we wrote drivers for many of the external boxes in the experiment such as function generators, pulse timing generators, power supplies, piezo actuators, etc. These drivers are wrapped into artiq in a way that interacting with the external hardware ends up looking identical to interacting with internal artiq hardware.




Figure 4.15: a) an image of the actual 2D MOT design installed on the experiment, including the 2D MOT coils wound around the enclosure. b) A stripped down CAD model showing the fiber collimator, mirrors and beamsplitters used to form the 2D MOT beams.

4.5.2 Compact 2D MOT design

Another change we had to make from the old experiment was the geometry of the 2D-MOT optics. The old experiment used large 6" lenses to expand the cylindrical 2D-MOT beams to fill the entire 2D-MOT glass cell. This required a lot of free-space optics for the telescopes. In the new experiment, the 2D MOT needed to be close to the wall of the pit, so we couldn't afford this large space. We ended up designing a very compact setup that used non-polarizing beamsplitter to split and redirect two incident circular beams in order to generate three different 2D-MOT regions in the cell with four counter-propagating beams each. The down-side of this design is that it has very few degrees of freedom to adjust beam alignments.

Fig. 4.15 shows the final design used on the experiment. The beam sizes, optical power, optical detuning (-7.5 MHz), and magnetic field strengths were all designed to be roughly the same as the old experiment. The fiber collimators from Schäfter + Kirchoff are 200mm focal length which gives roughly 31mm $1/e^2$ diameter beams. We typically have 250 mW of light after each fiber, which then splits six ways to roughly 40 mW per beam passing through the 2D MOT region. The beam is split three ways at each face of the 2D MOT cell using 2" non-polarizing beamsplitters. First a 30/70 (reflection/transmission) cube reflects 30% of the light, followed by a 50/50 cube which reflects 50% of the remaining light, followed by a mirror which reflects all of the remaining light. This gives roughly equal power in each of the three beams.

The magnetic coils were wound with approximately 120 turns with a coil separation of 7" and a coil minor diameter of 7". Each coil is independently driven at around 5A current, but they slightly differ based on the currents that maximize atom number in the 3D MOT. The discrepancy is likely compensating for poor optical alignment that we can't correct because of lack of degrees of freedom in the 2D MOT optics. According to the magnetic field gradient

equation in Section 3.3.1 in Brian's thesis (not numbered) [23], our new coil geometry would give a magnetic field gradient along the central axis of 0.75 G/cm compared to their quoted 10 G/cm. However, using their geometry and currents, I calculate their gradient was only 1.5 G/cm using Brian's equation for field gradient.¹³ Based on the scaling in geometry and currents, the new design at 5A only has a 2x lower magnetic field gradient than the old experiment. I originally intended to run more current in these coils, but the coils begin to get very hot and we've found that this local heating affects the 3D MOT atom number. Likely the vapor pressure is getting so high in the 2D MOT cell that the cooled atoms can't make it out of the 2D MOT region without extra collisions. To directly address this issue in the future, one would need to wind more turns or use a much larger diameter magnet wire - however, the size constraints inside the pit make it very difficult to fit larger coils around the 2D MOT assembly.

4.5.3 3D MOT

The 3D MOT was also designed with very similar physical parameters as the old experiment. The MOT beams are 1.6 cm waist beams with -15.5 MHz detuning from the $F = 4 \rightarrow F = 5'$ transition. We use around 40 mW/beam at the chamber compared to the old experiment's 25 mW/beam, because increasing the optical power continued giving us a larger atom number. We don't go any higher than 40 mW/beam because the tapered amplifier generating the 3D MOT light would begin having a shorter lifetime if we ran it at higher currents.

The MOT coils are 36 cm diameter and separated by 50 cm. Each coil is made from 35 turns of square copper tube with 8/mm outer width and 6mm inner width using 35 turns. With this geometry, 200A gives 7.5 G/cm magnetic field gradient at the 3D MOT region, which is about 15% less than the magnetic field gradient in the previous experiment. The wire is hooked up to water cooling from a large heat-exchange chiller so that we can run the water in the laminar flow regime. Andrew Neely led the design and construction of this so I won't say too much more about this. Note a couple things though:

- We had a lot of issues with leaks near the connection from hose to copper wire - if you see water in the pit, stop and fix the leak.
- We had issues with the coils shorting to the aluminum holders, since the Kapton coating on the wire was poorly done and has gaps in it. If you move the coils around a lot you could introduce a new short.
- It's very important to have an interlock on the chiller water flow - if the water flow stops and the MOT coils are left on, the MOT coils could overheat until e.g. the Kapton coating between wires starts burning. The coils were extremely annoying to wind and were done by hand in the machine shop, and the connections from wire to hose were also braised on by hand - you don't want to have to do this again. We have an interlock set up on the water flow, and as a backup I installed temperature sensors on the MOT coils that will sound an alarm if they start to overheat.

¹³Suspiciously close to a factor of 2π error.

- The MOT coil current switching circuit was designed for a very fast (roughly 1 ms) switching time. We still need to wait 50-100ms experimentally for the magnetic fields to settle before we can do very good PGC cooling. Probably the eddy currents in the chamber walls are limiting our switching time.

Because the chamber in the new experiment is so large, the alignment of the MOT was very difficult and it took us a couple of months to finally achieve a 3D MOT. A successful MOT requires all six MOT beams to intersect in the middle of the chamber, overlapped with the MOT coil B-field zero and the 2D MOT atomic beam. Before you get the first MOT, you don't have a readout in the chamber yet so there's no way to ensure any of this alignment - it ended up being mostly guess and check until we saw an initial signal. A very helpful debugging step was to first look for a MOT in the old mini-G chamber which used a pyramid mirror, since getting a MOT only requires a single beam and is very insensitive to alignment. We ended up identifying a couple frequency issues by doing this. I mention the alignment issue with the large chamber because we're still not sure that everything is aligned very well. I've walked the 3D MOT beams a fair amount around relative to one another, but there many degrees of freedom so it's not very clear whether you've reached an optimum. If you're really struggling with atom number in the future, some walking of the alignment of the 2D MOT and 3D MOT beams, and maybe even the position/relative current in the 3D MOT coils, could still give a notable increase in atom number.

4.5.4 Raman Sideband Cooling

The old experiment implemented Raman sideband cooling (RSC) in the detection chamber while the atoms had already traveled a foot or so since the MOT and molasses launch. In the new experiment, we implemented RSC only an inch or so above the MOT region. This has a few benefits, the main one being that the molasses launch alignment is mostly decoupled from the atomic fountain alignment. The atoms don't have time to travel very far transversely from a crooked molasses launch. Then, because RSC cools the atoms to the rest frame of the 3D RSC lattice, it effectively relaunches the atoms independently of the molasses launch. The atomic fountain alignment can now primarily be done with only the RSC beam pointings, instead of also needing to walk the angle of the 3D MOT beams to change the molasses launch angle. The close positioning to the MOT region also means the atomic cloud can't expand much before being RSC cooled, so it keeps the atom cloud more dense during the rest of the atomic fountain.¹⁴

Another difference with the old experiment is the polarizations of light used during RSC. Since there is no quarter waveplate near the retro-reflection mirror, RSC needed to be set up with a different polarization scheme compared to the old experiment. Thankfully, this polarization scheme had already been worked in the original Steven Chu implementation back in the day [85, 44].¹⁵ So, we just replicated the original scheme and it worked in the end.

¹⁴It's possible that this was a benefit in the old experiment. RSC only cools atoms aligned with the center of the RSC beam, so this could have meant that the hottest atoms were lost before interferometry began.

¹⁵If you want to spend a day or two spinning your wheels trying to understand all the little details of how RSC works in neutral atoms, see Andrew Kerman's thesis [44]. Or, you can take my word that the details in the thesis don't really help any understanding beyond what was in the original PRL [85].

Last, I'll note that I think the best long-term design for the atomic fountain will be using a lattice launch. This will be discussed in-depth in Section 4.6.4. With a lattice launch, RSC can be done exactly in the MOT region. A slight molasses launch (a couple hundred mm/s) will break the velocity degeneracy of RSC and of the lattice launch, then the lattice launch can be performed after RSC. This alignment scheme (mostly) eliminates the RSC lattice alignment and molasses launch alignment from the actual atomic fountain alignment. RSC could also be done over a longer duration with lower pump powers, so it might be able to get closer to the recoil temperature limit. This would help with transverse velocity spread as well as atom number after velocity selection.

4.5.5 Detection

The detection scheme is overall very similar to the old experiment. We use a light sheet resonant with the $F = 4 \rightarrow F = 5'$ transition, with some repump light mixed in to bring any leakage $F = 3$ atoms back to the $F = 4$ state. The main difference in the new experiment is that the large diameter of the new vacuum chamber meant we needed a much larger viewport to achieve similar solid angle for fluorescence detection - see Fig. 4.2 for details. A 200 mm diameter viewport (10" CF, 8" aperture) is installed on either side of the detection chamber so that light can be collected from both viewports. Each viewport is 300mm from the central axis of the chamber.

Currently, our main data taking comes from one of the viewports where light is re-imaged onto the same exactly photodiode and amplifier circuit as in the old experiment - it was taken off of the old apparatus and installed on the new one. Two 200mm diameter, 300mm focal length condenser lenses from Edmund optics are used for light collection. Based on the exact positions of the lenses we've estimated a magnification of the atom cloud of 0.8 at the detection photodiode. This lens geometry captures about 2.6% of the emitted fluorescence compared to the old experiment's 6.2%.¹⁶

Photodiode array

On the other side of the chamber, we've installed a photodiode array that is a collaboration with Azriel Goldschmidt at Lawrence Berkeley National Lab. Madeline Bernstein has led this collaboration from our side, so I'll only give a brief overview. The same imaging optics (200 mm diameter, 300 mm focal length lenses) are used to re-image the atomic fluorescence onto the photodiode array. Each photodiode is 6 mm vertically by 600 μm horizontally, and 16 total photodiodes are connected to readout electronics.

Fig. 4.16b) shows an image of the photodiode array mounted to the PCB with readout electronics. Fig. 4.16a) shows some preliminary ellipse data taken with the photodiode array, illustrating loss of contrast in the interferometer towards the edges of the atom cloud. This is expected from the atoms seeing lower laser intensity near the edges of the interferometry laser beam.

The photodiode array already helps with alignment of the atomic fountain since we can directly see which way the atom cloud moves. We hope that in the future the array can be

¹⁶This is an estimate of the old experiment collection efficiency based on the geometry described in Brian's thesis [23]

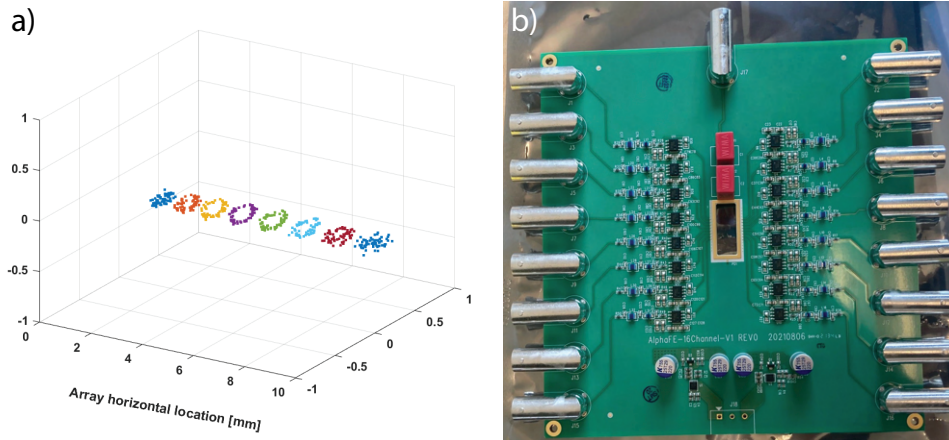


Figure 4.16: a) Ellipse data taken with the photodiode array, as a function of the horizontal position of the photodiodes. b) Image of the photodiode array mounted to a PCB containing the readout amplification circuit.

used to characterize systematic effects across the atom cloud, and potentially it is a new way to confirm our Monte Carlo understanding of spatially dependent phase shifts.

We’ve found that we can’t operate both the photodiode array and the single photodiode at the same time, because fluorescence reflects off of each photodiode face and makes its way to the other photodiode. It’s mostly an issue with light reflecting off of the single photodiode and then smearing out the signal on the photodiode array. Maybe there’s a way to fix this easily in the future by angling the chips, but for now it’s preventing both from being used at the same time.

4.5.6 Interferometry

The optics and electronics chain for generation of the interferometry beams is nearly identical to the old experiment. See Fig. 4.3 in Brian’s thesis for an illustration of the optics layout for interferometer frequency generation [23]. We use the same layout with the exception of the polarizing beamsplitter the combined beams in the old experiment has been replaced by a non-polarizing beamsplitter when combining the two optics paths ω_1 and $\omega_{2\pm}$. As a result, half of our optical power is lost to the other port of the combining beamsplitter. The light is sent to the chamber through a 15 m optical fiber, after which the exact same fiber port from the old α experiment is used to expand the beam to a 6.2 mm $1/e^2$ waist. We are not using the apodizing filter used for data taking in the 2018 α result [67] because the large clear aperture in the new chamber fixes the issue of power in the wings of the beam.

So far in the new experiment, we’ve gotten up to 7 Mrad of phase using a Bragg order $n = 3$, Bloch order $N = 100$, and interrogation time $T = 100$ ms. This is very close to the old experiment’s maximum total phase of 10 Mrad ($n = 5$, $N = 125$, $T = 80$ ms). Importantly, we’ve seen very good contrast still at these parameters - upwards of 20% contrast still at 7 Mrad phase. This is an indication that the larger clear-aperture in the new experiment has helped improve issues related to reflections off of the vacuum chamber walls, but more rigorous study of the contrast decay with T needs to be done to confirm this. Initial results

for Allan deviations show that we can average down to a few times 10^{-10} error in α after 24 hours of data taking. This is likely worse than the old experiment by a factor of a few, but we don't have good Allan deviation numbers from the old experiment for a direct comparison.

Our old MSquared TiSaph laser was dying ever since moving it to Campbell hall, and MSquared seems to be on the brink of going out of business... We bought a Sirah Matisse TiSaph laser that can output upwards of 6.5 W at 852 nm. It has a cavity included in the setup for short-term locking, but for long-term drifts we are still using the old experiment's strategy of phase locking to spectroscopy. This upgrade gave us an extra 30% optical power relative to the old setup.

Accounting for the factor of 2 loss in laser power from the old experiment due to not having a quarter waveplate in-vacuum, we'll likely never be able to drive $n = 5$ Bragg diffraction with a large enough detuning to also Bloch oscillations. We have some headroom now to go up to $n = 4$ Bragg though, and this will likely be implemented soon. For Bloch oscillations, we're limited by the parasitic lattice issues discussed in Section 4.6.3 below. If the launch velocity is increased, we may be able to increase to $N = 125$ or $N = 150$, but likely we won't be able to push too much above this. T can't be pushed out much further because the gravity gradient phase scales as T^3 and the gravity gradient systematic effect was already a large correction to the value of α in the old experiment at $T = 80$ ms [67]. As a result, the total phase achievable in the experiment will likely be very similar to the old experiment unless we get a significantly more powerful laser for interferometry. In order to make a better measurement of α , the best path forward is likely in decreasing shot-to-shot ellipse noise. Progress is already underway on this front - we're working on cooling the detection photodiode to decrease detection shot noise from photodiode dark currents, which we estimate is likely the limiting factor in the total detection noise.

4.6 New ideas

For lack of a better place to put these, I want to document some new experimental ideas that are in the works and not-quite-paper-worthy, at least not yet.

4.6.1 Systematic effect from velocity selection light shifts

While working on developing the Monte Carlo, I realized a potentially nasty systematic effect from light shift during Raman transitions. I haven't found any literature on this effect, but it would be very surprising if no one in our field has realized this since Raman transitions have been used in this way for 30 years...

In summary, the differential light shift during Raman transitions changes the atomic velocity class selected. Since the finite size laser beam intensity changes over the finite size of the atom cloud, the selected velocity class becomes a function of position of the atom within the laser beam. This leads to correlations between on-axis velocity and the position of the atom within the transverse plane of the laser. Since both the position and on-axis velocity of the atom affect diffraction phase from Bragg diffraction, this systematic has the potential to lead to phase shifts in an atom interferometer.




Figure 4.17: Diagram of relevant energy shifts in a typical 'lambda' Raman transition. The $F = 3$ and $F = 4$ ground states are energy shifted by their respective light shifts, leading to a two-photon detuning δ that is dependent on the local laser intensity.

Figure 4.17 illustrates a typical energy level diagram for a Raman transition with two counter-propagation lasers Ω_1 and Ω_2 . The two ground states are energy shifted by the respective light shifts of the two beam and respective detunings from the excited state. The excited state detuning Δ is assumed much larger than the hyper-fine state splitting so that the excited state manifold can be represented as a single state. The velocity resonance condition for the velocity-selective Raman transition is given by [54]:

$$2kv = \delta - 8\omega_r - \Delta + \frac{\Omega_2^2}{4\Delta} - \frac{\Omega_1^2}{4\Delta} \quad (4.1)$$

Ω_1 and Ω_2 vary with the local laser intensity, so the selected velocity depends on the location of the atom within the laser beam.

The actual experimental situation that we have is more complicated than what is shown in Fig. 4.17. Experimentally, we combine one beam with frequency ω_1 , with another that went through an EOM and has frequencies $\omega_2 + i * \delta_{EOM}$ where the EOM drive frequency δ_{EOM} is very nearly the Cesium ground state hyperfine splitting of ≈ 9.2 GHz. The optical power in each of the ω_2 frequency components varies as a Bessel function $|J_i(\beta)|^2$, where β is a dimensionless value related to the EOM RF drive power. All of these frequency components are retro-reflected in the chamber, so both the up-going and down-going passes of the beam interact with the atoms. An accurate calculation of the light shift and two-photon Rabi frequency therefore requires a careful summation of the contributions of all frequency components.

Figure 4.18 calculates the two-photon Rabi frequency, AC Stark shift, and single photon scattering for a range of detunings and EOM driving strengths. The calculation is based on parameters similar to those used in our experiment: 20 mW of light in the EOM beam, 2 mW of light in the non-EOM beam, a 6.2mm beam waist, and unless otherwise stated an EOM drive strength of $\beta = 1.8$ which roughly maximized the optical power in the ± 1 sidebands.

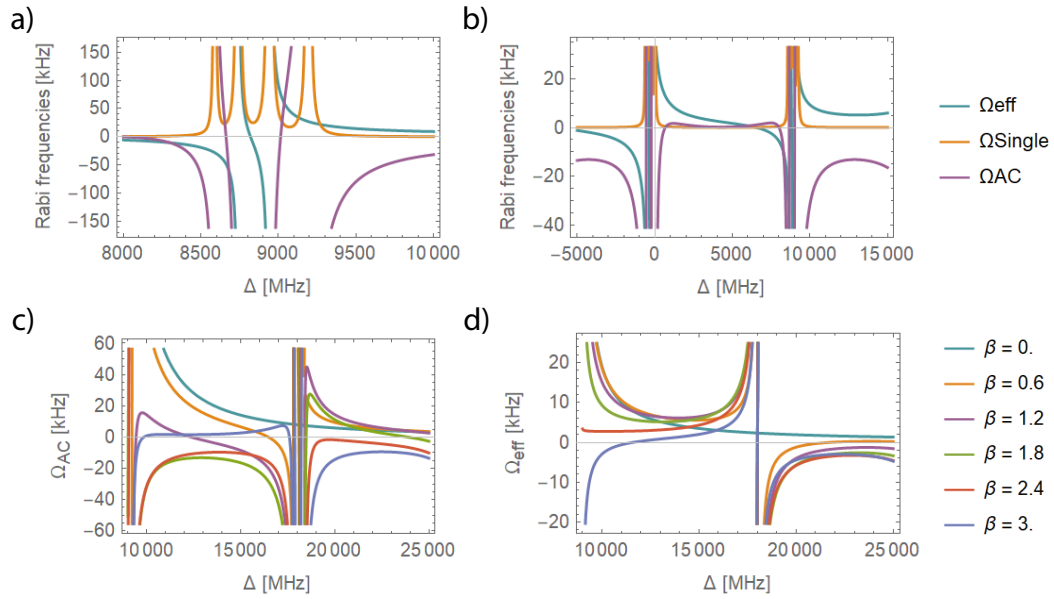


Figure 4.18: Numerical estimates of the AC Stark shift, effective two-photon Rabi frequency, and single photon scattering rate during Raman transitions. a,b) Rates for these processes for two different detuning ranges, a) zoomed in near a resonance with the excited state manifold and b) over a larger 20 GHz range. c) The AC Stark shift and d) the effective two-photon Rabi frequency as a function of detuning for a range of different EOM driving strengths β .

The frequency reference $\Delta = 0$ corresponds to the $F = 4 \rightarrow F' = 5$ transition. When one of the frequency components in the calculation nears single-photon resonance, the AC Stark shift, single-photon scattering rate and two-photon Rabi frequencies all diverge. In-between resonances, the summation of different processes leads to various zero-crossing for the AC Stark shift and the two-photon Rabi frequency. The portable gravimeter experiment in our group used this effect to find a detuning between the $F = 4 \rightarrow F' = 4$ and $F = 4 \rightarrow F' = 5$ where the AC Stark shift goes to zero, the scattering rate is relatively small, and the two-photon Rabi frequency is large [96]. This corresponds to roughly 9100 MHz detuning in Fig. 4.18a), though the ratio of powers between the two beams was much closer to one in the mini-g experiment so the graph does not represent their actual experimental setup.

Magic wavelength for Raman transitions

The systematic effect explored in this section relates to the differential AC Stark shift during a Raman transition. Figure 4.18c) shows that there are zero-crossings of the AC Stark shift where this effect would vanish. These so-called "magic wavelengths" are a function of the EOM drive strength β as well as the balance of optical powers between the EOM and non-EOM beams. Experimentally, we have been able to find magic wavelengths for a given set of parameters by varying the frequency of our TiSaph laser, equivalent to varying the x-axis on these plots. In the future, you will want the flexibility to choose any detuning. Fig 4.18c) and d) show that you can likely set a zero-crossing of the AC Stark shift at a given frequency by varying the EOM drive strength and/or the balance between optical powers, all while maintaining a non-zero two-photon Rabi frequency.

4.6.2 In-vacuum beam imaging

Since we can't put a camera into the vacuum chamber to image our laser beam, is there a way we can use our atoms to take a sensitive picture of the beam intensity? I was thinking about this during COVID-19 lockdown because I had nothing better to do. It turns out there is a very sensitive way to do this by turning the laser intensity information into phase information on the atoms using the AC Stark shift.

