

Shaken Lattice Interferometry

by

Carrie A. Weidner

B.S., Engineering Physics, University of Colorado at Boulder, 2010

B.S., Applied Mathematics, University of Colorado at Boulder, 2010

M.S., Physics, University of Colorado at Boulder, 2015

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

2018

This thesis entitled:
Shaken Lattice Interferometry
written by Carrie A. Weidner
has been approved for the Department of Physics

Prof. Dana Z. Anderson

Prof. James K. Thompson

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Weidner, Carrie A. (Ph.D., Physics)

Shaken Lattice Interferometry

Thesis directed by Prof. Dana Z. Anderson

Since the first demonstrations in 1991, atom interferometry has been a burgeoning field of research. The work done in this field is motivated by the potential sensitivity improvements that atom-based devices can have over the current state-of-the-art light- and MEMS-based devices. This dissertation presents a new and unique approach to atom interferometry in that we perform the basic interferometric sequence of splitting, propagation, reflection, reverse-propagation, and recombination with atoms trapped in a phase-modulated (shaken) optical lattice. In both simulation and experiment we demonstrate a one-dimensional shaken lattice interferometer configured as an accelerometer. The interferometry sequence is developed through the use of learning and optimal control algorithms that allow us to implement the desired state-to-state transformations and perform the desired operations, e.g. splitting and recombination of the atoms trapped in the lattice. This device has a sensitivity that scales as the square of the interrogation time and an ability to distinguish both the magnitude and sign of an applied acceleration signal. Furthermore we show that we can tailor the transfer function of the interferometer to be sensitive to a signal of interest, e.g. an AC signal of a given frequency. Finally, we explore the analytics of shaken lattice interferometry and offer some suggestions as to the future of this new technology.

Dedication

To Todd Weidner, Andrew Weidner, and Corey Tollerud, the three finest men I know.

Acknowledgements

This dissertation and the work that led up to it would not be possible without the love and support of an army of people who stood behind me throughout both the past six years and my life. I have been so absurdly lucky to land where I am after the tumult that came after my undergraduate years that I am often in awe of my good fortune.

My first acknowledgements must be to the three men to whom I have dedicated this work. My father Todd Weidner has been my greatest inspiration and role model. From the time he held me in his hand (yes, hand!) to now, he has always believed in me and encouraged me to follow my dreams. It is because of him that I have a solid work ethic, a caustic sense of humor, and a profound love for the Chicago Cubs (especially the 2016 World Series championship team). My brother Andrew is one of my best friends and closest confidants. It was such a pleasure to watch you grow during your undergraduate years at CU, and I truly treasure the love and friendship that we have. Finally, I want to thank Corey Tollerud. For the past 9.5 years, you have stood by me through the good and the bad. You make me laugh, you make me think, and you support every one of my hare-brained ideas and ambitions. No matter where life takes us, the time we've had together has been the best of my life, and I cannot wait to go on this next adventure together.

My mother, Beth, has always encouraged my ambitions. She fostered in me a love of books, and every now and again, I get to read one that doesn't have any equations in it. She has supported my writing through all of my schooling, and I can only hope that this dissertation brings