Using a Ramsey sequence of two microwave $\pi/2$ -pulses separated by a time T , one can turn the atomic phase information back into population information.¹⁷ If atoms are initialized in the $F = 4$ ground state, then the $\pi/2$ pulse will put the atom in a superposition of $F = 4$ and $F = 3$ ground states. If a single (detuned) frequency of light is incident on the atoms in the middle of the Ramsey sequence, the two ground states will see a differential AC stark shift and will phase-evolve at different rates depending on the local laser intensity. After the second microwave pulse, phase information on the wavefunctions will be turned into varying populations in the $F = 4$ versus $F = 3$ ground states. If the differential AC Stark shift resulted in many π radians of phase, then the population pattern across a Gaussian beam would resemble a circular fringe pattern.

Figure 4.19 shows a numerical calculation of the expected phase from differential AC Stark shift and the expected number of single photon scatters, as a function of the detuning

¹⁷You might be familiar with a Ramsey sequence as the sequence used in Cesium atomic fountain clocks which define the world time standard.

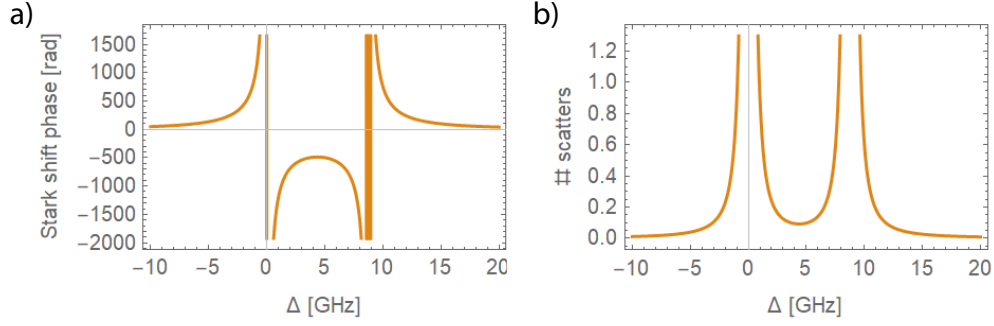


Figure 4.19: a) The total differential light shift phase and b) the number of single photon scatters as a function of laser detuning for a single frequency of light. Parameters used are 50 mW retro-reflected light over 0.5ms duration on a 6.2mm waist beam.

of the incident laser. $\Delta = 0$ corresponds to resonance with the $F = 4 \rightarrow F' = 5$ transition. Around $\Delta = 4.5GHz$, the $F = 3$ and $F = 4$ ground states see roughly equal and opposite laser detunings. In this region it's possible to accumulate thousands of radians of differential phase in a very short amount of time, with the probability of a single photon scattering event much less than one.

Figure 4.20 shows the expected population distribution for the same set of parameters. Across a $\sigma = 1.5mm$ atom cloud, the population distribution shows a very fine structure with population variation wavelengths less than $100\mu m$. Figure 4.20b) illustrates the dramatic sensitivity of this method in detecting very small intensity ripples in the beam - 0.2% intensity ripples are extremely clear in the population distribution.

There are a couple of challenges in using this method. First, one needs to ensure that the atom ensemble thermal velocity doesn't wash out the population distribution before it is imaged. After Raman sideband cooling, atoms have an average thermal velocity of around 5.5 mm/s in the axial direction, which is what we are interested in. In 5 ms, the atoms therefore travel on average $27.5\mu m$, which gives the length scale below which fringes begin to wash out. If the above scheme including imaging time can be accomplished in less than 5 ms, then thermal velocity won't be a limiting factor.

If the sequence is performed directly after PGC, the thermal velocity will be several times higher than after Raman sideband cooling, so either the imaging needs to be done more quickly, or the length scale of fringes that can be imaged will suffer. Moreover, the magnetic sublevel distribution of atoms will be spread roughly uniformly over all magnetic sublevels. If there is an applied magnetic field, the sublevels will split, and the microwave $\pi/2$ pulse will only perform a Ramsey sequence on one of the magnetic sublevels. If there is no magnetic field, all sublevels will contribute to the Ramsey sequence. Different magnetic sublevels will acquire different phases from small magnetic field gradients, which could mask the effects from the beam that are trying to be measured.

Last, it is difficult to get a fast, high-quality image of the atom cloud along the axis that the beam interacts with the atoms. After the Ramsey sequence is complete, one will need to perform either fluorescence imaging or absorption imaging on the atom cloud distribution. Absorption imaging likely won't work well because the $100\mu m$ length scale of the fringes will have a short Rayleigh range once imprinted on the laser beam, and may wash out before




Figure 4.20: Expected population distribution in the e.g. $F = 4$ state for a) Gaussian beam profile and b) Gaussian beam profile with an additional 0.2% intensity ripple with wavelength $\approx 300\mu\text{m}$. In reality the fringes should remain at 100% contrast over all of space, but the fluorescence will drop outside of the atom cloud. Here I represent this by weighting the contrast by the atom cloud population. Thermal motion and detection effects are not accounted for in this plot.

reaching the camera. Fluorescence imaging will be difficult because the camera will have to be at least ≈ 1 meter away from the atom cloud and therefore the solid angle of fluorescence detected will be very small.

I'll leave the exact details to the next generation of students who actually implement this, but care will need to be taken to ensure:

- Maximal detected photon count per pixel on the camera by using large diameter collection optics.
- Large depth of field so that signal from the top and bottom of the cloud doesn't start washing out the pattern. This will require a small aperture in the collection optics
- The camera alignment to the cloud needs to be very nearly on the same axis as the beam, so that the vertical columns of population remain aligned to the camera. This may require a 50/50 cube under the chamber so that the camera can be exactly on-axis
- All imaging needs to be done within several milliseconds to ensure thermal motion of atoms doesn't wash out image, as discussed already
- The 100 mW Stark-shifting beam probably shouldn't hit the camera directly in order to protect the CCD chip
- Probably you want to image the up-going and down-going beams separately using this method - to do this, the retro-reflection of the beam will need to be slightly misaligned so only one pass hits the atom cloud and the other pass doesn't cause any Stark-shift phase.




Figure 4.21: Illustration of the frequency pairs in the experiment during multi-frequency pulses, and our gold retro-reflecting mirror. We want atoms to only interact with the lattices formed by the upgoing ω_1 beam and the downgoing ω_{\pm} beams. All other lattices are parasitic - they cause unwanted perturbations to the Bragg diffraction or Bloch oscillation dynamics. See text for further discussion.

Even with all of these considerations, the numbers seem to indicate that we could still do a fairly good job at imaging part-per-thousand defects in the beam intensity. Jack Roth is currently working with Azriel Goldschmidt at Lawrence Berkeley National Lab on making an optimized setup for this imaging scheme.

4.6.3 Systematic effect from parasitic lattices

As discussed in the beginning of the chapter, the decision to omit a quarter waveplate from inside the vacuum chamber changes the polarizations used to drive Bragg diffraction/Bloch oscillations. All frequencies of light must have the same polarization of light, which means that any pair of frequencies will form an optical lattice that can interact with the atoms. In the old experiment, the orthogonal polarizations ensured that there were only a few parasitic lattices that were formed.¹⁸

This section will outline qualitatively how these parasitic lattices might impact the experiment, but further investigation will be needed in the future to make this more quantitative. See Fig. 4.21 for an illustration of the effect. Bloch oscillations and the second two Bragg diffraction pulses require three optical frequencies sent into the chamber, as opposed to the two frequencies required for Bragg diffraction or Bloch oscillations of one atomic velocity. I will commonly refer to these three-frequency pulses as 'multi-frequency' pulses from here on. The two lattices formed by the up-going ω_1 frequency and the down-going $\omega_{2\pm}$ frequencies interact with the lower and upper interferometers, respectively. These pairs are indicated by the solid lines connecting frequencies in Fig. 4.21. All dashed lines represent real lattices formed in the experiment that move at various velocities - they are not the lattices we intend to use for physics, so we call them parasitic. In the old experiment, the only parasitic lattices with non-zero matrix elements were the down-going ω_1 frequency interfering with

¹⁸See Chenghui's thesis Chapter 2 for a good discussion of the polarizations in the old experiment and which combinations of polarizations lead to non-zero matrix elements.

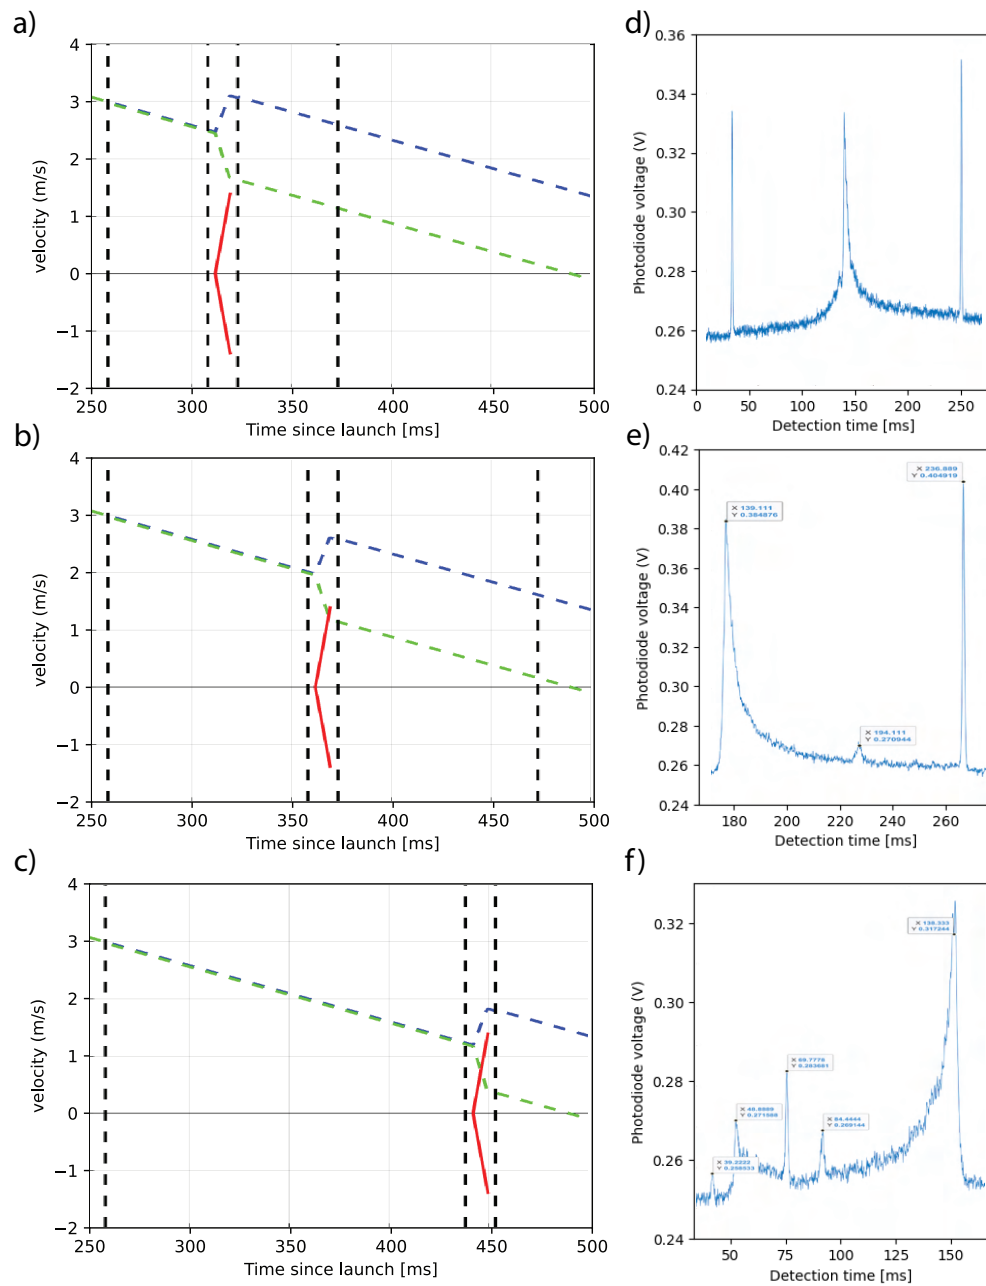


Figure 4.22:]

a-c) Illustrations of various moving lattices during the experimental sequence. Green and blue lines roughly represent the atomic trajectories accelerating under gravity, ignoring the (small) Bragg diffraction velocity splittings. Vertical dotted lines show the time of Bragg pulses. During Bloch oscillations, the atom velocities are accelerated away from one another - the solid red lines show the parasitic lattices that begin to cross paths with the desired lattices in b) and c). d-f) Experimental traces of the respective scenarios shown in a-c), where the parasitic lattices begin trapping population and perturbing dynamics.

Figure thanks to Yair Segev.

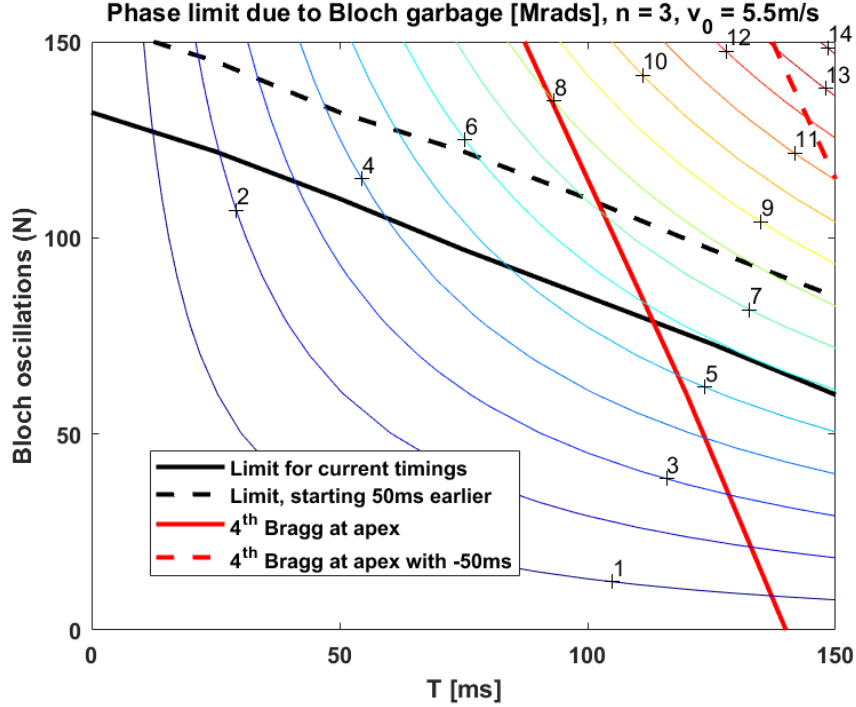


Figure 4.23: Limits to the total phase achievable in the experiment as a function of the number of Bloch oscillations and the interrogation time T , due to the parasitic lattices illustrated in Fig. 4.22. The colored gradient lines indicate contours of constant total phase (in Mrad), with cross-haired numbers indicating what phase each contour represents. The solid black line shows the limit when parasitic lattices start interfering with the Bloch oscillation pulse. Above this line, the Bloch oscillation dynamics will be highly perturbed. The red line shows when the fourth Bragg pulse occurs at the apex of the trajectories - it is not a limit to the total phase, but we need to stay clear of this line since standing waves in the lab frame will disrupt the Bragg diffraction dynamics.

the up-going $\omega_{2\pm}$ frequencies, since these were the only other combination of beams that had opposite polarizations. These would have been moving downwards in the chamber, and therefore would have stayed far off-resonant from the atoms while they were still upwards in their trajectories. In the new experiment, all dashed lines now interact with the atoms, so additional care must be taken to ensure that parasitic lattices don't perturb the dynamics.

Fig. 4.22 shows how these parasitic lattices interact with the atoms during the experimental sequence. If the pulse separation time T of the interferometer is short enough, the problem parasitic lattices formed by the interference of the up-going ω_{2+} and down-going ω_{2-} beams is far-detuned¹⁹ and does not grossly interfere with the Bloch oscillation dynamics. However, for longer T 's, the velocity of the parasitic lattices crosses the velocity of Bloch lattice(s) and severely perturbs the dynamics. Experimentally, we see total loss of contrast when this interference occurs and significant population loss from the original ground state

¹⁹I say 'far-detuned' in a velocity sense. This still translates to far-detuned in a frequency sense when working with the Schrodinger equation, were the frequency difference between optical frequencies is much greater than a recoil frequency.

Bloch lattice.

Note also that the $v = 0$ lattice is a (strong) standing wave in the lab frame - this lattice will also highly perturb dynamics for any processes occurring near zero velocity. Moreover, this standing wave lattice is the sum of the three lattices formed by the three frequencies in Fig. 4.21 interfering with themselves. The three superposed lattices have different spatial frequencies, which will cause a standing-wave structure in their interference. The frequency differences are on the order of tens of kHz to a few MHz, so the wavelength of the spatial structure will be hundreds of meters. Since the lattices have the same zero electric field boundary condition at the gold mirror, the interference pattern should add constructively near the gold mirror, so the spatial structure of this lattice shouldn't play a role in the systematic effects. I just wanted to make a note of this in case it's ever relevant.

Last, recall that in the old experiment, crossed polarizations ensured that these lab-frame standing wave lattices had zero matrix element to drive Bragg diffraction or Bloch oscillations, and therefore did not interact with the atoms.²⁰

One way to get around this issue would be to break Bloch oscillations into multiple steps - you could unload the lattice when the parasitic lattice gets close to the atomic velocity, let the parasitic lattice pass by the atomic velocity (from gravitational acceleration), then continue Bloch oscillation accelerations. This would make the experimental sequence much more complicated, and it might risk adding new systematic effects to the already complicated system. If you're not willing to break Bloch oscillations into multiple pulses, the constraints on T and N from the parasitic lattices will limit the total phase achievable in the experiment. Fig. 4.23 shows the limit to total phase in the new experiment due to this parasitic lattice issue. For reference, the old experiment used a maximum of $10Mrad$ phase in data taking for the measurement of α . The easiest way to increase total phase in the future will be to launch atoms faster so that the parasitic lattices are further detuned.

These parasitic lattices can also cause T-dependent diffraction phases, similar to those discussed in Section 4.2.1 but from a very different origin. Note that these effects are not new to our project, but this is an added source of these effects [66, 24]. Off-resonant lattices distort the typical Bragg diffraction dynamics, leading to phase shifts on the wavefunction. Symmetries in the SCI interferometer geometry ensure that the 1st and 4th pulse diffraction phases are cancelled or nearly cancelled.²¹ The diffraction phase from the third pulse ends up being the dominant contribution to the SCI interferometer since the multi-frequency pulses have a second perturbing lattice that leads to diffraction phase shifts. The old experiment found that by fixing the location of the second and third Bragg pulses in the atoms' time of flight, the diffraction phase drifts were substantially minimized. This is likely because the parasitic lattices that did affect the old experiment, discussed above, would have had a strong T-dependence since the velocity of the atoms in the lab frame dictates how off-resonant the parasitic lattices are. I suspect this is why they found fixing the location/time of the second and third Bragg pulses to make such a difference. In this generation of the experiment, similar shifts from parasitic lattices will likely be even worse - we might still be able to use the same trick of fixing the time/location of the second and third Bragg pulses

²⁰They still stayed away from zero velocity in the old experiment, likely because impure polarizations would have allowed for a weak interaction with the atoms that still interfered with dynamics.

²¹See Brian's thesis for a good discussion [23], as well as the paper form of this section of his thesis [66].

to eliminate or at least strongly suppress the effects.

Looking towards a next-generation experiment, I would recommend using two fibers instead of a retro-reflection mirror and one fiber. Using two fibers like this would save 2x in laser power and eliminate all parasitic lattices. This gives a lot more flexibility in interferometer geometries and total phase. It would also reduce single photon scattering by a factor of 2 which would further allow for more total phase, since you'll have more signal after many Bloch oscillations. If you go down this route, then you should also incorporate a well-designed system for characterizing the beam-beam alignment and the in-vacuum intensity profile of each laser beam, so that the systematic effects from two beams aren't unreasonably difficult to characterize. For the current generation of the experiment, we first want to get a handle on the single beam systematics before making the beam-related systematics more complex, so I don't think it was a huge mistake in our design decisions. We just didn't realize the extra limits on total phase that we would run into.

4.6.4 Lattice launch

The last-generation α experiment spent a great deal of time re-aligning the atomic fountain. Their setup required co-alignment of the molasses launch, RSC lattice, and interferometry beams all to align with gravity. Too many degrees of freedom and very strict alignment criteria made this very difficult to do. There was no alignment procedure that converged, so they had to guess and check.

In the new experiment, we have already reduced part of this issue by moving RSC to be much closer to the MOT region, thereby reducing the sensitivity to the molasses launch alignment. If alignment continues being an issue in the future, a lattice launch should fix the issue.²² By launching atoms with the same beam that is used for interferometry, the launch is automatically co-aligned with the interferometer beam. RSC can be done in the MOT region before launching with a very small molasses launch, which will further reduce the launch sensitivity to molasses launch and RSC alignment.

For an example of another experiment that uses a lattice launch, see Mark Kasevich's 10m fountain in e.g. this thesis [79]. They use slightly off-axis beams that interfere at the location of the MOT in a way that gives a vertical launch - this must be a terrible setup to have to align... For our system, it will be much more effective if we can launch atoms using the same beam as the interferometer beam, so that we automatically have co-aligned launch and interferometry laser. This will have to come at the cost of optical power, since we don't have a good way to combine beams before the last fiber without losing optical power.

The laser power requirements aren't too bad for this system. Figure 4.24 shows a numerically calculated survival fraction of the atomic ensemble after such a lattice launch, which integrates over the finite size of the atom cloud and laser beam.²³ In order to launch to 6 m/s

²²At the cost of added complexity and probably a lower launched atom number...

²³If these plots look familiar, it's because my code is based on Matt Jaffe's code that he used for plots in his thesis. My code can be found in the git repo `alpha_random/atom_laser_interactions/Finite size LZ tunneling and scattering.nb`. Also, note that there are many many places to have factor of two errors in this calculation. I've spent a ton of time searching for factors of two in this script, so don't start making changes unless you're 100% certain you're correct or have experimental data to back it up.

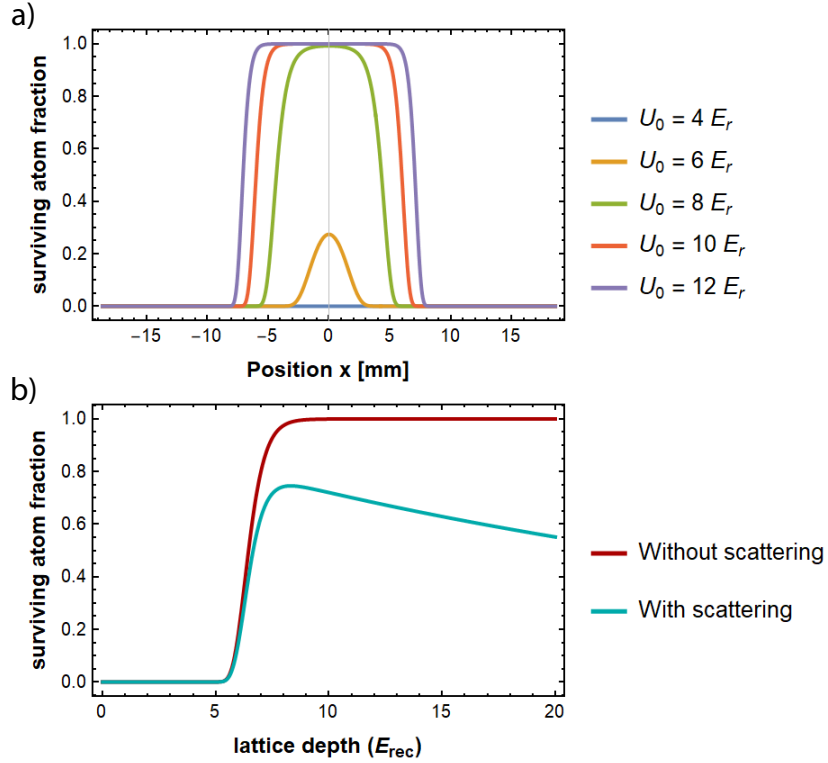


Figure 4.24: a) Landau-Zener survival probability for 940 Bloch oscillations with the given lattice depths, as a function of the position of the atom relative to the center of the beam. For saturated Bloch oscillations, tunneling creates a sharp cutoff for surviving atoms. b) The survival probability of a lattice launch with and without single photon scattering included for a detuning of $2\pi * 100$ GHz and an acceleration rate of $10g$.

with an acceleration rate of 10g, you need around 940 Bloch oscillations.²⁴ Other parameters used in the simulation are a beam waist of 6.2 mm, atom cloud waist of 1.5mm, and a blue detuning of $2\pi * 100$ GHz. The optimal lattice depth is around $8.5E_r$, which corresponds to about 250 mW of light in each frequency component. For someone trying to design a lattice launch laser system, this is a lot of power but nothing crazy. The important feature is that you need a very large detuning. The laser will be interacting with the atom cloud for 60 ms for these parameters, or even longer if you need to ramp more slowly.

²⁴Recall that the atoms actually experience a force of 11g because gravity adds to our lab-frame acceleration.

Chapter 5

Monte Carlo overview

5.1 Introduction

Most quantities measured by atom interferometers require very precise knowledge of the k -vector of the laser. As discussed in Chapter 1, our experiment measures the recoil frequency of Cesium $\omega_r = \hbar k^2/(2m)$, so turning our measurement of ω_r into a measurement of \hbar/m requires knowledge of k^2 . For ideal plane waves laser beams, the k -vector is given by $k = \omega/c$, where ω is the optical frequency and c is the speed of light. If we had a plane wave laser beam and measured its frequency ω using an optical frequency comb, then we would know k to very high precision.

However, finite sized laser beams cause the magnitude of the k -vector to deviate from ω/c , and even the direction of the k -vector rotates away from the axis of propagation. The well-know ‘Gouy shift’ captures deviations in the k -vector from a Gaussian laser beam[27]. This result has recently been generalized to a local calculation of δk based on the local amplitude curvature and phase gradient of the electric field of the laser [6], which allows for a general treatment of momentum transfer from distorted non-Gaussian laser beams. This will be discussed quantitatively in Section 5.3.2

While many models have been used to capture various aspects of beam-related effects in atom interferometers, no single model has been used to demonstrate a complete understanding of the resulting effects[89, 91, 33, 67, 6]. The goal of this Monte Carlo package is to capture all of the relevant physics that might shift our measurement of α .

Beam-related systematics are likely to be the most difficult to characterize in our entire experiment moving forward, so having a model of beam-related systematic effects that we trust will be crucially important the a next-generation measurement of α . To this end, it’s also crucial that we verify the model experimentally. Part of the aim of this project is demonstrate an understanding of beam-related effects by intentionally adding large distortions to the laser beam in order to match our model with experimental results - larger distortions on the beam should make it easier for us to measure the effects.

By treating the physics in our Monte Carlo as generally as possible, we have identified a couple new sources of phase error that needs to be considered in the treatment of inhomogeneous laser beams in atom interferometers. First, deviations δk in the k -vector of the laser can cause the interferometer arms to not overlap exactly at the closing of the interferometer.




Figure 5.1: Plotted here is the $\delta k/k$ Gouy shift near the focus of a $w_0 = 3.25$ mm $1/e^2$ radius beam, roughly what the old α experiment used.

We have found that the resulting separation phase needs to be taken into account in order to accurately capture the resulting phase shifts. We have also found that off-axis components of the k-vector must be accounted for in accurately modeling contrast loss, another affect that has not yet been included in systematics studies. In total, this model presented in this chapter will hopefully serve as the foundation for beam-related systematics analysis in our next published measurement of the fine-structure constant.

5.2 Motivation

Given the desire for such a model, why does the α experiment need a Monte Carlo simulation to execute it, as opposed to analytic estimates? Some systematic effects in our experiment can be calculated analytically and are insensitive to the finite size/ temperature of our atom cloud and the finite size of our laser beam. For example, the gravity and magnetic field gradients change over length scales much larger than the atom cloud. These types of systematic effects can be treated analytically by calculating a correction to the measurement.

Other systematic effects, however, are highly dependent on the size and temperature of the atom cloud, and the size of the laser beam. The Gouy shift is the primary example - it varies $O(100\%)$ as you go from the center of Gaussian laser beam to one waist off-axis. The 2018 α publication [67] corrected their measurement at -2.6 ± 0.03 ppb, meaning they claimed knowledge of this effect at the 1% level. Fig. 5.1 shows (very roughly) what the Gouy shift looked like in their experiment.¹ For an atom cloud with a 1σ waist of 1.5 mm, the Gouy shift varies by 10's of percent across the extent of the atom cloud. Moreover, the atom cloud is expanding during the time of flight and between the different interferometer pulses, so the ensemble-averaged effect is changing pulse-to-pulse.

¹Their laser beam was slightly flatter in the center than a typical Gaussian beam, so their Gouy shift would have been slightly lower than what's plotted in the center region of the beam. Also, since α ends up being proportional to the k , the average $\delta k/k$ can be directly interpreted as the correction made to the value of α .

One could imagine calculating an expectation value for the Gouy shift analytically by using a Gaussian integral between your atom cloud distribution and the Gaussian laser beam. However, the varying laser intensity also couples to the survival probability of an atom through the interferometer, which leaves correlations between the Gouy shift and the contribution to the final interferometer contrast. Let's say you also take this into account in your expectation value, and can still get closed-form integral for a single Bragg pulse. In reality, there are four Bragg pulses and one Bloch pulse, and the atoms are moving around thermally between each of the pulses. Somehow, maybe you could still get an analytic solution written down when accounting for all of these effects. The kicker is that our experimental laser beam is not an ideal Gaussian beam, and it turns out that the non-ideal nature of the beam can have significant effects on the ensemble-averaged Gouy shift. All of these effects together, and many more that I don't mention here, make the problem far too complicated to treat analytically - we need to study these systematic effects with a Monte Carlo simulation.

The previous generation alpha experiment had two Monte Carlo packages that were used in their final publication [67]. One was initially made in Matlab by Brian Estey, and then Richard Parker continued developing it after Brian left. The other package was made in Python/C by Chenghui Yu from scratch in the final years of his thesis. The old codes worked as needed for publication of a measurement of α [67], however there are a number of improvements that can be made moving forward.

The new code:

- has much more documentation so that new users can build off of the existing code instead of writing a new code package
- treats Bloch oscillations generally through the Schrodinger equation, instead of assuming all phase shifts from Bloch oscillations cancel
- includes loss of contrast effects from transverse momentum kicks due to light shifts and angled k-vectors, giving a more accurate model of contrast loss
- runs much much faster by using a look-up table for Schrodinger equation evolution, instead of integrating the Schrodinger equation directly for each atom in the Monte Carlo
- calculates local lattice depth and light shifts from a general 3D model of the laser propagation, instead of using only a few CCD images

To follow up more on the last two points, one concerning aspect of the previous Monte Carlo work was that an image filter was used to smooth experimental CCD images of the laser beams, since it's very difficult to take accurate images of a laser beam with a CCD due to dust on optics and etaloning within the camera or within the CCD chip itself. They varied the pixel width of the image filter to match experimental contrast to predicted Monte Carlo contrast - if their Monte Carlo model for contrast loss was incorrect, this would could cause issues in estimating the Gouy shift δk from small scale intensity inhomogeneities. The current Monte Carlo model aims to develop a more complete model of contrast loss so that potential issues like this won't arise in the future. Moreover, we are working to obtain

accurate CCD images of the beam so that we don't need to use an image filter to remove imaging artifacts on the beam profile. A recent result from the MAGIS-100 collaboration might be very useful in obtaining defect-free images of the beam [40], where they take a series of images on different parts of the CCD and with different angles of the CCD, then use principle component analysis (PCA) to identify the actual modal profile of the beam.

This chapter gives a high-level overview of the physics built into the Monte Carlo model, followed by an overview of the code package that simulates this physics. I then discuss some initial results using the Monte Carlo model to study systematic effects from laser beams with significant intensity profile distortions. The work to compare this model to experiment is ongoing, so we don't yet have experimental results at the time of publishing this thesis. Stay tuned for more results in the near future.

5.2.1 Monte Carlo Design Philosophy

Before diving into the weeds, I'd like take this opportunity to strongly encourage future students on this project to put effort into building a shared Monte Carlo code instead of re-writing a code package from scratch, as has been done for the last ten years or so.² There is a strong culture in atomic physics of doing everything yourself, and I'm sure any of you reading this would be interested in writing your own Monte Carlo. Of course, doing everything yourself is how you learn and understand things the best. However, the scope of the α project is just too large for any one student to do all of it - if you scrap my code, it'll take you at least a year of work to catch up, likely more. Instead, you should spend a month learning how to use the code, then start building on it and improving it, and using the code for studying interesting physics. If you can start a culture of everyone building on the same code, you'll have a much more powerful Monte Carlo and a much more efficient team.

In addition, we've learned over the last ten years that beam-related systematic effects are some of the nastiest and hardest to characterize effects in our experiment. The Monte Carlo analysis will only get more complicated as the experiment develops, so having a shared complex code package will become increasingly important in the future.

5.3 Monte Carlo physics

Let's first look at the physics that the Monte Carlo model contains. At it's core, the code simulates a single atom going through an SCI interferometer. The atomic wavefunction is evolved through Bragg diffraction and Bloch oscillations by numerically integrating the Schrodinger equation, which solves for the amplitude and phase of atoms along each of the interferometer trajectories. These Schrodinger equations were derived in Chapter 2. A model of propagation of the laser beam electric field, including retro-reflection, is used to calculate the local lattice depth, light shift, laser phase, etc. At the end of the interferometer, overlapping wavefunction trajectories are interfered and phase information on the wavefunction is converted to population information in the output ports. These populations are used to parametrically plot the output on an ellipse, mimicking our experimental data analysis

²I can't resist a good opportunity to preach some unsolicited advice...

pipeline. The ellipse is fit to extract the differential phase between two interferometers as well as the interferometer contrast.

Most of the physics used in the code is wrapped in boolean parameters so that the user can easily turn on/off different physical effects, for example turning on/off gravity, simulating zero temperature on-axis atoms (no random initial conditions), or turning on/off the Gouy shift. However, the main application of the code is to simulate as closely as possible our exact experimental conditions. For these simulations, the atoms are moving upwards at around 3 m/s near the interferometry region of the chamber when interferometry begins, with an on-axis velocity spread of around $0.05 v_r$ ³ and an off-axis velocity spread of around $1.5 v_r$. The 1σ radius of the atom cloud is about 1.5 mm and the laser beam waist is around 6 mm. Accurate distances for retroreflection of the beam are taken into account. There are many more parameters included in the model but I won't list them all here. The point is that the code is written in a way that allows the user to easily switch from highly simplified situations that isolate individual physical effects to using the full generality of the physics in the model to predict real experimental outcomes. The former is very useful for debugging or for gaining intuition for different parameters, but the latter is where the true predictive power of the Monte Carlo code lies.

5.3.1 Laser Electric Field Models

To begin, many physical effects in the Monte Carlo model require knowledge of the electric field propagation of the laser beam, not just the intensity propagation. For example, the laser phase imparted to the atom depends on the phase of the electric field. As another example, deviations in the local k -vector depend on the gradient of the phase of the electric field. To keep the model as general as possible, the code is currently built on analytic functional forms of the propagation structure of the laser beam electric field $E(x, y, z)$.

There are three important models I'll outline here: 1) a standard Gaussian laser beam, 2) a Gaussian beam with a small 1D sinusoidal or square wave amplitude modulation, and 3) a model for a Gaussian beam that passes dust on optics. Some data is shown in Section 5.9 that experimentally verifies propagation of the square wave model, which is one of the more involved models used in simulations.

Gaussian beam

The evolution of a Gaussian laser is very well known (see e.g. [80]), so I will simply state the electric field evolution here as a reference:

$$E(x, y, z) = E_0 \frac{w_0}{w(z)} \exp\left(\frac{-(x^2 + y^2)}{w(z)^2}\right) \exp\left(-i\left(kz + k\frac{x^2 + y^2}{2R(z)} - \psi(z)\right)\right). \quad (5.1)$$

In this expression, E_0 is the peak electric field amplitude at the center of the beam focus $E(0, 0, 0)$. $w(z) = w_0\sqrt{1 + (z/z_R)^2}$ is the beam waist as a function of propagation distance z , where the Rayleigh range $z_R = \pi w_0^2 n / \lambda$ with an index of refraction n and optical wavelength

³ $v_r = 3.5\text{mm/s}$ is the Cesium recoil frequency for 852 nm light

λ .⁴ $1/R(z) = z/(z^2 + z_R^2)$ captures the radius of curvature $R(z)$ of the phase fronts as the beam expands, and it is commonly written as an inverse like this to avoid infinities at $R(0)$. $\psi(z) = \arctan(z/z_R)$ is the Gouy phase, which captures a shortening of the laser k -vector near the beam focus [27].

A Gaussian beam is used by default in the Monte Carlo code since it is a typical first approximation to an experimental laser beam profile. One thing to note is that for 6 mm beams used in our experiment, the Rayleigh range is around 130 m. This is much longer than the ≈ 10 m of propagation that the laser undergoes inside the chamber, so the laser beam is very nearly collimated over the entire extent of the experiment.

Sinusoidal modulation

Next, we derive the propagation of a model for a sinusoidally modulated Gaussian beam so that we can study the resulting systematic effects as a function of the spatial frequency of the modulation. The initial beam is assumed to have a small sinusoidal amplitude modulation on top of an overall Gaussian amplitude profile. For a modulation along one transverse axis, it is sufficient to solve the Helmholtz propagation only along that transverse axis. The orthogonal transverse axis will undergo standard Gaussian evolution.

The amplitude for the laser electric field at $z = 0$ is assumed to be:

$$u(x, 0) = u_0 e^{-x^2/w_0^2} [1 + (A_q + i\phi_q) \sin(qx)] \quad (5.2)$$

where waist w_0 is the $1/e^2$ intensity radius of the beam, A_q is the initial amplitude modulation and ϕ_q is the initial phase modulation, and q is the spatial frequency of the modulation.

To solve for the propagation of the beam including in the near-field, we can use the Fresnel diffraction integral:

$$u(x, z) = \frac{1}{i\lambda} \int u(x, 0) \frac{e^{ipikr}}{r} \frac{z}{r} dx' \quad (5.3)$$

which is valid as long as the beam has propagated much more than a beam waist, in our case 6 mm. Provided this is true, it's a very good approximation for the true propagation of the beam after solving the full Helmholtz. equation. Applying this to our initial conditions in Eq. 5.2, we arrive at

$$u(x, z) = u_0 \frac{e^{ikz}}{i\lambda z} \left[I(0) + \frac{1}{2i} (A_q + i\phi_q) (I(q) - I(-q)) \right], \quad (5.4)$$

where

$$I(q) = \frac{\sqrt{2\pi}}{\sqrt{\frac{2}{w_0^2} - \frac{ik}{z}}} \exp\left(\frac{-w_0^2(kx - qz)^2}{k^2 w_0^4 + 4z^2}\right) \times \exp\left(\frac{i(4kx^2 z - kq^2 w_0^4 z + 2k^2 q w_0^4 x)}{2k^2 w_0^4 + 8z^2}\right). \quad (5.5)$$

⁴We use $n = 1$ because the experiment operates in vacuum, and $\lambda = 852$ nm

The expression for $I(q) - I(-q)$ is given by

$$I(q) - I(-q) = \frac{\sqrt{2\pi}}{\sqrt{\frac{2}{w_0^2} - \frac{ik}{z}}} \exp\left(\frac{-w_0^2(k^2x^2 + q^2z^2)}{k^2w_0^4 + 4z^2}\right) \exp\left(\frac{4ikx^2z - ikq^2w_0^4z}{2k^2w_0^4 + 8z^2}\right) \\ \times \left(\exp\left(\frac{2iqxk^2w_0^4}{2k^2w_0^4 + 8z^2}\right) \exp\left(\frac{2kqzxw_0^2}{k^2w_0^4 + 4z^2}\right) - \exp\left(\frac{-2iqxk^2w_0^4}{2k^2w_0^4 + 8z^2}\right) \exp\left(\frac{-2kqzxw_0^2}{k^2w_0^4 + 4z^2}\right) \right). \quad (5.6)$$

The above expression can be simplified assuming we're looking at distances much less than a Rayleigh range $z \ll z_R = kw_0^2/2$, and assuming $qx \ll kz$ (which is easily true in our experiment except for locations very close to $z = 0$). One arrives at

$$I(q) - I(-q) = \frac{2i\sqrt{2\pi}}{\sqrt{\frac{2}{w_0^2} - \frac{ik}{z}}} \exp\left(\frac{-w_0^2(k^2x^2 + q^2z^2)}{k^2w_0^4 + 4z^2}\right) \exp\left(\frac{4ikx^2z - ikq^2w_0^4z}{2k^2w_0^4 + 8z^2}\right) \\ \sin\left(\frac{2qxxk^2w_0^4}{2k^2w_0^4 + 8z^2}\right) \quad (5.7)$$

Or, written relative to the overall Gaussian profile of the beam:

$$\frac{|I(q) - I(-q)|}{|I(0)|} = \exp\left(\frac{-iq^2w_0^2}{4} \frac{z}{z_R}\right) \exp\left(\frac{-q^2w_0^2}{4} \frac{z^2}{z_R^2}\right) \sin(qx). \quad (5.8)$$

We now have a nice expression for $u(x, z)$:

$$u(x, z) = u_0 e^{-x^2/w_0^2} (1 + \delta) \quad (5.9)$$

where

$$\delta = \exp\left(\frac{-iq^2w_0^2}{4} \frac{z}{z_R}\right) \exp\left(\frac{-q^2w_0^2}{4} \frac{z^2}{z_R^2}\right) \sin(qx) \quad (5.10)$$

and $z_R = \pi w_0^2/\lambda$ is the Rayleigh range of the overall Gaussian beam. The first exponential term in Eq. 5.10 is a phase term that causes the modulation to oscillate from amplitude modulation to phase modulation and back, over a characteristic length scale $4z_R/(q^2w_0^2)$. Note that this exactly matches the Talbot length for Talbot patterns on laser beams [90]. The second exponential term captures exponential decay of the modulation amplitude with a characteristic length scale $4z_R/(qw_0)$, similar to a Rayleigh range for the effect.

This model also can be used to describe a square wave modulation on the beam as well. A square wave with spatial frequency q is a Fourier series summation over components of the sinusoidal terms (odd-multiples of q) in Eq. 5.10:

$$\delta_{square} = \sum_{n=1,3,5..} \frac{4}{n\pi} \delta(nq) \quad (5.11)$$

where $\delta(nq)$ is to be interpreted as δ from Eq. 5.10 where q replaced by nq . In this case the overall beam profile is given by

$$u(x, z) = u_0 e^{-x^2/w_0^2} (1 + \delta_{square}) \quad (5.12)$$




Figure 5.2: (blue) The intensity profile of light 2 m after passing through a $100 \mu\text{m}$ diameter aperture. (orange) A Gaussian approximation to the Airy pattern that has equal area under the curve.

Both the sinusoidal model and the square wave model have two important features. First, because of the first exponential term in Eq. 5.10, higher spatial frequency modulations change rapidly between amplitude modulation and phase modulation along the propagation axis. This leads to important physical effects in atom interferometers due to large gradients along the propagation axis. Secondly, the second exponential term in Eq. 5.10 causes very high frequency spatial noise to decay away before reach atoms in the interferometer. This smoothing of the beam helps limit effects from the highest spatial frequency components, which also have the highest curvature and gradients and hence the largest potential for systematic effects. Experiments with larger propagation distances between the last optic and the atom sample will have a lower cutoff spatial frequency above which spatial perturbations on the beam will be filtered out before reaching the atoms.

Dust on optics

Another useful model solves for the propagation of a beam after passing dust particles that diffusely scatter incident light. The dust effectively creates a circular 'hole' in the beam. Propagation of this hole can equivalently be thought of as the electric field of the overall Gaussian laser beam *minus* the electric field of a plane wave passing through a circular aperture.

Diffraction of light through a circular aperture is well-known to give Airy rings on a beam in the near-field. In the far-field limit, the intensity distribution of a plane wave passing through a circular aperture is given by [84]

$$I(r) = I_0 \left[\frac{2J_1(\beta)}{\beta} \right]^2 \quad (5.13)$$

where $J_1(\beta)$ is the Bessel function of the first kind, and $\beta = \pi dr / (\lambda z)$ where d is the diameter of the dust particle/aperture, r is the distance away from the axis of the beam, and z is the




Figure 5.3: Simulated propagation of a 6.2mm waist beam with roughly 20 dust particles, all $50\ \mu\text{m}$ diameter, in the region $(-w_0, w_0)$. The beam model is shown at a) at the plane of the dust, b) 0.4 m from the dust, and c) 1.6 m from the dust. For anyone that has imaged a large laser beam with a CCD camera, this will look eerily familiar.

propagation distance. In our experiment the beam propagates over 2 m before reaching the atoms. As a Rayleigh range, 2 m corresponds to a dust size of $730\ \mu\text{m}$ - this means that dust particles of size $d \ll 730\ \mu\text{m}$ can be treated as ‘far-field’.

In the far-field, the beam profile is very well approximated by a Gaussian beam. Fig. 5.2 shows in blue the intensity distribution of light 2 m after passing through a $100\ \mu\text{m}$ diameter aperture. The pattern is very nearly Gaussian in the far-field - a Gaussian waist of

$$w_0 = 0.9 \frac{\lambda z}{d} \quad (5.14)$$

fits the curve very closely fits the Airy pattern and also contains the same total power integrated across the beam [104].

The dominant effect on the beam intensity profile comes from interference between the Gaussian field from the main laser beam interfering with (minus) the electric field emanating from the location of the dust, which is well approximated by a Gaussian with a waist given by Eq. 5.14. The different wavefront curvatures between the two Gaussians interfere to create a ring structures on the propagation.

The Monte Carlo models dust by subtracting the Gaussian electric field of the above dust model from the main Gaussian beam. The amplitude on the dust electric field contribution is proportional to the amplitude of the initial Gaussian beam at the location of the dust, and the electric field of the dust originates at the location of the dust particle. The model allows the user to enter an array of dust locations, where all dust particles are assumed the same diameter (though this assumption could easily be changed), and the total field sums over the contribution from all pieces of dust. An example of a beam propagation with randomly placed dust is shown in Fig. 5.3.

5.3.2 Physical effects from laser beam intensity inhomogeneities

Equipped with models for how laser beams propagate, how do these beams affect atoms in the interferometer? Three primary ways are included in the model: 1) the local lattice depth changes, which changes the Schrodinger equation dynamics, 2) momentum is transferred to

the atoms which changes the resulting trajectories, and 3) these both result in phase changes on the wavefunctions which are measured at the end of the interferometer.

Effects from changes in lattice depth

The most direct effect of inhomogenous beam intensity is a spatially varying lattice depth. For resonant processes such as Raman transitions, Bragg diffraction, or single photon transitions, differences in laser intensity will directly affect the Rabi frequency of the process. This leads to amplitude and phase shifts on the diffracted wavefunctions - for example during Bragg diffraction, these phase shifts are commonly referred to as ‘diffraction phase’ as discussed in Chapter 2. Higher-order processes are more sensitive to deviations in intensity - for example, n^{th} order Bragg diffraction is typically n -times more sensitive to intensity variations than two-photon Raman transitions or first-order Bragg diffraction.

For adiabatic processes such as Bloch oscillations, changes in laser intensity affect the probability of tunneling from the lattice, as also cause phase shifts on the surviving wavefunctions. Chapter 2 showed that in the weak-lattice limit⁵, Landau-Zener tunneling from the lattice can be described by a probability of survival per Bloch oscillation of [65]:

$$P_{\text{surv}} = 1 - e^{-\frac{2\pi\Omega^2}{4r}} \quad (5.15)$$

where Ω is the lattice depth, and r is the frequency ramp rate used experimentally to accelerate the lattice. This equation was also presented in Eq. 2.38. The fractional change in survival probability per fractional change in lattice depth is given by

$$\frac{\Omega}{P_{\text{surv}}} \frac{dP_{\text{surv}}}{d\Omega} = \frac{\pi\Omega^2}{r} \frac{e^{-2\pi\Omega^2/(4r)}}{1 - e^{-2\pi\Omega^2/(4r)}} \quad (5.16)$$

This quantity saturates at a maximum value of two as the lattice depth $U_0 \rightarrow 0$. On the other extreme, the quantity quickly goes to zero in the limit $U_0 \gg \sqrt{r}$. The correlation between lattice depth and survival probability is therefore stronger at lower lattice depths, and is suppressed at higher lattice depths. Two recent measurements of the fine structure constants used this scaling to bound systematic effects from spatial intensity inhomogeneities [67, 55].

Note that the lattice depth is proportional to the product of the single-photon Rabi frequencies from the up-going and down-going beams. In experiments with retro-reflected beams such as ours, the lattice depth at a given location depends on the laser beam profile at two different propagation distances. For experiments using two independent fiber ports for the up-going and down-going beams, the propagation of beams from each fiber port will need to be independently studied and the lattice depth will depend on both beam profiles.

Momentum transferred to atom

Next, we consider the momentum imparted on the atoms while interacting with the laser, There are two main effects - first from the direct momentum transfer from photon absorption

⁵ $\Omega \ll 8E_r$

and emission, and second from gradients in the AC stark shift from the laser beam which leads to dipole forces on the atom.

When photons are directly exchanged with the atom, the atom receives a momentum $\Delta\vec{p} = \hbar\vec{k}$ for each photon transferred. This occurs during velocity-sensitive Raman transitions, Bragg diffraction, Bloch oscillations, and resonant single-photon transitions in atoms with long-lived excited states. The k-vector of the laser beam is defined locally as the gradient of the phase of the electric field, $\vec{k} = \vec{\nabla}\phi$.

A general expression for the local variation of the k-vector along the propagation axis δk can be derived according to Helmholtz propagation of the beam [6]:

$$\delta k = -\frac{1}{2} \left\| \frac{\vec{\nabla}_{\perp}\phi}{k} \right\|^2 + \frac{1}{2k^2} \frac{\Delta u}{u} \quad (5.17)$$

where u is the local amplitude of the electric field, ϕ is the local phase of the electric field, and $\vec{\nabla}_{\perp}\phi$ denotes the transverse phase gradient. The first term describes a rotation of the k-vector away from the propagation axis, which causes the projection of the k-vector onto the propagation axis to get smaller. The second term is the source of the well-know Gouy phase shift on-axis of a Gaussian laser beam - the k-vector is shorter even at the center of the beam. Compressing electric field amplitude into a smaller transverse region requires a spread in transverse momentum, and conservation of momentum therefore requires a shortening of the k-vector. There is a very nice paper describing this intuition behind the Gouy shift [27].

AC Stark shifts are another important effect in the interaction between atoms and the interferometry laser beam. The AC Stark shift changes the internal energy by an amount

$$U_{Stark} = \hbar \sum_i \frac{\Omega_i^2}{4\delta_i} \quad (5.18)$$

where Ω_i and δ_i are the single photon Rabi-frequency and detuning from the excited state, respectively, for the i^{th} beam. Each frequency of light contributes its own AC Stark shift.

When atoms are interacting with a standing wave of light, the wavefunction tends to the low-energy region of the lattice, so the average light shift may differ from the above equation, which assumes a uniform light shift spatially. This is why the α project has historically used blue-detuned laser beams - the atoms are trapped at the low-intensity region of the standing wave, which minimizes single photon scattering and light shift effects. This can be quantified using the dimensionless Lamb-Dicke factor $\eta = (4\Omega/E_r)^{1/4}$, where the reduction in single photon scattering is given by η^2 [70]. I have assumed that the same scaling is true for the light shift, since both are proportional to the light intensity. The code currently takes this into account for Bloch oscillations but not Bragg diffraction, since most of the light shift in the interferometer as a whole is coming from Bloch oscillations. Note that this only applies to optical lattices interacting with the atom - other off-resonant optical lattices will still contribute the full light shift.

Gradients in the AC Stark shift lead to forces on the atoms, identical to the forces in optical dipole traps. This force is given by:

$$F_{Stark} = \vec{\nabla}U_{Stark} \quad (5.19)$$

For a general interaction with a laser, the atom's momentum is therefore changed by an amount

$$\Delta\vec{p} = \int F_{Stark} dt + \hbar \sum_i \vec{k}_i \quad (5.20)$$

where the Stark force is integrated over the duration of the pulse, and each photon's k-vector is added to the atom's momentum.

Effect on interferometer phase

The previous section outlines how the laser beam can impart momentum to atoms, which then modifies the atoms' trajectories. There are three ways that the atomic trajectories then lead to differential phase shifts in the interferometer.

First, atoms acquire phase from kinetic energy and gravitational energy based on their position and velocities. This results in a free evolution phase (to first order):

$$\phi_{free} = \int L dt = \int dt \left(\frac{p^2}{2m} - mgx \right) \quad (5.21)$$

where L is the classical Lagrangian, x and p are the position and momentum of the atom at time t , m is the mass of the atom, and g is the local gravitational acceleration. When atoms receive a modified k-vector kick δk , their resulting positions and momenta are modified, leading to different kinetic and potential energy phases and a different free-evolution phase shift in the interferometer.

Next, each time an atom absorbs or emits photons from the laser beam, the local phase of the laser beam is imprinted onto the wavefunction. For Raman transitions or Bragg diffraction, which are two photon transitions, the phase shift received at time t is given by (to first order):

$$\phi_{laser} = n(k_1 z(t) - \omega_1 t) - n(-k_2 x(t) - \omega_2 t) \quad (5.22)$$

where k_1, ω_1 are the values of the k-vector and frequency of the first photon, and k_2, ω_2 are the values of the k-vector and frequency of the second photon. For Raman transitions, $n = 1$, and for higher-order Bragg diffraction, n will be greater than 1. This was also discussed in Chapter 1.

Eq. 5.22 is an appropriate treatment of a plane wave, however when laser wavefronts are not plane waves the laser phase more generally is given by the local phase argument of the electric field:

$$\phi_{laser} = n(\phi_1(x, y, z) - \omega_1 t) - n(-\phi_2(x, y, z) - \omega_2 t) \quad (5.23)$$

where $\phi_i(z) = \arg(E_i(x, y, z))$.

In regards to modified atomic trajectories from modulated laser beams, there are two ways in which laser phase contributes to the final interferometer phase. First, when atoms receive a modulated momentum kick $k + \delta k$, the atomic trajectories are slightly different and the atoms therefore sample the laser phase at different points in space. To first order this change in laser phase can be captured by Eq. 5.22. Second, the modulation of the

laser beam itself causes a modulation in the phase-front of the laser that deviates from Eq. 5.22. Directly taking the argument of the laser electric field as in Eq. 5.23 at the location of the unperturbed atomic trajectories captures this second effect. Both effects are important in accurately modeling the resulting phase shifts in the interferometer. The Monte Carlo treats this all generally, where the momentum kicks are vector quantities that affect the vector velocities of the atomic trajectories, and the laser electric field is sampled at the exact location of the atoms at the time of each pulse.

The last source of phase shifts come from the separation phase at the closing of an interferometer. Separation phase is a well-known effect that appears when an interferometer's trajectories do not perfectly close at the time of the final beamsplitter pulse, and is given by⁶ [43]:

$$\phi_{sep} = \bar{\mathbf{p}} \cdot \Delta \mathbf{x} \quad (5.24)$$

where $\bar{\mathbf{p}}$ is the average momentum between the two output ports being considered, and $\Delta \mathbf{x}$ is the relative displacement between the two interferometer arms at the time of the final closing pulse in the interferometer. Note that there is a strange sign convention in defining separation phase this way. $\Delta x = x_l - x_u$ is the position of the center of the lower wavepacket minus the upper wavepacket, which is opposite to how we typically define differences in atom interferometry calculations [43].

While separation phase has been considered in many contexts, it has not been considered in relation to the local k-vector variations from intensity inhomogeneities in a laser beam. For a smooth Gaussian laser beam profile with a Rayleigh range much longer than the length scale of the experiment, the k-vector deviations given by Eq. 5.17 are very nearly identical between different interferometer pulses, so the interferometer closes nearly exactly. However, when there is spatial intensity noise on a laser beam, the k-vector changes along the propagation occur on a much shorter length scale as described by Eq. 5.10. Atoms receive different δk momentum kicks from different laser pulses, which leads to the interferometer not closing and a resulting separation phase.

Taken together, the dominant phase shift in the interferometer comes from a correlation between the δk momentum kicks an atom receives and the probability that that atom contributes to the final contrast in the interferometer. This δk kick affects free evolution phase, laser phase and separation phase - all three need to be considered to accurately capture the effect.

For any small-scale modulation on a Gaussian laser beam, the average δk from the modulation across the laser beam is the zero - the effects are oscillatory about δk from the large-scale Gaussian curvature. This averaging doesn't occur perfectly across the atom cloud because the contribution of each atom to the total contrast of the interferometer also depends on the intensity of the beam at each pulse. As noted in [6], Bloch oscillations cause preferential tunneling in lower intensity regions of the beam, leaving a correlation between the surviving atoms and the δk received from the modulated beam during Bloch oscillations. Similarly, the local Rabi frequency and diffraction phases [24] of resonant processes are affected by the intensity modulation, leaving a correlation between the contribution to total contrast and the δk received. Last, non-overlapping trajectories, which cause loss of

⁶Eq. 1.4 is identical but it is reproduced here for continuity

contrast, and the related separation phases, similarly contribute to the correlation between measured contrast the δk received. The overall effect on the interferometer's differential phase is therefore a result of correlations between a number of different processes.

5.3.3 From wavefunction amplitude to ellipses

The above describes how to get complex wavefunction amplitudes and classical wavefunction positions for all of the various interferometer trajectories at the end of the interferometer. How do you then turn this information into a point on an ellipse?

First, when wavefunctions that are being interfered that don't overlap exactly, there is a resulting separation phase and loss of contrast. The separation phase has already been mentioned in e.g. Eq. 5.24, and this equation is used directly in the Monte Carlo code to add separation phase to the atomic wavefunctions.

The length scale relevant for loss of contrast is the Heisenberg limited position spread associated with a given velocity spread. However, one might be concerned about the fact that we don't necessarily have Heisenberg-limited atomic wavefunctions in our experiment, like I was a few years ago. Similarly, even if you started with a Heisenberg-limited atomic wavepacket, time evolution of the wavepacket causes the wavefunction to expand. The following argument proves that any amount of time t later, no matter how large the wavefunction extent becomes, the overlap integral between two identical wavefunctions will still only be non-zero over a length scale of the Heisenberg limited position spread of the wavefunction.

Let's derive this explicitly starting with a normalized Gaussian wavepacket at time $t = 0$:

$$\psi(x, t = 0) = \left(\frac{1}{\pi\sigma_x}\right)^{1/4} \exp\left(\frac{-(x-x_0)^2}{2\sigma_x^2}\right) \quad (5.25)$$

where σ_x is the 1σ position width of the atomic population distribution, which it is related to the 1σ momentum width σ_p by the Heisenberg uncertainty principle $\sigma_x = \hbar/(2\sigma_p)$. The free particle time evolution of the Gaussian wavefunction in Eq. 5.25 is analytically solvable:

$$\psi(x, t) = \left(\frac{1}{\pi\sigma_x}\right)^{1/4} \left(1 + \frac{i\hbar t}{m\sigma_x^2}\right)^{1/2} \exp\left(\frac{-(x-x_0)^2}{2\sigma_x^2\left(1 + \frac{i\hbar t}{m\sigma_x^2}\right)}\right) \quad (5.26)$$

The quantity we're interested in is the overlap integral between two wavefunctions that are displaced by an amount Δx . We obtain:

$$\langle\psi_1(x + \Delta x, t)|\psi_2(x, t)\rangle = \exp\left(\frac{-\Delta x^2}{4\sigma_x^2}\right) \quad (5.27)$$

which is found to be independent of time. The above convention assumes ψ_1 and ψ_2 are parts of the same atomic wavefunction so that they interfere. The purpose of this exercise is to show that the overlap integral between two atomic wavefunctions depends only on the initial Heisenberg limited position spread of the atom. Even if the atomic wavepacket grows in size during free evolution, the length scale relevant for the overlap integral remains the same.

An equivalent but simpler way to see this is to show that the overlap integral at time $t = 0$ doesn't change for later t 's.

$$\begin{aligned} \langle \psi_1(x + \Delta x, 0) | \psi_2(x, 0) \rangle &= \langle \psi_1(x + \Delta x, 0) | U^\dagger(t) U(t) | \psi_2(x, 0) \rangle \\ &= \langle \psi_1(x + \Delta x, t) | \psi_2(x, t) \rangle. \end{aligned} \quad (5.28)$$

This uses a resolution of the identity as the unitary time evolution operators $U(t)$ to show in operator form what was shown in Eq. 5.27.⁷ One can make this general to three-dimensions - you find the same expression for each dimension independently, where each axis can have its own position/velocity spread. This general expression is used in the Monte Carlo code to treat the transverse overlap separately from the on-axis overlap, since the on-axis temperature is much lower due to velocity selection pulses.

For a given interferometer with wavefunctions ψ_{u1} and ψ_{u2} interfering in the upper output port, and ψ_{l1} and ψ_{l2} interfering in the lower output port,

For the upper interferometer, the Monte Carlo code calculates the populations in the upper and lower output ports to be

$$P_{ui} = a |\psi_{ui1} + e^{i(\phi_{\text{diff}} + \phi_c + \phi_{\text{sep}})} \psi_{ui2}|^2 + (1 - a) (|\psi_{ui1}|^2 + |\psi_{ui2}|^2) \quad (5.29)$$

where

$$a = \exp\left(\frac{-\Delta x^2}{4\sigma_x^2}\right) \quad (5.30)$$

is the overlap integral between the two wavefunctions, as derived in Eq. 5.27. P_{ul1} , for example, denotes the population in the upper interferometer, lower output port, from the first wavefunction being interfered. ϕ_{diff} is a user-controlled differential phase that can be used to open the ellipse. ϕ_c is the common mode phase to both upper and lower interferometers, and is a random number between 0 and 2π in the simulation to mimic strong vibration noise. ϕ_{sep} is the separation phase.

The population in the lower interferometer output ports is determined by the same equations but without adding the phase shift ϕ_{diff} to the second wavefunction. As expected, the common mode phase is common to both the upper and lower interferometers, but the differential phase differs.

Eq. 5.29 in words says that the fraction of the wavefunctions that is coherently interfering is used to calculate the interference contribution to the population in the output port. The non-overlapping fraction of the wavefunctions don't interfere, but only contribute their individual populations to the output ports. For example, if the wavefunction trajectories didn't overlap at all, there would be no coherent interference term and the population in each output port would be roughly the same, assuming the Bragg diffraction $\pi/2$ pulses were splitting the population near 50%. Similarly, if the wavefunctions overlap perfectly, you can get 100% contrast in the interference between output ports because the second term vanishes.

The simulation also optionally includes the population from waste ports, which are illustrated in Fig. 5.4. The population in the waste ports is added incoherently to the output

⁷I asked Holger about this years ago, and apparently he asked his undergrad quantum mechanics class this exact question on a homework assignment. A few students provided the above answer with unitary operators - that's where I got this nice solution from.

ports, which typically limits the maximum contrast to around 50% (see Brian’s thesis Section 5.1).

The actual point (x, y) on the ellipse is calculated via:

$$x = \frac{P_{lu} - P_{ul}}{P_{lu} + P_{ul}} \quad (5.31)$$

and

$$y = \frac{P_{uu} - P_{ul}}{P_{uu} + P_{ul}} \quad (5.32)$$

where e.g. P_{ul} indicates the population in the upper interferometer of the SCI, lower output port.

One important thing to note here is that a single atom is creating a point on an ellipse. This is in contrast to our experimental reality where an atom is projected into one of the output ports in an interferometer - we can’t just split atoms apart. The Monte Carlo doesn’t include atom number shot noise because we have direct access to the wavefunction information and can therefore bypass this extra noise source. As far as I know, this only makes the statistical averaging of the Monte Carlo code better without affecting any of the results.

The Monte Carlo code also (optionally) averages many atoms to create a single ellipse point. This 1) reduces the shot-to-shot noise on an ellipse for very noisy beam profiles, but more importantly 2) it mimics the experimental scenario we have where a single shot of the experiment averages thousands of atoms to create a single point on the ellipse.

To average many atoms into a single ellipse point, we perform a weighted average of the individual ellipse points, where the weighting is done by the total population in the upper and lower interferometers. The averaged ellipse point $(x_{\text{tot}}, y_{\text{tot}})$ becomes

$$x_{\text{tot}} = \frac{1}{P_{\text{tot},l}} \sum_{i=1}^{N_{\text{atoms}}} (P_{i,ll} + P_{i,lu})x_i \quad (5.33)$$

where $P_{i,ll}$ is the i^{th} atom’s population in the lower interferometer, lower port, and $P_{\text{tot},l} = \sum P_{i,ll} + P_{i,lu}$ is the total population of all of the atoms in the lower interferometer ports. An analogous equation is used for the y-axis using the upper interferometer populations. This weighted average is set up to mimic what happens experimentally with our data. For example, some atoms tunnel more during Bloch oscillations compared to other atoms depending on the local lattice depth that each atom sees. The atoms that tunnel more won’t contribute as much to the total measured interferometer signal.

Note that this doesn’t include effects of the probability of detecting an atom based on the experimental detection system - in the future, you may want to add this in to simulation e.g. the effects of a pinhole filter in front of the detection photodiode.

5.4 Overview of the Monte Carlo code

Section 5.3 above describes all of the physics that is wrapped into the Monte Carlo model. This section describes how the actual code is organized to incorporate all of this physics

while still being very fast and user-friendly. The code is located in the Mueller group GitHub organization with the repository name ‘alpha_monte_carlo’. The repository is private to the Mueller group, but for anyone outside of our group that is interested in the code please contact someone in our lab and we can arrange some code sharing.

For a very spatially noisy beam, it can require upwards of a billion atoms to integrate a single ellipse’s statistics down to 1 mRad level error bars, since there is so much phase noise atom-to-atom. To make a useful plot, one needs to sweep some parameter(s) and simulate many ellipses, potentially on the order of tens of billions of atoms. The code needs to be fast enough that this can be done in a reasonable amount of time, say a day or less, so that a user can iterate on the code to properly study a problem.

In addition, the code must be functional and user-friendly enough that the user doesn’t waste a bunch of time coding/ copy pasting, and so that the code is easily understandable and passed down to future users. It must be documented well enough that a new user can successfully learn how to use it significantly faster than it would take them to write the package from scratch.

The Monte Carlo is doing fairly well on the speed front, with some possible speed-ups discussed in section 5.6.1 that would require a large overhaul of the code. Compared to the previous Monte Carlo packages, the largest speed-up improvements in this code are from using pre-computed integration tables and from purchasing a simulation computer that can run 128 processes in parallel.

The old Monte Carlo used for the α data analysis never fully included Bloch oscillations because it was computationally too expensive. Richard’s code never actually had a functioning Bloch simulation, it just put the wavefunction amplitude in a different momentum state (with no phase shifts or loss of amplitude), then continued simulating the final two Bragg pulses.⁸ Chenghui had a working simulation of Bloch oscillations but the integration was too slow to use in the ppb-level systematics analysis. Instead he used a look-up table of the surviving wavefunction population after Bloch oscillations, but this didn’t include phase shifts from the Schrodinger equation integration. The Bloch simulations in this new Monte Carlo are similarly slow to Chenghui’s, since they are both written with similar integration algorithms in python compiled down to C. However, by pre-computing the results of Bloch integration in the new code package, the Monte Carlo itself can run much much faster after the precompute tables have been generated.

In terms of functionality, the structure of the Monte Carlo code is designed to give the higher-end user the most amount of control possible with the fewest lines of code. A good example of this is the Sweep and Sweep2D classes – a user could create a two-dimensional sweep over parameters with billions of atoms, and plot and save the outputs, all with only 10 or so lines of code. The params dictionary is also central to the code: a single dictionary of all simulation parameters is passed all around the code, allowing the user to easily change any parameter in the simulation (at the time of writing there are over 100 parameters). Parameters almost never need to be explicitly passed between objects or methods which makes for much cleaner coding. I was initially worried that passing around a large dictionary like this would cause memory or speed issues, but operations related to the params dictionary

⁸The code was primarily used to study diffraction phase effects the Bragg pulses, so for this application it was acceptable to omit Bloch oscillation dynamics.




Figure 5.4: Ordering of the wavefunctions as kept track of within an atom object during an SCI interferometer. The lower trajectory is always first in order. Dashed vertical lines indicate Bragg diffraction pulses, solid vertical lines indicate Bloch oscillations. All wavefunction trajectories are kept track of, including ‘junk’ ports that do not participate in coherent interference. Multiple numbers in the final step indicate overlapping wavefunctions - these are the wavefunctions that are manually interfered at the end of the sequence. The time step simple indicates where in the sequence the ordering corresponds to.

are very far from being the bottleneck in the code so this was a non-issue.

The class structure breaks the code into intuitive buckets, although there are some grey areas between classes where I had to put some code in weird spots – if you dive deep enough into the code you’ll find examples of this. For the most part though, the code divides up very nicely between classes.

As for a coding language, I wanted to use Julia as the core language because it is naturally so fast. In the end we decided to use Python because there was only a factor of ≈ 2 difference in the integration speed, and Python is much more documented and higher-level developed. Similarly, we made some basic integration code in C which was essentially the same speed as Julia, so we didn’t actually need to code directly in C. Matlab was also an option for coding languages since it’s inherently very fast, but when given the choice it’s much better to stick to an open-sourced language.

5.4.1 Introduction to classes

Broadly, the Monte Carlo is designed with the following logic from lowest-level to highest-level:

Atom objects are used to capture all dynamics of a single atom throughout the interferometer. They are the heart of the Monte Carlo code. An atom object contains the values of the wavefunction along different branches of the interferometer, as well as the corresponding position, velocity, and local laser electric field information along the different branches. Fig.

5.4 gives a visual representation of what the ordering of wavefunctions looks like in an SCI interferometer.

Time Evolution objects operate on Atom objects. They take in an atom object, time evolve the different trajectories through the Schrodinger equation and other classically-treated updates, and spit out a new wavefunction object (or two new wavefunction objects, in the case of Bragg diffraction). We have a general TimeEv class, and then Bragg, Bloch, and FreeEvolution child classes for our SCI interferometer. If you ever want to simulate more general interferometers in the future, you can likely just make a new TimeEv child class, or even build off of the existing classes.

The next higher-level class in the code is the Ellipse class, with the closely related EllipsePoint class. An EllipsePoint is a collection of atoms simulated with the same initial conditions, then averaged together in a contrast-weighted average. This mimics one experimental ‘shot’ in our experiment. The Ellipse object simulates a number of EllipsePoints, and then these EllipsePoints are fit to an ellipse in order to extract fit parameters such as differential phase and contrast.

The real power of the Monte Carlo, however, is in being able to sweep parameters and see how the differential phase and contrast are affected. The Sweep class creates a very easy interface to generate a 1D or 2D sweep over parameters, simulate an Ellipse for each parameter, and then plot the results and save data to an HDF5 file. The Sweep class has proved extremely useful in studying beam-related systematics, and will continue to be very useful in studying the relevance of all parameters in the Monte Carlo.

There are some other useful classes that will be described in the subsequent sections. The Beam class returns the local electric field of the laser beam, and the related analysis functions calculate a number of effects locally from the electric field: δk and the 3D k-vector, laser phase, lattice depth, light shifts, and dipole forces. One of the most useful functions in the code is the params() function - it generates a params dictionary of all parameters in the simulation, each of which can be over-ridden by the user. There is also an engine for converting quantities between dimension-ful and dimension-less units. This is implementation of the dimesionless unit system described in Chapter 2.

5.4.2 Atom Class

With that background, let’s look a little more in depth at each of the classes.

‘Atom’ objects track the relevant information of the atomic state along each branch of the interferometer by storing information as (pre-allocated) attributes of the object. The most important attribute is the wavefunction values themselves, but other information such as the local position and velocity of the atom along that trajectory are also stored. In addition, I also try to store any information that requires a larger 2D array. For example, the code stores the value of the laser electric field at 14 different locations - the location of the atom and a point in each of the 6 directions from that point, as well as the same for the retro-reflected beam. These are all needed to take derivatives of the lattice depth, light shift and laser phase. These 14 electric field values are stored separately for each branch of the interferometer (up to 16 trajectories), so I pre-allocate a 16x14 array.

Data is overwritten as the atom is evolved in order to conserve memory allocations. Previously, I used to have a “Wavefunction” class, which stored the state of a single branch

of the interferometer, then this object would be split and copied to new objects during time evolution. It turned out copying the Wavefunction object was very slow, so I re-wrote the code to instead store all information directly on the Atom object and update data by overwriting pre-allocated arrays. It was very convenient to store all sorts of information on the Wavefunction objects, but the current use of atom objects makes the code much simpler and streamlined.

Direct integration of the atomic wavefunction through the Schrodinger equation requires knowledge of the wavefunction over the entire discrete momentum basis. This was discussed in detail in Chapter 2. In contrast, the standard Monte Carlo mode of operation only stores the value of the wavefunction on the momentum state of the interferometer branch being tracked. All other wavefunction values are dropped, since they don't directly interfere when closing the SCI interferometer. This functionality is controlled by the Atom 'type' attribute. An 'SCI' type will drop the neighboring momentum states, as discussed above, and is the default for full SCI simulations/ ellipse plotting. An Atom with 'integration' type will store the full wavefunction information, which is typically only used for debugging integration or one-off scripts looking at physics within a single Bragg/Bloch pulse. Note that Wavefunction amplitude on neighboring momentum states causes parasitic interferometers [66], so this systematic effect is being ignored currently. In the future, the Monte Carlo code will need to be generalized to include parasitic interferometers, but this should only require small modifications to the Atom class and TimeEv classes.

Each operation in an interferometer (Bragg diffraction and Bloch oscillations) propagates each wavefunction based on the local electric field of the laser at the respective interferometer branch. The atom's stored wavefunction values are then overwritten. For Bragg diffraction, the two output states are stored so that each branch of the wavefunction is effectively 'split' in two on each pulse, which are then propagated forward independently. It's necessary to track wavefunction values separately along each interferometer arm in order to use local laser electric field information during each pulse, and also to only interfere spatially overlapping trajectories at the end of the interferometer.

5.4.3 TimeEvolution Classes

At their core, the TimeEv classes propagate atoms through the Schrodinger equation. For the FreeEvolve class, this is very simple - free-space momentum states are diagonal in the kinetic energy operator, so the wavefunctions just get a phase shift proportional to the kinetic energy of each state. For Bragg diffraction and Bloch oscillations, the Schrodinger equation needs to be numerically integrated. If using precomputed tables, the results of numerical integration are looked up from a look-up table instead of integrating each time. Precompute tables are discussed in more detail in their own section below. If not using precomputed tables, the Schrodinger equation is numerically integrated.

The numerical integration is implemented with Numba in order to get extremely fast C-level performance. The numerical integration is just a giant for-loop with linear algebra inside, and Numba is designed for exactly this application. The integration is performed with a Crank-Nicolson (CN) algorithm that gives an unconditionally stable conservation of probability. There are many types of algorithms to choose from for numerical integration, but very few conserve probability - the CN method is therefore very popular for Schrodinger

equation simulations [86].

Numerically integrating the Schrodinger equation requires iteratively solving for the wavefunction at the next discrete time-step. The Schrodinger equation reads:

$$i\hbar \frac{\partial \psi}{\partial t} = \hat{H}\psi(t) \quad (5.34)$$

When taking a small discrete time-step such that the Hamiltonian \hat{H} is assumed independent of time, the problem reduces to solving:

$$\psi(t + dt) = e^{-i\hat{H}dt/\hbar}\psi(t) \quad (5.35)$$

where the exponential operator is a unitary operator. Since the Hamiltonian contains position and momentum operators, it's not straightforward to compute the matrix exponential. Moreover, if you Taylor-expand this operator, you end up with something non-unitary. $1 - i\hat{H}dt$ has a norm slightly more than 1, so probability diverges when integrating.

If you instead apply the operator symmetrically (following the derivation in reference [86]):

$$e^{i\hat{H}dt/(2\hbar)}\psi(t + dt) = e^{-i\hat{H}dt/(2\hbar)}\psi(t) \quad (5.36)$$

then you can Taylor expand each side and arrive at a unitary integration algorithm, the CN algorithm:

$$\left(1 + \frac{i\hat{H}dt}{2\hbar}\right)\psi(t + dt) = \left(1 - \frac{i\hat{H}dt}{2\hbar}\right)\psi(t) \quad (5.37)$$

This conserves probability because each side of the equation multiplies the wavefunction by a quantity with the same norm.

Typically, the down-side to the CN algorithm is that in order to solve for $\psi(t + dt)$, one needs to invert the matrix on the left-hand side of the equation. Fortunately for us, our matrices are tri-diagonal and tri-diagonal matrix inversion can be done in $O(3N)$ time instead of $O(N^2)$ time for a general $N \times N$ matrix. The Monte Carlo code uses the Thomas algorithm which is the fastest way to solve this problem numerically [81].

The Hamiltonian for atoms in an optical lattice was shown in matrix form in Chapter 2. Building from there, we can re-write Eq. 5.37 above as:

$$U_1\psi(t + dt) = U_2\psi(t) \quad (5.38)$$

The matrices U_1 and U_2 in dimensionless units are given by:

$$U_1 = \begin{pmatrix} \ddots & \ddots & \ddots & & \ddots & & \\ \ddots & \alpha_{-2} & \beta & 0 & & & \ddots \\ \ddots & \beta^* & \alpha_{-1} & \beta & 0 & & \ddots \\ & 0 & \beta^* & \alpha_0 & \beta & 0 & \\ \ddots & & 0 & \beta^* & \alpha_1 & \beta & \ddots \\ & \ddots & & 0 & \beta^* & \alpha_2 & \ddots \\ & & \ddots & & \ddots & \ddots & \ddots \end{pmatrix} \quad (5.39)$$

and

$$U_2 = \begin{pmatrix} \ddots & \ddots & \ddots & & \ddots & & \\ \ddots & \gamma_{-2} & \lambda & 0 & & & \ddots \\ \ddots & \lambda^* & \gamma_{-1} & \lambda & 0 & & \ddots \\ & 0 & \lambda^* & \gamma_0 & \lambda & 0 & \\ \ddots & & 0 & \lambda^* & \gamma_1 & \lambda & \ddots \\ & \ddots & & 0 & \lambda^* & \gamma_2 & \ddots \\ & & \ddots & & \ddots & \ddots & \ddots \end{pmatrix} \quad (5.40)$$

with α_m and γ_m defined as

$$\alpha_m = 1 + \frac{i(2m)^2 dt}{2} \quad (5.41)$$

$$\gamma_m = 1 - \frac{i(2m)^2 dt}{2} \quad (5.42)$$

and β and λ defined as

$$\beta = \frac{idtV}{2} \quad (5.43)$$

$$\gamma = -\frac{idtV}{2} \quad (5.44)$$

where V is the off-diagonal coupling matrix element in the Hamiltonian. m is the index of the basis state with momentum $p = 2m\hbar k$.

The CN algorithm implemented in the code solves this linear algebra problem with the Thomas algorithm iteratively at each time step. The coupling off-diagonal terms are updated at each timestep based on a numba function in the TimeEv class that defines the pulse shape based on the local electric fields from the beam model. Since off-diagonal elements (e.g. β and γ above) are constant across the matrix, only scalar arguments are passed

around. The solution of the matrix equation $U_1\psi(t+dt) = U_2\psi(t)$ is reduced to one forward and one backward for loop iterating through the ψ vectors - the algorithms for matrix multiplication (one for loop) and the Thomas algorithm (two for loops) are combined to reduce the total algorithm to two for loops. The α_m and λ_m vectors are only defined once since they don't change each time-step. The algorithm is about as stream-lined as possible using numba/python - to further improve integration speed, you would probably need to move to a different coding language.

The above summarizes the numerical methods used for integration of the Schrodinger equation. However, the TimeEv classes do more than just Schrodinger equation integration, and actually these extra functionalities are some of the major physics improvements over the old code. The user can turn on/off many different physical effects by changing boolean values in the 'prms' dictionary:

- Including/removing kinetic energy phase from time evolution
- Including/removing integrated light shift phase from the optical potential
- Gouy shift modified (3-D) k-vector from local amplitude curvature/ phase gradient of the laser electric field
- Momentum kicks from forces from light shift gradients
- Separation phase when closing the interferometer, from imperfect trajectories
- Phase from the modulation frequency, essentially cancels the total interferometer phase like what we do experimentally
- Include interferometer waste ports so that max possible contrast is 50%

In total, these boolean options give the user a ton of control over which physics you're including in the simulations versus intentionally ignoring. I've found this to be extremely useful, even essential, in helping disentangle which physical processes cause which effects in the ellipse phase and contrast.

5.4.4 Ellipse and EllipsePoint Classes

Moving upwards in the code, the Ellipse class is in charge of generating Ellipses. I was running into dramatic memory issues when running millions of atoms, so for memory reasons I split some of this work into the EllipsePoint class. This helps to minimize the data being stored from each Atom when generating an ellipse.

As noted in Section 5.3.2 above, each atom gives us a data point for an ellipse without shot-noise - this is opposed to what would happen experimentally, where a single atom must project into one of the two output ports. For conditions with lots of atom-to-atom noise, however, you still need to average over many atoms to create a single point on the ellipse, otherwise the ellipse data will be too noisy to fit. This is also coincidentally⁹ the same as what we do experimentally with $O(10,000)$ atoms per experimental shot.

⁹not a coincidence

The code generates some number of `EllipsePoints`, each of which describe a single point on the ellipse as the ensemble-average of $O(10000)$ atoms. Each `EllipsePoint` is assigned a random common mode phase to mimic random vibration noise in our experiment. Once all `EllipsePoints` are simulated, the final ellipse data is least-squares fit¹⁰ to an ellipse equation. The differential phase and interferometer contrast are extracted from the fit, among other parameters, as well as the error bars on these quantities. Typically we are most interested in the differential phase of the ellipse, but contrast information is also important to match with experiment. Other outputs like diffraction phase or analytical Gouy shift can also be studied. The raw data and fitted ellipse can be optionally plotted with booleans in the `params` dictionary.

5.4.5 Sweep and Sweep2D Classes

Moving another level up in the code, the `Sweep` and `Sweep2D` classes allow the user to sweep one or two parameters in the `prms` dictionary and plot the resulting differential phases and contrasts as a function of the swept variable. This is a super powerful class that can very easy control a lot of machinery. Output data is stored to HDF5 files. For 2D sweeps, the data is loaded to the HDF5 file after each 1D sweep is completed, so you can check on the output data without waiting for the entire sim to complete. Note that once you're working with large simulations at this level, you want to think carefully about how to correctly parallelize the code for optimal performance. See Section 5.5 for further discussion.

5.4.6 Precompute Class

One major upgrade of this Monte Carlo versus old *alpha's* Monte Carlo is that the results of the Schrodinger equation integration are precomputed and stored in one HDF5 file. The user can create a large set of precompute tables once, and then forever be able to interpolate over the precompute table to lookup integration results. This enables a massive speedup of the simulations, especially for very large Bloch oscillation orders where convergence of a single integration could take days.

The `Precompute` class is responsible for everything related to the precompute tables. It generates the precompute data and saves to HDF5 files in an organized way (part of the beauty of using HDF5 files). The class also loads data from the precompute files, generates interpolation functions over this data, and uses the interpolation functions to generate the output wavefunction values. The `Bragg` and `Bloch` time evolution objects will look up results from a precomputed interpolation function if `params['precompute'] == True`.

I have used symmetries of the Hamiltonian to reduce the precompute tables to the minimum number of tables needed. With these symmetries, we reduce the necessary number of precompute tables by a factor of 2. I've used the following identities for Bragg diffraction:

- $\langle n | U_{Br,SF,v0} | 0 \rangle = \langle 0 | U_{Br,SF,-v0} | n \rangle$

¹⁰Stephanie Bie, an undergraduate working on our project for a year or two, wrote a package for Bayesian ellipse fitting. She also laid the initial groundwork for integrating the Bayesian ellipse fitting into the Monte Carlo code, on her own branches of the Github project. There's some low-hanging fruit here for anyone that wants to work on speeding up Stephanie's code and fully integrating it into the Monte Carlo code.

- $\langle -N - n | U_{Br,MF,v0} | -N \rangle = \langle N + 2n | U_{Br,MF,-v0} | N + n \rangle$
- $\langle -N | U_{Br,MF,v0} | -N - n \rangle = \langle N + n | U_{Br,MF,-v0} | N + 2n \rangle$

and the following for Bloch oscillations:

- $\langle -N | U_{Bl,MF,v0} | 0 \rangle = \langle N + n | U_{Br,MF,-v0} | n \rangle$

where $U_{x,y,z}$ indicates the unitary time-evolution operator for an entire pulse of type x , with y indicating single- or multi-frequency pulse, and a velocity offset from the lattice of z . The initial state is on the right and target state on the left, as in standard bra-ket notation. The Precompute class automatically uses these symmetries when looking up integration results.

The file organization in the precompute machinery is hard-coded but in a relatively neat way. One quirk of the code that's actually turned out to be somewhat useful is that the simulation computer will save the newly generated precompute tables inside the scripts folder instead of immediately overwriting the already existing precompute table in the precompute_data

folder. The precompute tables are backed up by git, so in principle this doesn't matter but it's still nice to have a chance to review the file before it overrides the existing functional one. The worst possible scenario would be accidentally overwriting a not-backed-up precompute table that took a month to run...

The main parameters that define the tables are the Bragg and Bloch order, however, the tables also depend on the Bragg pulse width (sigma of the pulse and the number of sigmas used in integration), the Bloch ramp duration and the Bloch ramp rate. If any of these parameters change, you need to rerun your integration tables which can take weeks.

Also, note that the tables are generated over a range of lattice depths and on-axis velocities with respect to the lattice, so there are fixed boundaries to the interpolation. If your lattice depth or on-axis velocity lies outside of these bounds, then the interpolation functions will extrapolate to get the final wavefunction amplitude and phase, and this extrapolation is most likely not accurate.

In Fig. 5.5, I show what some of the data stored in these tables looks like for single-frequency Bragg diffraction of order $n = 5$. The population plot shows the magnitude squared of the amplitude in the target state, and the phase plot shows the unwrapped phase difference between the initial and final states - this second plot is the definition of what we mean by 'diffraction phase' as discussed in Chapter 2. The y-axis is the range of lattice depths used in computing the precompute table, where the $\pi/2$ pulse intensity corresponds to $\Omega_{\pi/2} = 29.7E_r$. The x-axis is the initial (z) velocity offset from the lattice during integration. For a scale reference, the old alpha experiment used a one-sigma velocity width of 0.05 recoil velocities and fifth order Bragg diffraction, so you can see that their finite velocity width would have been enough to start causing small reductions in Rabi-flopping efficiency. Their atom cloud transverse width was comparable to the beam waist, and since the diffracted population is very sensitive (see e.g. [57]) to lattice depth, the finite cloud size would have been a more dramatic limiting factor in the Rabi-flopping efficiency than the on-axis velocity spread.

One main takeaway from Fig. 5.5 is that there is a singularity in the phase of the final wavefunction at zero velocity and the π -pulse intensity. This *can* occur because the




Figure 5.5: (left) Wavefunction population in the target state and (right) wavefunction phase difference (in rad) between the target and initial states. Note that the unwrapped phase difference has a singularity where the population in the target state reaches a maximum, due to the population in the initial state going to zero. Each pixel is one precomputed data point in the tables.

population in the initial state goes to zero at this point, however I'm not sure why it *does* occur. This is potentially the reason why old alpha didn't use precompute tables in their Monte Carlo simulations, because interpolating around the singularity and branch cut could lead to incorrect results. The reason I have moved forward using these look-up tables anyways is because this singularity occurs around the π -pulse intensity, whereas we use $\pi/2$ -pulses in our experiment and during simulations. For $n = 5$, the intensity would have to be 20% too large in order to have any issues in the precompute look-up, and for lower Bragg orders the intensity would have to be even larger still. Even for the exaggerated beam distortions I'm looking at in this thesis, $\pm 10\%$ distortions are pretty much the worst case scenario I consider, so even with random initial conditions atoms would never be taking values from these problem areas. Even so, anyone using these precompute tables should know about this issue in case you try to use the code for new situations that break this assumption. For example, if you want to simulate systematic effects in a dual Mach-Zehnder gravity gradient measurement, you will need to simulate pi-pulses, which will definitely result in interpolation errors.

5.4.7 Beam class

The Beam class is responsible for everything related to the electric field $E(x, y, z)$ structure of the laser beam, including the electric field and local laser phase, lattice depth, light shift, and derivatives of these quantities. As the code is written now, everything is based off of the local electric field of the laser beam. This assumes knowledge of the phase of the electric field everywhere, however in lab we only have knowledge of the intensity of the beam where we take CCD images. Some remarks about this:

First, and most importantly, you can easily change the structure of this class in the future since it is all disentangled from the rest of the code. For example, you could change it to load from CCD images of the experimental beam. This would look quite different than how the code currently operates, but mostly you would only need to change functions in the `libs/models/beam.py` and `libs/models/analysis.py` files (and then update the function calls in the rest of the code as needed).

Second, writing the code in terms of the electric field was useful for the experiment where we intentionally modulate the laser beam to compare Monte Carlo to experiment. We strongly modulated the beam in a controlled way such that we effectively knew the phase front of the electric field, even without measuring it directly.

Third, it's a more direct physical way to derive other quantities, making the coding much more physical and natural. It also allows one to study effects much more carefully. For example, the Gouy shift causes the k -vector of the laser to change. This causes a local phase shift on the electric field of the laser, since the wavelength is locally different, which leads to a different laser phase imparted on the atom. The deviation in the k -vector also changes the recoil velocity of the atom, leading to a different accumulated kinetic energy phase after propagation. Laser phase must be included in the laser beam model in order to accurately calculate these effects. A simple model of only the intensity of the laser would prevent you from implementing some of these physical effects, or would effectively force them to be zero thereby ignoring potentially important systematic effects.

Fourth, one of my goals/ our goals with this code is to bound systematics from the laser

phase front well enough that we don't need to directly measure it in lab, but instead can just bound its effects by looking only at the intensity propagation of the beam with CCD images. I hope a future student is able to use this Monte Carlo package to demonstrate that laser phase doesn't need to be measured experimentally, since this would make the beam characterization much more complicated.

As the code currently operates, calling the local electric field of the laser beam is the most time consuming part of the code, especially for more complicated beams. This is discussed in more detail in Section 5.6.

5.5 Typical Monte Carlo Workflows and Intended Use

First, you should begin with already precomputed precompute tables for the simulations you intend to run. Next, assuming you have your precompute tables, you typically work on some file in the 'scripts' folder. The most important scripts used for studying the system or doing science required revolve around Sweep or Sweep2D objects, however there are plenty of other scripts as well for things like testing Bragg diffraction or Bloch oscillations numerical integration, testing that Gouy shifts in simulation match the analytical estimate, testing convergence of the precompute tables, etc. I highly recommend you learn familiarize yourself with these test scripts when first exploring the code - it will be by far the fastest way to start seeing how things work.

One of the most useful scripts for general debugging or playing with the code is `scripts\general_tests\test_ellipse.py` - it generates and plots one ellipse. You can use this script to get fast feedback on how any changes you're making to the code affects the ellipses. You can also change any parameter and see its' effects on the ellipse. For initially building an intuition of how the parameters affect the entire code, this script will be your best friend. Also, you can use `cProfile.run('ell.run_SCI_ellipse()')` instead of `ell.run_SCI_ellipse()` to profile the speed of the code. It will print out an extensive summary of which function calls are taking how much time, so you can see where the bottle-necks are. If you want an accurate estimate of timing for larger runs, make sure the simulation itself is at least tens of seconds since the initial compiling of the code and loading imported modules takes O(5-10s).

If you're running a major simulation that requires lots of statical integration, or if you're computing new precompute tables, you'll want to run these files on the remote simulation computer. The sim computer can multi-thread 128 processes in parallel, which is extremely useful for off-loading huge simulations (the sim computer also has an almost 2x faster clock than my laptop). I use CyberDuck to transfer files between my laptop and the sim computer, WireGuard to VPN into the lab internet network, and then I ssh into the simulation computer once on the same network. Slurm is used for submitting and managing all simulations run on the sim computer. From the command line, you navigate to the correct folder then type e.g. `'sbatch submit_slurm_example.sh'` which runs a shell script that actually calls the python code, after allocating cores and memory. Huge thanks to Jack Roth for setting this all up and teaching me how to use it. If you have any questions about how it works, you'll have to ask Jack as I'm fairly clueless here.

The code is set-up to be able to parallelize simulations in several different ways. At the highest level, if `prms['sweep_parallel'] == True`, then each ellipse in your Sweep or Sweep2D

will be sent to different cores. For example, if you are doing a 1D sweep with 20 sweep points (which simulates 20 ellipses, each with different parameters), it will assign each of the 20 ellipses to 20 different cores. Note that we have a 128 core simulation computer and we would only be using 20 cores, so this is not an optimal use of computing resources. Note that if `prms['sweep_parallel'] == True`, then `prms['ellipse_parallel']` is set to `False` automatically to avoid clashes in the code with over-parallelization.

To circumvent the issue of not optimally using all cores, you can also parallelize over each point on each ellipse instead of over the ellipses themselves. To do this, set `prms['sweep_parallel'] == False` and `prms['ellipse_parallel'] == True`. This way, you can use 128 points on your ellipses (or integer multiples of 128) in order to optimally use all cores on the sim computer. The code will go through each ellipse linearly, and parallelize over the ellipse points within each ellipse. This is the recommended way to fully utilize the sim computer cores. This would not be effective if you only had a few atoms per ellipse point, since it's $O(\text{seconds})$ expensive to set up the parallelization each time. But when we have typical runs of thousands of atoms per ellipse point, where each ellipse point takes minutes or hours to run, this set-up time doesn't matter.

The last way that the code is parallelized is when precomputing the precompute tables. The code linearly goes through the combinations of `n_bragg` and `N_bloch`, then within this will parallelize over the table of `Omegas` and `velocities`. Note this is parallelization over a 2D table which uses `starmap()` instead of `pmap()`, but other than that essentially works the same. Typical precompute tables have $O(200)$ points, so we're doing OK on making optimal use of the sim computer's 128 cores, but not entirely optimal. Maybe in the future it would be worth parallelizing the entire precompute run to make more optimal use of the 128 cores, but for now it's functional and less than a factor of 2 from optimal.

When running on your personal computer, make sure to set `prms['n_cores']` to 8 or whatever your computer has (or less than 8 if you don't want to clog up your whole computer). When on the sim computer, set to 128 (or 64 or 32 if you want to run 2 or 4 sims in parallel, etc.). Make sure to also assign the number of cores and RAM per core on the `submit_xxx.sh` slurm file, otherwise you might not get access to the number of cores that you put in `prms['n_cores']`.

Some comments about file paths - the default file paths are different on my laptop versus on the sim computer. On my laptop, it typically defaults to the overall `'alpha_monte_carlo\'` folder. For example, when it saves HDF5 files or images to the results folder, they go to `'alpha_monte_carlo\results\'` ... When on the sim computer, if you run a script in `'alpha_monte_carlo\scripts\example\'` ..., it will save results to `'alpha_monte_carlo\scripts\example\results\'` ... This has proven to be quite convenient, since it keeps the sim computer runs separate from the local runs, and since then you only need to synchronize one folder between the sim computer and your local computer. It also stores new precompute tables in `'alpha_monte_carlo\scripts\precompute\temp\'` ... instead of overwriting the existing precompute table. This could probably be changed relatively easily in the future if you want to, but as long as you know how it behaves it should be fine. The only exception would be if e.g. different people's computers do different things, then it's probably worth standardizing it.

5.6 Current Monte Carlo Speed

The state of the code as I'm graduating is pretty close to as fast as you could make it in python using the Atom class structure I've laid out. An example profiling of the 'test_ellipse.py' script is broken down in Table 5.1 where the most time consuming functions are tabulated. For this simulation, all of the 'physics' is turned on except for light shift phase and light shift transverse momentum kicks, because these are cancelled experimentally when we use light shift compensation. The most time-consuming call is 'atom.get_E_fields' at 25%, which is the function where a grid of local E-fields is stored to the Atom object so that the light shift, lattice depth, laser phase, and dk_vec can all later be calculated. This simulation was run with a 'Gaussian' beam type, which is the simplest beam type. More complicated beam profiles very quickly become the bottleneck in the code as it is currently written and are the primary area that needs to be sped-up moving forward.

The second-most time-consuming functions are the 'precompute_get_bragg_final_wf' and 'precompute_get_bloch_final_wf', which is where the wavefunction values are determined from the (pre-loaded) interpolation functions over the precompute data. Currently, the code loops over wavefunctions and loads the final wavefunction values from the correct precompute table, so this function is *NOT* vectorized. One possible speed-up in the future would be to vectorize this code so that many wavefunction values can be loaded at the same time, however this would likely require substantial changes to the underlying datastructures of how the code is built.

Next, the 'analysis.get_dk_vec' function is quite time-consuming because it requires taking the local laplacian and therefore involves a lot of numpy operations. The 'analysis.get_light_shift_derivative' function is not shown here but required even more time/numpy operations, since it is taking a derivative of the already complicated 'analysis.get_light_shift' function.

I put some example numpy operations in this table for reference - this simulation code is already extremely fast if 'np.angle' is 5% of the total simulation time. Note that the numpy times are already included in some of the above function calls, so this table is not meant to sum to 100%.

On the simulation computer with 128 threads in parallel, I can run about $8 * 10^8$ atoms per day with these parameters and a more realistically complicated 'dusty' beam with 10 pieces of dust. The statistics of how an ellipses averages down phase noise depend on the beam structure and the contrast of the ellipse. For a dusty beam (10 pieces of dust around $75\mu m$ diameter within the region $(-w_0, w_0)$), 24 hours of running averages down to about 0.7 mRad error bar on a single ellipse. In this example, the total phase of the SCI interferometer was around $7 * 10^7$ rad, so the sensitivity to alpha is roughly $0.5 * 10^{-11}$ in 24 hours of Monte Carlo integration. The final alpha measurement will probably aim for an error bar around this level, and experimentally we will aim to have this level of integration in 24 hours, so this very roughly means the current Monte Carlo can keep up with real-time analysis of the experiment. Obviously we would want the code as fast as possible, but this isn't a terrible place to be for final data analysis. However, the speed gets more frustrating when you're trying to use the code for studying a new process or gaining intuition for which parameters are important. For example, if you want to sweep e.g. launch velocity of the atoms to see the sensitivity to alpha at the 10^{-10} level, each point in your sweep can take 24 hours, depending

Function	Total time (s)	Total time (%)
Ellipse.ellipse_pt	50.7	100.0%
atom.load_all_E_fields	12.6	24.9%
atom.split_all_atom_data	1.5	3.0%
atom.get_SCI_point	2.6	5.1%
analysis.get_dk_vec	7.1	14.0%
analysis.get_laser_phase	1.1	2.2%
analysis.get_lattice_depth	0.7	1.4%
precompute.get_bragg_final_wf	7.8	15.4%
precompute.get_bloch_final_wf	1.2	2.4%
timeev.add_kinetic_energy_phase_atom	4.2	8.3%
timeev.update_position_and_velocity	4.4	8.7%
np.angle	2.7	5.3%
np.linalg.norm	1.6	3.2%
np.ones	1.5	3.0%
np.outer	1.6	3.2%

Table 5.1: Ellipse timing for 16 ellipse points, 1000 atoms per point, on a single core of my laptop. Data is shown for a 'Gaussian' beam, which is the fastest to run in the simulation. More complicated beams make the atom.load_E_fields() call longer such that it quickly dominates the total simulation time. See text for more details.

on how noisy the beam is. Waiting a week for each plot makes iteration extremely difficult.

5.6.1 Potential Monte Carlo Speed-ups

With the exception of some needed improvements in how slow it is to call the local electric field, you're going to struggle to get a factor of 2 improvement in the speed of the code. Instead of chasing after factor of 2 improvements in the future, you would be much better off learning how to run code on a computing cluster where you could get 1000 cores to use, or purchasing a larger simulation computer to sit in the corner of lab. Our current computer can run 128 processes in parallel, and there are rumors that next-next-generation AMD processors might be able to run 512 processes in parallel.

As for the code itself, since we're not too far from limited by the speed of calling the actual interpolation function over the precompute tables, there's not too much more we can do with the current structure of the code. The biggest win would be speeding up the 'get_E_field' function call. For example, in the 'dust_spot_E_field' beam model, the code is vectorized over (x, y, z) positions in the 'atom.load_all_E_fields' function call (up to 200 positions max), but then the code does a for loop to sum over the contributions of each piece of dust. This could be improved by vectorizing the entire function call, including the sum over each piece of dust (do instead with `np.sum` along the desired axis). However, I don't know how much this will help in practice because in order to vectorize over many variables, each variable would have to become a multi-dimensional object, and allocating/ updating this much memory would be very time-consuming. Maybe a better way to do this is to create a three-dimensional look-up table for the value of the electric field so that the code only needs to call an interpolation function. There is a good chance that calling the interpolation function is quite a bit faster than summing terms in the analytic function. On the other hand, making an interpolation over such a large region of space on a complicated beam propagation structure will likely run into memory issues, so I'm not sure this is a viable solution moving forward.

If you really need an order of magnitude improvement, I don't think you can do this with the code written as-is. Your best bet would be to redo the core of the code to use e.g. 'AtomCloud' objects instead of 'Atom' objects, so that instead of vectorizing code within an Atom (over 16 wavefunctions/positions/velocities), you could vectorize the code over the entire atom cloud (thousands of atoms or more). This way, instead of for-loops over each atom in the ellipse, you just run through the time evolution code once, evolve all atoms at the same time, and then process the results. Since most of the code is numpy operation limited, this should definitely speed up the code, but I don't know by how much. The downside of this change would be that the code is much less general and intuitive. For someone who knows the code well, in a few weeks you could probably implement this enough to get a sense of how much speed-up you would see.

A terrible option that I don't recommend would be to re-write the entire package using numba 'jit(nopython=True, cache=True)' decorators on every function. Numba has terrible support for classes, so you probably wouldn't be able to keep using class structures, but instead would need to transition to just function calls. The convenience of class structure organization would be hard to give up. You would also need to give up the 'prms' dictionary, although possibly you could replace it with some sort of fixed list or data structure within numba. These changes would require re-writing pretty much all of the TimeEv, Atom, and

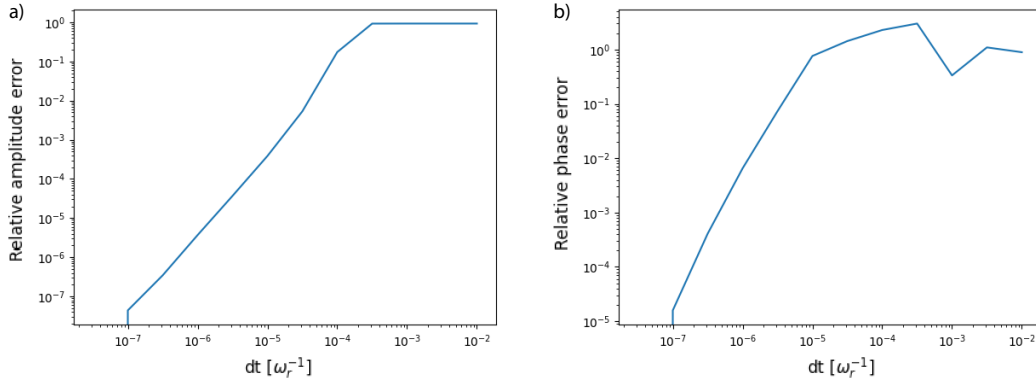


Figure 5.6: a) the amplitude and b) the phase error during integration of the Schrodinger equation, as a function of the integration time step 'dt'.

Beam classes, and much of the Ellipse class. It's hard to say whether this would be useful or not. It would give you a speedup, but at the cost of the code being much less readable, and much more difficult to use/code new things.

All that said, you're probably better off making the 'atom.load_all_E_fields' call as fast as possible, getting access to a large computing cluster, then focusing on getting the experiment working...

5.7 Monte Carlo convergence

There's a few areas where we have to think about how well the Monte Carlo is converging: integration of the Schrödinger equation, interpolation over our finite precompute table, and statistical convergence when running simulations.

5.7.1 Integration convergence

First, integration of the Schrodinger equation requires a finite time step in our discrete Crank-Nicolson integration scheme. The errors in a Crank-Nicolson scheme are known to scale as dt^2 [71], so typically a 10x smaller time step should give you a 100x smaller error.

Figure 5.6 shows an example convergence plot I used for determining integration convergence. The integration time step dt is swept, and the error of each simulated wavefunction amplitude and phase (in the target state) is plotted relative to the smallest time step in the sweep. We can see amplitude and phase both converge as dt^2 once the errors start falling below $O(1)$.

Since we are ultimately interested in 1 mRad errors in the Monte Carlo, I use a criteria for integration convergence of 0.1 mrad relative errors. The most important thing here is that the errors between the precompute tables are less than 1 mrad relative to one another, and this convergence criteria should safely avoid any 1 mrad relative errors.

When setting up the precompute tables, I would run a convergence plot like this for each Bragg and Bloch order, since higher order processes need smaller timesteps for convergence.

One major issue I've been running into is that very high Bloch orders ($N=100$ or more) require exceedingly small time steps for convergence, particularly for multi-frequency pulses. Multi-frequency pulses have the $\cos(\omega_m t)$ term on the coupling terms in the Hamiltonian, and therefore convergence doesn't even start until your time step is much less than $1/\omega_m \sim 10^{-3}$ for $N=100$. However, these simulations are also very long in time ($t_{tot} \sim 100\omega_r^{-1}$). Therefore we typically don't start convergence until we have well over 10^6 time steps. To get good convergence, I'm finding we need $O(10^{10})$ time steps, for which a single integration can take days.

A general takeaway is that the dt needed for convergence of multi-frequency Bloch pulses scales roughly like $1/N^3$, so going from $N = 5-$ to $N = 100$ is extremely expensive. In any case, I've still found convergence for $N=50$ on the order of a few days. It's possible we can overcome some of these convergence issues with a higher-order integration method beyond the Crank-Nicolson method that converges as e.g. dt^4 . We would still need time steps smaller than the modulation frequency before we see convergence, but if we can save a decade or two in time steps it could save lots of time, even if the integration method itself is more expensive per step.

5.7.2 Precompute table convergence

The next area we need to study convergence is in the precompute tables. To generate the precompute tables, wavefunctions are simulated over a 2D array of initial (on-axis) velocity and lattice depth (see Fig. 5.5). The grid spacing on this 2D space must be chosen such that interpolation between grid sites still gives an accurate wavefunction to better than 0.1 mrad anywhere on the grid, the same convergence criteria as in the integration. When setting up the initial grid spacing for the table in this project, I was surprised that a rather coarse grid spacing can still give 0.1 mrad errors anywhere on the grid. The grid spacing will need to be tested more rigorously in the future though, since there are many tables that should all individually be checked for errors.

5.7.3 Statistical convergence

For beam profiles with lots of intensity noise, the statistical convergence of the Monte Carlo can be very poor. Since the Monte Carlo doesn't include atom number shot noise, I believe the poor convergence is coming from lots of phase noise atom-to-atom. Many of the plots show that the beam distortions cause large losses in contrast, and if this contrast loss was from phase noise atom-to-atom it would indicate $O(2\pi)$ phase noise or more.

More work needs to be done here to set up a rigorous system for confirming statistical convergence. Initial results to this end are not showing convergence as I would expect, which might be an indication that outliers in the data are preventing successful convergence for very large atom numbers. Possibly this is related to the finite boundaries of the precompute tables such that extrapolation beyond the boundaries is leading to outliers in the data.

For generating the plots in the section below, I was able to qualitatively see that the plots converge after running the plots with different atom numbers. In the very near future, I plan to adopt the code to store the errors on the ellipse fit, which should give a good measure of statistical convergence on the ellipse data. I could also implement something like an Allan

deviation by simulating millions of atoms, then grouping them in different ways and plotting the relative error.

5.8 Monte Carlo Results

We have very recent experimental data where we have not seen shifts that were predicted by the Monte Carlo code, so as a disclaimer, the results in this section are likely quantitatively incorrect. We're just beginning to match experiment to Monte Carlo, so I expect there will be quite a bit of debugging before the two agree. Also, we need a better method for confirming statistic convergence of the plots. That said, it's still very instructive to qualitatively look at what the Monte Carlo model is predicting in order to get some intuition for which physical effects matter and which don't, and to get an order of magnitude estimate of the size of these effects.

Unless otherwise stated, the Monte Carlo simulations below use the same initial conditions that are used experimentally. Some important parameters are listed here. $n = 5$ Bragg diffraction is used, and $N = 50$ Bloch oscillations accelerate the atoms at an acceleration $a = 10g$, where g is the acceleration due to gravity. Atoms are moving upwards at 3.3 m/s at the time of the first Bragg pulse. The thermal velocity spread is $0.05v_r$ vertically, and $1.5v_r$ transversely, where $v_r = 3.5$ mm/s is the recoil velocity of Cs. The atom cloud is assumed spherical with a 1σ width of 1.5 mm. The laser beam has a $1/e^2$ radius of $w_0 = 6.2$ mm. The focal plane of the beam is 2.3 m before the atom location during for the first pulse, and the retroreflection mirror is 4.7 m after the focal plane.

One important result from the Monte Carlo mode is that transverse components of the local k-vector of the laser are an important mode of loss of contrast from spatial intensity noise. As far as the authors are aware, these effects also have not been included in Monte Carlo simulations of atom interferometers. Fig. 5.7 shows the contrast of the interferometer as a function of the spatial frequency of the square wave modulation from Eq. 5.11. The three curves show the effect of different modes of loss of contrast, showing that transverse components of the k-vector are an important factor in capturing the contrast loss seen in atom interferometers.

When all modes of loss of contrast are included, we see severe loss of contrast for most spatial frequencies. When light shift momentum kicks are omitted, there is a large revival in contrast, especially for higher-frequency modulations. This is the case experimentally when light shift compensation is used. Notably, the effects of imperfect k-vector momentum transfer still do cause substantial loss of contrast. The final curve omits any transverse kicks from the simulation. The remaining effects on contrast seen in the plot are from non-ideal Rabi frequencies during Bragg diffraction. These results are consistent with the experimental observation that light shift compensation can revive some but not all loss of contrast [66, 46].

Next, we show some important scalings in Fig. 5.8, in particular the differential phase shift dependence on the distance from the source of intensity noise and the atoms, as well as the lattice depth during Bloch oscillations. Fig. 5.8a) shows that interferometers become much less sensitive to higher frequency spatial noise the further they are from the source of intensity modulation on the beam. This is a result of the exponential decay with propagation distance from Eq. 5.10. Additionally, the peak differential phase shift is roughly proportional




Figure 5.7: Contrast loss as a function of the spatial frequency of the modulation. When transverse momentum kicks from light shifts gradients and transverse k-vector components are included, there is broad loss of contrast. When light shift gradients are omitted (or cancelled experimentally), there is still broad loss of contrast from transverse k-vector kicks, especially at lower spatial frequencies. When all transverse kicks are omitted, we see the effect from lattice depth variations on Rabi frequencies during beamsplitters. The apparent 'noise' on this plot is actually from structure in the propagation of the beam.




Figure 5.8: a) Differential phase shift in units of ppb in α as a function of the spatial modulation frequency. Different curves represent different propagation distances z from the dust to the first pulse of the interferometer. b) Same axes, but the different curves represent different lattice depths during Bloch oscillations, in units of the recoil energy $E_r = \hbar^2 k^2 / 2m$.




Figure 5.9: Effects of dust size on a) the differential phase and b) the contrast of an SCI interferometer. The legend in b) applies to both figures. See text for more discussion and for parameters used.

to $\propto 1/z^2$, where z is the distance between the source of the intensity noise and the first pulse of the atom interferometer.

Fig. 5.8b) shows that lower lattice depths during Bloch oscillations lead to larger phase shifts in the interferometer. This is a result of atoms preferentially tunneling from lower intensity regions of the noisy beam, leaving a correlation between the surviving atoms and the location in the beam. This correlation is stronger when Bloch oscillations are further from saturation, ie. when the overall lattice depth is lower. For reference, a Bloch lattice depth of $\Omega_b = 9.4E_r$ corresponds to roughly 50% survival rate. For a detailed analysis of this effect, see [6]. Note that this figure was generated with a sine wave modulation *without* the term that oscillates amplitude modulation to phase modulation and back. The purpose of the figure was to gain a qualitative understanding of the different scalings with a simpler beam model.

In addition to the phase shift scaling with the Bloch lattice depth, there is an additional structure that is unchanged for all Bloch powers. These effects are from Bragg diffraction - they are a result of correlations between the Bragg diffraction efficiency and the δk momentum kick received. These effects warrant a more careful analysis before a next-generation measurement of the fine-structure constant is published. The old experiment used the scaling with Bloch power to experimentally bound the potential effects from these δk variations, but this method would not have captured the remaining correlations from Bragg diffraction.

The model for dust on optics presented in section 5.3.1 is used to study the differential phase shifts and contrast loss caused by a random scattering of dust on an optic before the vacuum chamber. In Fig. 5.9, a density of 20 dust particles are randomly distributed over the region $(-w_0, w_0)$ in both the x- and y-dimensions. For our beam waist of $6.2mm$, this corresponds to roughly one piece of dust per $(2.8mm)^2$. An SCI interferometer is simulated with a Bragg order $n = 5$ and a Bloch oscillations order $N = 50$, using the same initial conditions as are used experimentally, and light shift compensation is assumed. With the dust locations fixed, the size of the dust particles is varied and the interferometer differential phase and contrast are plotted in Fig 5.9a) and 5.9b) respectively. Each line on the plots corresponds to a different initial distance z between the dusty optic and the first interferometer pulse.

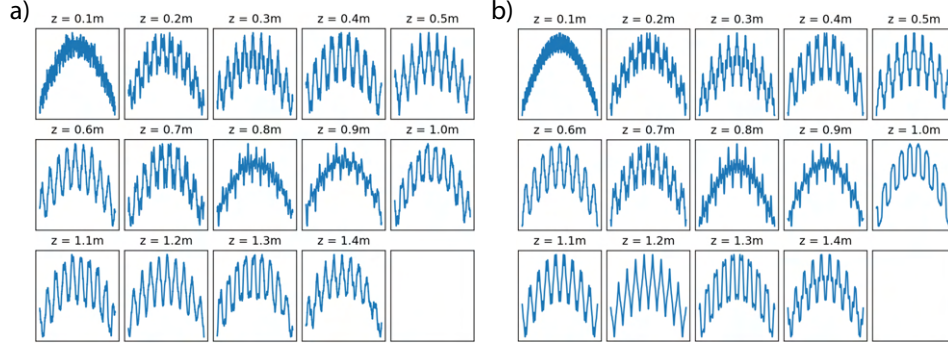


Figure 5.10: Propagation of a square wave intensity modulated laser beam showing strong Talbot patterning. a) Experimental data with images taken every 10cm - the images are vertically integrated to average over noise from other optical distortions e.g. from dust on optics. b) The corresponding theoretical prediction from Eqs. 5.11 and 5.9. The only fit parameters from the model are the absorption and phase shift as the beam goes through the titanium thin film.

This result demonstrates that a moderately dusty optic can easily cause ppb-level phase shifts in a SCI interferometer, and large pieces of dust can cause 10ppb effects or larger. Allowing a laser beam to propagate for several meters in vacuum can suppress the effects somewhat, however the effects clearly remain. The apparent 'noise' in Fig. 5.9a) is a result of the highly structured propagation of the beam's electric field. When the dust size is swept, the entire propagation of the beam is changed because each dust particle's effects now have a different Rayleigh range. The five discrete times when the laser interacts with the atoms during the interferometer then sample a seemingly random set of local effects, causing seemingly random resulting phase shifts.

5.9 Initial experimental results

The first step in verifying our model is verifying our calculations of how a complicated laser beam propagates. We've been collaborating with Sam Peana, a researcher at Purdue, who is fabricating masks of thin titanium patterns deposited onto an anti-reflective coated glass slide. Deposition thicknesses around 2nm give an absorption around 10%, which gives us a somewhat strong modulation on the beam but not so strong as to break our assumption that the modulation is small relative to the overall Gaussian amplitude. The first ≈ 0.5 nm of the titanium oxidizes from exposure to air, which tends to phase shift the laser beam but not absorb it.

This model was verified experimentally, see Fig. 5.10. A pattern of titanium was deposited onto an AR-coated glass wafer in a periodic pattern creating a square wave amplitude modulation on the beam. The titanium reflects some of the electric field and also phase shifts the transmitted field. The mask shown in Fig. 5.10 very closely matches the model in Eq. 5.12 with $A_q = 0.034$, $\phi_q = 0.72$, the only two fitted parameters. The spatial period of the modulation $q = 2\pi/(600\mu m)$ was fixed by the mask fabrication.

After the above mask was used to confirm the propagation model of the beam, Sam Peana

fabricated a series of masks with various noise spatial wavelengths $q = (300 \mu\text{m}, 400 \mu\text{m}, 600 \mu\text{m}, \text{ and } 800 \mu\text{m})$. This second set of masks had approximately twice the modulation depth A_q and a similar phase modulation ϕ_q , which is consistent with the fact that the phase modulation came from the outer oxide layer, independent of the overall titanium thickness. Based on initial Monte Carlo estimates, this mask series seemed to cover a range of spatial frequencies where we could expect to see experimental phase shifts and loss of contrast with this strong modulation depth.

Additionally, Sam fabricated masks with psuedo-dust. A random series of dust positions with 100 particles/in² was fabricated on a 2" x 2" glass slide, where circles of titanium fully absorbed incident light. This density of dust particles is roughly one dust particle per 6 mm², or equivalently around 6 dust particles per beam waist squared when using a 6 mm beam waist. Some areas of the mask have higher dust density than others, so depending on the desired density one can move the beam to a different region of the mask to change the effect. The dust particle sizes were uniform on each mask, and four different masks were fabricated with dust diameters 40 μm , 80 μm , 120 μm , 160 μm . At 160 μm diameter dust, the Monte Carlo model predicted dramatic loss in contrast of the interferometer, and possibly an associated phase shift around the 1 ppb level.

Experimentally, we have used these masks to try to measure loss of contrast and differential phase shifts in the interferometer. So far, the experimental results indicate that the Monte Carlo is overestimating effects - where the Monte Carlo predicts we should see contrast loss and phase shifts, we don't yet see any effects.

The first main data set was a several day campaign with and without a mask in front of the main laser beam. During data taking, we were blind to the comparison with/without the mask so that we could verify the data was self-consistent over many days before checking to see if we had seen a phase shift. The data was consistent with no phase shift at the roughly 1 ppb level, while the Monte Carlo predicted a larger shift. The mask used was the same mask that was fit to the model in Fig. 5.10.

The second data set was a much quicker test - we put many masks in front of the beam so that the beam was very highly modulated, then we looked to see if there was any measurable contrast loss in the interferometer. This data set used some of the new set of masks from Sam Peana that had double the modulation depth, in addition to the 160 μm diameter dust mask. The experimental result was that there was no measurable contrast loss. For both sets of data, parameters used were $n = 3$ Bragg diffraction, $N = 100$ Bloch oscillations, and $T = 100$ ms interrogation time.

The fact that the second test saw no measurable contrast loss indicates that the contrast loss model is incorrect. The predicted phase shifts are closely related to the predicted loss of contrast, since the dominant phase shifts come from a correlation between k -vector kicks an atom receives and the probability that the atom contributes to the overall contrast. However, there are far too many parameters in the model to draw conclusions from a only a couple data points.

I think the most likely scenario is that there are factor of 2 or π errors in the code that have resulting in overestimating the effects we should see experimentally. The physics included in the Monte Carlo is based on first-principles so should be very accurate. One way in which the Monte Carlo model differs from how experimental data is taken is that the Monte Carlo model only runs simulations at a single interrogation time T and looks for

relative shifts as some parameter is swept. Experimentally, we take data at several different T 's in order to extract the diffraction phase, since the diffraction phase experimentally is seen to change over time. This discrepancy would likely change the Monte Carlo estimate somewhat, but I don't think it would eliminate the effect we're looking for.

Another possibility is that there is some physics omitted from the model that ends up being relevant experimentally. For example, the Monte Carlo doesn't include effects from single photon scattering, even though we see strong scattering effects experimentally. It's possible that including scattering effects in the model would somewhat change the predictions, but again I don't think this effect dramatically affect the predicted experimental shifts.

In total, much more work still needs to be done to match the Monte Carlo results to experiment. The model is very complicated, so even with careful coding and lots of double-checking, there are surely still errors in the code. Once we are reliably comparing experiment to the model, we can see if there are any other physical effects that need to be added to the model in order to reproduce experimental data. More to come!

5.10 Future Monte Carlo Work

While I tried to include as many effects as possible in the code, there are still some physical effects that are left out. There are also ways you could try to speed the code up in the future, if you end up needing better statistics.

5.10.1 Monte Carlo Assumptions

First, parasitic interferometers are not included in the simulations. I saved the entire wavefunction array in the precompute tables, so the wavefunction information is still contained in the precompute tables, just that the code does not make use of it. In order to include this in the future, you'll have to a) store additional data on the atom object for amplitudes following parasitic paths, and b) modify the TimeEv operators to also capture the evolution of the parasitic states. This probably won't be too hard to do, just make a good plan for the modifications before starting. It'll be a good learning exercise for how the meat of the code works.

Next, single-photon scattering is not included in the simulation. It wouldn't be too hard to track a probability of scattering at each pulse based on the local laser intensity, but the hard part is in determining whether this atom is then seen by our detection photodiode. When an atom single-photon scatters, it absorbs a photon from either the up-going or down-going beam and receives a single recoil velocity kick along the axis of the beam. Then it emits a photon in a random direction when it decays from the excited state, giving another recoil velocity kick in a random direction. The time-of-flight after the pulses is typically around 0.5 - 1 second, so for a recoil velocity of 3.5mm/s the atom moves around 2-3mm away from center of the initial atom cloud. The experiment uses a $\tilde{3}\text{mm} \times 3\text{mm}$ photodiode and a one-to-one light collect scheme¹¹, so to first order we're only capturing a $\tilde{3}\text{mm} \times 3\text{mm}$ region of the center of the cloud. To make matters more complicated, the probability of detection isn't a discrete function, it's a complicated transfer function where you can still detect an

¹¹actually the magnification is 0.8

atom far from the center of the cloud, just with a decreased probability. In principle, you could try coding all of this up, but in practice I imagine it's going to be hard to get the model agreeing with the experiment. Not impossible, but hard.

Note that we've never looked into systematic effects from single-photon scattering in our experiment, we've just assumed that they're zero. If the population in the different peaks is roughly 50/50 during the interferometer, the scattered atoms will create a uniform background in each time-of-flight experimental peak. Adding a constant background like this will reduce contrast but not rotate the ellipse. However, our Bragg pulses are far from perfect due to finite size of the atom cloud with respect to the finite laser beam, so the populations are *not* 50/50 during the interferometer. This will cause an asymmetric scattering background across the various time-of-flight peaks. It's possible that this asymmetric background gives a systematic phase shift when performing ellipse fitting - the effect should be further investigated.

I'm also going to document here a number of small assumptions that I think are negligible but still good to document. I'll first make a list then follow up with a more detailed explanation.

1. Light shift gradients on-axis are ignored
2. The problem is only treated quantum-mechanically in 1D
3. The k-vector is assumed constant across the wavefunction when solving the Schrodinger equation
4. All atoms are assumed to be in the same F internal state the whole time
5. Similarly the atoms are assumed to be in the $m_F = 0$ magnetic sublevel
6. The gravity gradient systematic is omitted
7. The classical propagation during Bragg/Bloch pulses is approximate, not exact
8. I assume uncorrelated initial conditions
9. Some small assumptions go into the Precompute tables.

1) In taking the light shift gradients for the light shift forces, I don't account for the gradient along the axis of the beam because it will be many orders of magnitude smaller than the transverse effects. I don't see numerically how this would matter, but maybe I'm over-looking something.

2) I only treat the problem quantum-mechanically in 1D, along the beam propagation axis. This is because the optical potential changes on a shorter length-scale than the thermal deBroglie wavelength along that axis: $\lambda_{light} \ll \lambda_{dB_z} \sim 5\mu m$. Any transverse changes in the potential are on a much larger length scale - any optical distortion with wavelength less than $50\mu m$ or so is washed out by the time it reaches the atoms, and the transverse thermal deBroglie wavelength is very small: $\lambda_{dB_{transverse}} \sim 100nm$. Therefore any transverse potential gradients are essentially linear over the extent of the atom's wavefunction and can be treated classically.

3) Similarly, I assume the laser k-vector is constant when solving the Schrodinger equation, since we discretize momentum in integer multiples of k . In reality, it has spatial dependence (e.g. local Gouy shift), but on length scales much longer than the thermal deBroglie wavelength of the atom. You'd have to integrate the wavefunction in real-space to simulate these effects, and I don't have a good intuition for what to expect, but the effect should be very small. Maybe it would effectively look like 'heating' of the wavefunction, since you're diffracting with a spread in k-vectors.

4) I assume that the initial atom cloud is entirely in the same F state/ sees the same lattice beam detuning. If you don't have good state prep, these other atoms could in principle still go through the interferometer. Since it's an incoherent mixture and the probability of forming a coherent interference from the wrong state will likely be very small, this should have a small effect.

5) Similarly, I assume that atoms remain in the same $m_F = 0$ state. Imperfect state prep or B-field gradients could lead to atoms in other m_F states. These atoms could still form interferometers, but the magnetic field phase would be very large and likely have a large spread ($> 2\pi$) across the atom cloud, meaning the contribution would look like an incoherent loss of contrast. If somehow there were enough atoms in a different m_F state that had a well-defined phase shift, it's possible these atoms could cause a systematic phase shift when fitting ellipses.

6) I don't include gravity gradient systematic, since you can calculate this very well analytically (or even numerically if you map out gravity gradient as a function of height). This effect should be independent of random initial conditions of the atom cloud since it varies on a much larger length scale, so you don't need Monte Carlo to characterize it.

7) The way I treat classical propagation of velocity and position of the atom within a Bragg or Bloch pulse is approximate, and not exact. That is, the momentum kicks are modeled discretely during the center of the pulse. For Bragg pulses, I classically propagate to the center of the pulse (in time), then do the Schrodinger equation evolution, transverse momentum kicks, laser phase, etc. based on this location, then finish classically propagating the second half of the pulse. For Bloch, I propagate to the center of the parabolic trajectory (in time), otherwise everything else is treated the same as Bragg. Bragg pulses are very short in time (tens of μs) compared to typical T's (tens or 100 ms), so for Bragg this shouldn't matter. Bloch pulses can be several ms though, so there might be large physical effects. Even if we're off on the total Mrad phase, we'll still capture the beam-related systematics to O(1%) or better, which is probably as good as we'll ever be able to compare to experiment anyways. Also, this way of modeling assumes the lattice depth is constant for the entirety of the e.g. Bloch oscillation pulse, whereas in reality the atom is traveling during the pulse and sampling different parts of the beam. The atoms during Bloch oscillations might travel O(5mm) during a Bloch pulse - the beam propagation is so minimal over 5mm that I assume that using the center location for the local lattice depth is sufficient to first-order. It's still worth noting that this modeling assumption is baked into the model.

8) I don't propagate the atom cloud before starting the interferometer. Experimentally, the cloud propagates for hundreds of ms, which allows the position and velocity of the atoms to become more correlated. Moreover, our Raman sideband cooling is likely most effective

in the center of the cloud and less effective away from the center.¹² This might actually have measurable effects on e.g. average Gouy shift of the ensemble, and should also be studied in future work. You can use `prms['T_evolve_before_start']` for generating some correlation between position and velocity from thermal expansion before the start of the interferometer.

9) There's a couple assumptions I make in using the precompute tables that I expect to have minimal impact on results. First, if you plot the phase of the final Bragg wavefunction unwrapped over the precompute table. See further discussion the Precompute section above.

¹²Minimally the temperature is position dependent. The cavity interferometer experiment in our group found that the hotter atoms in the outside of the Raman sideband cooling beams were limiting the contrast of their lattice hold interferometer.

Chapter 6

Conclusions and next steps

6.1 Conclusions

The most important conclusion drawn from Chapter 2 is that systematic effects from the Schrodinger equation evolution of Bloch oscillations cancels out in our experiment to a very high degree of accuracy. The reason is that the differential effects between the upper and lower interferometers cancels gradients in the optical potential to first order, such that only higher order gradients remain. This, combined with the fact that the interferometers are very nearly overlapped at the time of Bloch oscillations in an SCI interferometer, ensures that systematic effects from the Bloch dynamics are not likely to be an issue moving forward.

This does not mean that there aren't any effects from Bloch oscillations, only that the effects directly from the Schrodinger equation are small. There are still large effects from deviations in the local k -vector of the laser since most of the momentum transfer in the SCI interferometer happens during Bloch oscillations.

Some conclusions from Chapter 3 are as follows. First, symmetric Bloch oscillations were proven to be an efficient way of generating extremely large momentum splittings in an atom interferometer. Second, the finite-temperature of an atom along the beam propagation axis breaks the momentum symmetry of the problem. This limits how efficiently you can drive symmetric Bloch oscillations for a given cloud temperature. Third, the reflection and recombination pulses are very sensitive to lattice depth fluctuations because they rely on coherent interference between different basis states of the wavefunction. This will likely lead to systematic effects. For all of these reasons, the technique might warrant further investigation.

Chapter 4 outlined the design and construction of the new vacuum system. Hopefully the notes about UHV design of a very large vacuum system will be useful to others in the future. One conclusion from this chapter was that extra parasitic lattices will likely limit the total phase achievable in this version of the experiment, and may introduce systematic effects.

Although Chapter 5 presents on-going work, there are some useful conclusions buried inside. First, it was possible to make a number of large improvements compared to the old Monte Carlo packages. These include dramatic speed-ups by using precomputed integration tables, the creating of a much more general simulation that includes Bloch oscillations and

new loss-of-contrast mechanisms, and that careful design of the structure of the Monte Carlo code will hopefully allow the code to be a foundation for future work on the subject. Initial experimental results also indicate that the Monte Carlo has been overestimating the size of potential systematic effects. Although this makes our current project of comparing Monte Carlo to experiment more difficult, it is promising for the future of the α experiment since distorted beams might be less of an issue than we thought.

6.2 Next steps

In the immediate future, we plan to continue working on comparing the Monte Carlo model to experimental data. This will be a continuation of the initial results from Chapter 5. Getting Monte Carlo to match the experimental data will likely take some time to get right, both on the experimental side and the Monte Carlo side.

The experiment is roughly within a factor of 2 in sensitivity compared to the old experiment, and my labmates are currently working to quantify this more precisely. Any easy improvements in the sensitivity of the experiment are worthwhile, because a factor of two improvement in sensitivity would mean 4x less data taking for a given error bar.

On the Monte Carlo side, more work needs to be done testing the sensitivity of simulation results to different parameters so we can see which parameters might need to be very carefully calibrated. In order to further debug possible factor of 2 errors in the code, the most useful next steps will be simple comparisons between the Monte Carlo and the experiment, and as many comparisons as possible that isolate different aspects of the code. More work also needs to be put into studying convergence of the simulations more carefully - the different types of convergence were outlined in Section 5.7.

Beyond the near-term goal of building a Monte Carlo model to match experimental data for severely distorted beams, we will need to start looking forward to an actual measurement of the fine-structure constant. There are two things that need to come together - statistical sensitivity, and systematic uncertainty. Very likely we'll focus initial efforts on optimizing the statistical sensitivity as much as possible in the existing system. After this is accomplished, we can start measuring or bounding various systematic effects and building up an error table. The gravity gradient item in the error table will be particularly difficult to measure accurately. It will likely require measurements of the gravity gradient as a function of height in the chamber, as was explored briefly in Fig. 4.1. Systematic effects related to intensity inhomogeneities of the laser beam will of course be another very important item in the error budget. Initial results from experimental measurements of severely distorted beams indicate that the effects are likely much smaller than we expected from the Monte Carlo model. This might mean that the effects are in the worst-possible scenario - large enough that they need to be measured, but small enough that it takes lots of integration time to be able to resolve the effects. Or, it could mean the effects are small enough that it might be relatively easy to bound the effects based on measurements of the laser beam intensity profile.

Additional steps along the path to measuring α might include any of the 'new ideas' presented in Section 4.6. I'm excited data from an optimized setup to measure the laser beam profile in-vacuum using Ramsey interferometry. I'm also nervous about results related to extra parasitic lattices in the chamber - this systematic will need to be explored quantitatively

in the future in order to gauge how big of an issue it might be.

In the very long-term, when people are thinking about designing a new experiment, I'm sure issues with the current design will inform future work. I would consider using two independent laser beams from the top and bottom so that 1) you don't lose a factor of two in laser power by combining beams on a 50/50 beamsplitter, 2) you don't have extra single photon scattering from retro-reflection, 3) you eliminate all of the parasitic interferometer issues in the current design, and 4) you might be able to find other ways to improve the overall efficiency of the laser frequency generation scheme. These are a lot of wins that will enable much more total phase accumulation in the interferometer. The cost of this will be that the difficulty of bounding systematic effects related to the spatial profile of the beams is now twice as difficult due to having two beams, and additional systematic effects related to the beam-beam alignment now need to be taken into effect.

I would also take the time to make the MOT and RSC regions of the vacuum chamber with much smaller diameters than in the current design, while still maintaining the clear aperture of the experiment. The large lever-arm the 3D MOT beams, the large travel path for the 2D MOT atomic beam, and the large distance between the 3D MOT coils all make co-alignment of the MOT very difficult. I think it's likely this is reducing our atom number by quite a bit, and as a result we're close to atom number shot noise when the old experiment wasn't. Similarly but not as critical is the large lever arm on the RSC beams, which makes alignment a bit more touchy. In order to make these chambers smaller, you'll need some sort of reducing flange between the interferometry region and the lower chamber, but if you have a tall enough room to work with this shouldn't be an issue.

One might also consider switching to e.g. Rubium instead of Cesium, since 1560nm fiber lasers can be doubled to get lots of power at 780nm. It's notable that using Rubidium and two independent laser beams would make our experiment much more similar to our French competitors [56]...

Strontium atom interferometry is still too new to have a good understanding of how useful it might be, but possibly this is a new direction in the future. It would be such a large change to the experiment that I don't think it'll ever happen in the Mueller group, but maybe it's worthy of a research proposal for a new faculty member trying to hop on the Strontium train.

Bibliography

- [1] Sven Abend. *Atom-chip gravimeter with Bose-Einstein condensates*. Phd thesis, Leibniz University Hannover, 2017.
- [2] Sven Abend, Martina Gebbe, Michael Gersemann, H Ahlers, Hauke Müntinga, E Giese, Naceur Gaaloul, Colja Schubert, Claus Lämmerzahl, Wolfgang Ertmer, Wolfgang P Schleich, and Ernst M Rasel. Atom-chip fountain gravimeter. *Physical Review Letters*, 117(20), Nov 2016.
- [3] Tatsumi Aoyama, Toichiro Kinoshita, and Makiko Nio. Theory of the anomalous magnetic moment of the electron. *Atoms*, 1(28), Feb 2019.
- [4] Peter Asenbaum, Chris Overstreet, Tim Kovachy, Daniel D. Brown, Jason M. Hogan, and Mark A. Kasevich. Phase shift in an atom interferometer due to spacetime curvature across its wave function. *Physical Review Letters*, 118(18), May 2017.
- [5] G. Audi, F.G. Kondev, M. Wang, B. Pfeiffer, X. Sun, J. Blachot, and M. MacCormick. The nubase2012 evaluation of nuclear properties. *Chinese Physics C*, 36(12):1157–1286, Dec 2012.
- [6] Satyanarayana Bade, Lionel Djadaojee, Manuel Andia, Pierre Cladé, and Saïda Guellati-Khelifa. Observation of extra photon recoil in a distorted optical field. *Physical Review Letters*, 121(7), 2018.
- [7] Felix Bloch. über die quantenmechanik der elektronen in kristallgittern. *Zeitschrift für Physik*, 52(7-8):555–600, 1929.
- [8] Rym Bouchendira, Pierre Cladé, Saïda Guellati-Khélifa, François Nez, and François Biraben. New determination of the fine structure constant and test of the quantum electrodynamics. *Physical Review Letters*, 106(8), Feb 2011.
- [9] W.L. Bragg. The diffraction of short electromagnetic waves by a crystal. *Proceedings of the Cambridge Philosophical Society*, 17:43–57, 1913.
- [10] Matthias Büchner, R. Delhuille, Alain Miffre, Cécile Robilliard, Jacques Vigué, and Caroline Champenois. Diffraction phases in atom interferometers. *Physical Review A*, 68(1), Jul 2003.

- [11] Malo Cadoret, Estefania de Mirandes, Pierre Cladé, Saïda Guellati-Khélifa, Catherine Schwob, François Nez, Lucile Julien, and François Biraben. Combination of bloch oscillations with a ramsey-bordé interferometer: New determination of the fine structure constant. *Physical Review Letters*, 101(23), Dec 2008.
- [12] Paolo Chiggiato. Outgassing properties of vacuum materials for particle accelerators. *Proceedings of the 2017 CERN–Accelerator–School course on Vacuum for Particle Accelerators, Glumslöv, (Sweden), 2017.*
- [13] Sheng-wei Chiow, Tim Kovachy, Hui-Chun Chien, and Mark A. Kasevich. $102\hbar k$ large area atom interferometers. *Physical Review Letters*, 107(13), Sep 2011.
- [14] Pierre Cladé. Bloch oscillations in atom interferometry. In *International School of Physics "Enrico Fermi"*, 2014.
- [15] Pierre Cladé, Manuel Andia, and Saïda Guellati-Khélifa. Improving efficiency of bloch oscillations in the tight-binding limit. *Physical Review A*, 95(6), 2017.
- [16] Pierre Cladé, Thomas Plisson, Saïda Guellati-Khélifa, François Nez, and François Biraben. Theoretical analysis of a large momentum beamsplitter using bloch oscillations. *The European Physical Journal D*, 59(3):349–360, Jul 2010.
- [17] Muon g-2 Collaboration. Measurement of the positive muon anomalous magnetic moment to 0.20 ppm. 2023.
- [18] Maxime Dahan, Ekkehard Peik, Jakob Reichel, Yvan Castin, and Christophe Salomon. Bloch oscillations of atoms in an optical potential. *Physical Review Letters*, 76(24):4508–4511, Jun 1996.
- [19] Savas Dimopoulos, Peter Graham, Jason Hogan, Mark Kasevich, and Surjeet Rajendran. Gravitational wave detection with atom interferometry. *Physics Letters B*, 678(1):37–40, Jul 2009.
- [20] Michele Dolfi, Adrian Kantian, Bela Bauer, and Matthias Troyer. Minimizing nonadiabaticities in optical-lattice loading. *Physical Review A*, 91(3), Mar 2015.
- [21] R. Driben, V. V. Konotop, T. Meier, and A. V. Yulin. Bloch oscillations sustained by nonlinearity. *Scientific Reports*, 7(1), Jun 2017.
- [22] Ethan R. Elliott, Markus C. Krutzik, Jason R. Williams, Robert J. Thompson, and David C. Aveline. Nasa’s cold atom lab (cal): system development and ground test status. *npj Microgravity*, 4(1), Aug 2018.
- [23] Brian Estey. *Precision measurement in atom interferometry using bragg diffraction.* PhD thesis, Dec 2016.
- [24] Brian Estey, Chenghui Yu, Holger S. P. Müller, Pei-Chen Kuan, and Shau-Yu Lan. High-resolution atom interferometers with suppressed diffraction phases. *Physical Review Letters*, 115(8), Aug 2015.

- [25] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse. Measurement of the electron magnetic moment. *Physical Review Letters*, 130(071801), Feb 2023.
- [26] J. Feldmann, K. Leo, J. Shah, D. A. B. Miller, J. E. Cunningham, T. Meier, G. von Plessen, A. Schulze, P. Thomas, and S. Schmitt-Rink. Optical investigation of bloch oscillations in a semiconductor superlattice. *Physical Review B*, 46(11):7252–7255, Sep 1992.
- [27] Simin Feng and Herbert G. Winful. Physical origin of the gouy phase shift. *Optics Letters*, 26(8):485, 2001.
- [28] Florian Fitzek, Jan-Niclas Kirsten-Siemß, Ernst Rasel, Naceur Gaaloul, and Klemens Hammerer. Accurate and efficient bloch-oscillation-enhanced atom interferometry. *arxiv*, Jun 2023.
- [29] Florian Fitzek, Jan-Niclas Siemß, Stefan Seckmeyer, Holger Ahlers, Ernst M Rasel, Klemens Hammerer, and Naceur Gaaloul. Universal atom interferometer simulation of elastic scattering processes. *Scientific Reports*, 10(1), Dec 2020.
- [30] J. B. Fixler, G. T. Foster, J. M. McGuirk, and M. A. Kasevich. Atom interferometer measurement of the newtonian constant of gravity. *Science*, 315(5808):74–77, Jan 2007.
- [31] Dong-Feng Gao, Jin Wang, and Ming-Sheng Zhan. Atomic interferometric gravitational-wave space observatory (aigso). *Communications in Theoretical Physics*, 69(1):37, Jan 2018.
- [32] Martina Gebbe, Jan-Niclas Siemß, Matthias Gersemann, Hauke Müntinga, Sven Herrmann, Claus Lämmerzahl, Holger Ahlers, Naceur Gaaloul, Christian Schubert, Klemens Hammerer, Sven Abend, and Ernst M Rasel. Twin-lattice atom interferometry. *Nature Communications*, 12(1), May 2021.
- [33] Martina Gebbe, Jan-Niclas Siemß, Matthias Gersemann, Hauke Müntinga, Sven Herrmann, Claus Lämmerzahl, Holger Ahlers, Naceur Gaaloul, Christian Schubert, Klemens Hammerer, and et al. Twin-lattice atom interferometry. *Nature Communications*, 12(1), 2021.
- [34] E. Giese, A. Roura, G. Tackmann, E. M. Rasel, and W. P. Schleich. Double bragg diffraction: A tool for atom optics. *Physical Review A*, 88(5), Nov 2013.
- [35] Peter Graham, Jason M Hogan, Mark A Kasevich, and Surjeet Rajendran. Resonant mode for gravitational wave detectors based on atom interferometry. *Physical Review D*, 94(10), Nov 2016.
- [36] Rudolf Grimm, Matthias Weidemüller, and Yurii B. Ovchinnikov. Optical dipole traps for neutral atoms. *Advances In Atomic, Molecular, and Optical Physics*, page 95–170, 2000.
- [37] P. Hamilton, M. Jaffe, P. Haslinger, Q. Simmons, H. Müller, and J. Khoury. Atom-interferometry constraints on dark energy. *Science*, 349(6250):849–851, Aug 2015.

- [38] D. Hanneke, S. Fogwell, and G. Gabrielse. New measurement of the electron magnetic moment and the fine structure constant. *Physical Review Letters*, 100(12), Mar 2008.
- [39] F. Heiße, S. Rau, F. Köhler-Langes, W. Quint, G. Werth, S. Sturm, and K. Blaum. High-precision mass spectrometer for light ions. *Physical Review A*, 100(2), Aug 2019.
- [40] Joseph Jachinowski, Natasha Sachdeva, and Tim Kovachy. Beam profiling with noise reduction from computer vision and principal component analysis for the magis-100 experiment. *2021 IEEE Conference on Antenna Measurements Applications (CAMA)*, 2022.
- [41] Matt Jaffe, Philipp Haslinger, Victoria Xu, Paul Hamilton, Amol Upadhye, Benjamin Elder, Justin Khoury, and Holger Müller. Testing sub-gravitational forces on atoms from a miniature in-vacuum source mass. *Nature Physics*, 13(10):938–942, Jul 2017.
- [42] Alan O. Jamison, Benjamin Plotkin-Swing, and Subhadeep Gupta. Advances in precision contrast interferometry with yb bose-einstein condensates. *Physical Review A*, 90(6), Dec 2014.
- [43] David Johnson. *Long baseline atom interferometry*. PhD thesis, Jun 2011.
- [44] Andrew Kerman. *Raman sideband cooling and cold atomic collisions in optical lattices*. Phd thesis, Jan 2002.
- [45] Andrey R Kolovsky and Hans Jürgen Korsch. Bloch oscillations of cold atoms in optical lattices. *International Journal of Modern Physics B*, 18(09):1235–1260, Apr 2004.
- [46] T. Kovachy, P. Asenbaum, C. Overstreet, C. A. Donnelly, S. M. Dickerson, A. Sugarbaker, J. M. Hogan, and M. A. Kasevich. Quantum superposition at the half-metre scale. *Nature*, 528(7583):530–533, 2015.
- [47] Tim Kovachy, Sheng-wey Chiow, and Mark A. Kasevich. Adiabatic-rapid-passage multiphoton bragg atom optics. *Physical Review A*, 86(1), Jul 2012.
- [48] Tim Kovachy, Jason M. Hogan, David M. S. Johnson, and Mark A. Kasevich. Optical lattices as waveguides and beam splitters for atom interferometry: An analytical treatment and proposal of applications. *Physical Review A*, 82(1), 2010.
- [49] Sascha Kulas, Christian Vogt, Andreas Resch, Jonas T Hartwig, Sven Ganske, Jonas Matthias, Dennis Schlippert, Thijs Wendrich, Wolfgang Ertmer, Ernst M Rasel, Marcin Damjanic, Peter Weßels, Anja Kohfeldt, E. Luvsandamdin, Max Schiemangk, Christoph Grzeschik, Markus Krutzik, Andreas Wicht, Achim Peters, and Sven Herrmann. Miniaturized lab system for future cold atom experiments in microgravity. *Microgravity Science and Technology*, 29(1-2):37–48, Oct 2016.
- [50] Robert Littlejohn. Physics 221ab, 2020-21 home page.
- [51] Albrecht Mager. Magnetic shields. *IEEE Transactions of Magnets*, 6(1), May 1970.

- [52] Vladimir S Malinovsky and Paul Schiff Berman. Momentum transfer using chirped standing-wave fields: Bragg scattering. *Physical Review A*, 68(2), Aug 2003.
- [53] Katherine McAlpine, Daniel Gochnauer, and Subhadeep Gupta. Excited-band Bloch oscillations for precision atom interferometry. *PHYSICAL REVIEW A*, 101(2), Feb 2020.
- [54] Kathryn A Moler, David J Weiss, Mark A Kasevich, and Steven Chu. Theoretical analysis of velocity-selective Raman transitions. *Physical Review A*, 45(1):342–348, Jan 1992.
- [55] Léo Morel, Zhibin Yao, Pierre Cladé, and Saïda Guellati-Khélifa. Determination of the fine-structure constant with an accuracy of 81 parts per trillion. *Nature*, 588(7836):61–65, 2020.
- [56] Léo Morel, Zhibin Yao, Pierre Cladé, and Saïda Guellati-Khélifa. Determination of the fine-structure constant with an accuracy of 81 parts per trillion. *Nature*, 588(7836):61–65, Dec 2020.
- [57] Holger Müller, Sheng-wei Chiow, and Steven Chu. Atom-wave diffraction between the Raman-Nath and the Bragg regime: Effective Rabi frequency, losses, and phase shifts. *Physical Review A*, 77(2), 2008.
- [58] Holger Müller, Sheng-wei Chiow, Quan Long, Sven Herrmann, and Steven Chu. Atom interferometry with up to 24-photon-momentum-transfer beam splitters. *Physical Review Letters*, 100(18), May 2008.
- [59] Holger Müller, Achim Peters, and Steven Chu. Müller, Peters, Chu reply. *Nature*, 467(7311):E2–E2, Sep 2010.
- [60] Holger Müller, Achim Peters, and Steven Chu. A precision measurement of the gravitational redshift by the interference of matter waves. *Nature*, 463(7283):926–929, Feb 2010.
- [61] Holger S. P. Müller, Sheng-wei Chiow, Sven Herrmann, and Steven Chu. Atom interferometers with scalable enclosed area. *Physical Review Letters*, 102(24), Jun 2009.
- [62] Gheorghe Nenciu. Dynamics of band electrons in electric and magnetic fields: rigorous justification of the effective Hamiltonians. *Reviews of Modern Physics*, 63(1):91–127, Jan 1991.
- [63] Qian Niu and M. G. Raizen. How Landau-Zener tunneling takes time. *Physical Review Letters*, 80(16):3491–3494, Apr 1998.
- [64] Eric B Norrgard, N. Sitaraman, John Barry, D J McCarron, Matthew Steinecker, and David DeMille. In-vacuum scattered light reduction with black cupric oxide surfaces for sensitive fluorescence detection. *Review of Scientific Instruments*, 87(5):053119–053119, May 2016.

- [65] Zachary Pagel, Weicheng Zhong, Richard H. Parker, Christopher T. Olund, Norman Y. Yao, and Holger Müller. Symmetric bloch oscillations of matter waves. *Physical Review A*, 102(5), 2020.
- [66] Richard H. Parker, Chenghui Yu, Brian Estey, Weicheng Zhong, Eric Huang, and Holger Müller. Controlling the multiport nature of bragg diffraction in atom interferometry. *Physical Review A*, 94(5), 2016.
- [67] Richard H. Parker, Chenghui Yu, Weicheng Zhong, Brian Estey, and Holger Müller. Measurement of the fine-structure constant as a test of the standard model. *Science*, 360(6385):191–195, Apr 2018.
- [68] Ekkehard Peik, Maxime Dahan, Isabelle Bouchoule, Yvan Castin, and Christophe Salomon. Bloch oscillations of atoms, adiabatic rapid passage, and monokinetic atomic beams. *Physical Review A*, 55(4):2989–3001, Apr 1997.
- [69] A. Peters, K. Y. Chung, B. Young, J. Hensley, and S. Chu. Precision atom interferometry. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 355(1733):2223–2233, Dec 1997.
- [70] Hannes Pichler, Andrew J Daley, and P Zoller. Nonequilibrium dynamics of bosonic atoms in optical lattices: Decoherence of many-body states due to spontaneous emission. *Physical Review A*, 82(6), Dec 2010.
- [71] William Press, Saul Teukolsky, William Vetterling, and Brian Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, third edition, 2007.
- [72] Avinoam Rabinovitch and J Zak. Does a bloch electron in a constant electric field oscillate? *Physics Letters A*, 40(3):189–190, Jul 1972.
- [73] Ernst M. Rasel, Markus K. Oberthaler, Herman Batelaan, Jörg Schmiedmayer, and Anton Zeilinger. Atom wave interferometry with diffraction gratings of light. *Physical Review Letters*, 75(14):2633–2637, 1995.
- [74] G. Rosi, F. Sorrentino, L. Cacciapuoti, M. Prevedelli, and G. M. Tino. Precision measurement of the newtonian gravitational constant using cold atoms. *Nature*, 510(7506):518–521, Jun 2014.
- [75] Gabriella Rosi, Guglielmo D’Amico, Luigi Cacciapuoti, Fiodor Sorrentino, Marco Prevedelli, Magdalena Zych, Āaslav Brukner, and Guglielmo M Tino. Quantum test of the equivalence principle for atoms in coherent superposition of internal energy states. *Nature Communications*, 8(1), Jun 2017.
- [76] D. M. Stamper-Kurn, G. E. Marti, and H. Müller. Verifying quantum superpositions at metre scales. *Nature*, 537(7618):E1–E2, Aug 2016.
- [77] Daniel A. Steck. Cesium d line data. Jan 1998.

- [78] J. K. Stockton, K. Takase, and M. A. Kasevich. Absolute geodetic rotation measurement using atom interferometry. *Physical Review Letters*, 107(13), Sep 2011.
- [79] Alex Sugarbaker. *ATOM INTERFEROMETRY IN A 10 M FOUNTAIN*. Phd thesis, Aug 2014.
- [80] Orazio Svelto. *Principles of lasers*. Springer, New York, 2016.
- [81] L. H. Thomas. Elliptic problems in linear differential equations over a network. *Watson Sci. Comput. Lab Report, Columbia University*, 1949.
- [82] Eite Tiesinga, Peter J. Mohr, David B. Newell, and Barry N. Taylor. Codata recommended values of the fundamental physical constants: 2018. *Reviews of Modern Physics*, 93(2), Jun 2021.
- [83] Le Tuan and Liviu F Chibotaru. A simple derivation of the landau–zener formula. *Physical Chemistry Chemical Physics*, 16(15):6942–6942, Jan 2014.
- [84] David George Voelz. *Computational Fourier optics : a MATLAB tutorial*. Spie Press, Bellingham, Wash., 2010.
- [85] Vladan Vuletic, Cheng Chin, Andrew J Kerman, and Steven Chu. Degenerate raman sideband cooling of trapped cesium atoms at very high atomic densities. *Physical Review Letters*, 81(26):5768–5771, Dec 1998.
- [86] C. Wachter. *Numerical Solution of the Time-Dependent 1D-Schrödinger Equation using Absorbing Boundary Conditions*. Bachelor’s thesis, University of Graz, 2017.
- [87] Gregory H. Wannier. Wave functions and effective hamiltonian for bloch electrons in an electric field. *Physical Review*, 117(2):432–439, Jan 1960.
- [88] Christian Waschke, Hartmut G. Roskos, Ralf Schwedler, Karl Leo, Heinrich Kurz, and Klaus Köhler. Coherent submillimeter-wave emission from bloch oscillations in a semiconductor superlattice. *Physical Review Letters*, 70(21):3319–3322, May 1993.
- [89] D. S. Weiss, B. C. Young, and S. Chu. Precision measurement of \hbar/m_{Cs} based on photon recoil using laser-cooled atoms and atomic interferometry. *Applied Physics B Lasers and Optics*, 59(3):217–256, 1994.
- [90] Jianming Wen, Yong Zhang, and Min Xiao. The talbot effect: recent advances in classical optics, nonlinear optics, and quantum optics. *Advances in Optics and Photonics*, 5(1):83, Mar 2013.
- [91] A. Wicht, E. Sarajlic, J. M. Hensley, and S. Chu. Phase shifts in precision atom interferometry due to the localization of atoms and optical fields. *Physical Review A*, 72(2), 2005.
- [92] S. R. Wilkinson, C. F. Bharucha, K. W. Madison, Qian Niu, and M. G. Raizen. Observation of atomic wannier-stark ladders in an accelerating optical potential. *Physical Review Letters*, 76(24):4512–4515, Jun 1996.

- [93] Eric Wodey, Dorothee Tell, Ernst M Rasel, Dennis Schlippert, R Baur, U Kissling, B. Kölliker, Michael Lorenz, M. Marrer, U. Schläpfer, Mariana Widmer, Christian Ufrecht, S Stuibler, and Peter Fierlinger. A scalable high-performance magnetic shield for very long baseline atom interferometry. *Review of Scientific Instruments*, 91(3):035117–035117, Mar 2020.
- [94] Peter Wolf, Luc Blanchet, Christian Borde, Serge Reynaud, Christophe Salomon, and Claud Cohen-Tannoudji. Classical and quantum gravity does an atom interferometer test the gravitational redshift at the compton frequency? *Classical and Quantum Gravity*, 28(145017), Jun 2011.
- [95] Peter Wolf, Luc Blanchet, Christian J Bordé, Serge Reynaud, Christophe Salomon, and Claude Cohen-Tannoudji. Atom gravimeters and gravitational redshift. *Nature*, 467(7311):E1–E1, Sep 2010.
- [96] Xuejian Wu, Zachary Pagel, Bola S. Malek, Timothy H. Nguyen, Fei Zi, Daniel S. Scheirer, and Holger Müller. Gravity surveys using a mobile atom interferometer. *Science Advances*, 5(9), 2019.
- [97] Victoria Xu. *Lattice atom interferometry in an optical cavity*. Phd thesis, 2020.
- [98] Victoria Xu, Matt Jaffe, Cristian D. Panda, Sofus L. Kristensen, Logan W. Clark, and Holger Müller. Probing gravity by holding atoms for 20 seconds. *Science*, 366(6466):745–749, Nov 2019.
- [99] Chenghui Yu. *Measuring the fine structure constant with a state-of-the-art atom interferometer*. PhD thesis, Aug 2018.
- [100] Chenghui Yu, Weicheng Zhong, Brian Estey, Joyce Lok, Richard Parker, and Holger S. P. Müller. Atom-interferometry measurement of the fine structure constant. *Annalen der Physik*, 531(5):1800346–1800346, May 2019.
- [101] Alexey V Yulin, Rodislav Driben, and Torsten Meier. Bloch oscillations and resonant radiation of light propagating in arrays of nonlinear fibers with high-order dispersion. *Physical Review A*, 96(3), Sep 2017.
- [102] J Zak. Comment on “time evolution of bloch electrons in a homogeneous electric field”. *Physical Review B*, 38(9):6322–6323, Sep 1988.
- [103] Clarence Zener. Non-adiabatic crossing of energy levels. *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, 137(833):696–702, Sep 1932.
- [104] Bo Zhang, Josiane Zerubia, and Jean-Christophe Olivo-Marin. Gaussian approximations of fluorescence microscope point-spread function models. *Applied Optics*, 46(10):1819, Mar 2007.

- [105] Weicheng Zhong, Richard Parker, Zachary Pagel, Chenghui Yu, and Holger Mueller. Offset simultaneous conjugate atom interferometers. *Physical Review A*, 101(053622), May 2020.
- [106] L. Zhou, Z. Y. Xiong, W. Yang, B. Tang, W. C. Peng, K. Hao, R. B. Li, M. Liu, J. Wang, and M. S. Zhan. Development of an atom gravimeter and status of the 10-meter atom interferometer for precision gravity measurement. *General Relativity and Gravitation*, 43(7):1931–1942, Mar 2011.
- [107] Lin Zhou, Shitong Long, Biao Tang, Xi Chen, Fen-Biao Gao, Wencui Peng, Weitao Duan, Jiaqi Zhong, Xiong Zongyuan, Jin Wang, Yuan-Zhong Zhang, and Mingsheng Zhan. Test of equivalence principle at 10^{-8} level by a dual-species double-diffraction raman atom interferometer. *Physical Review Letters*, 115(1), Jul 2015.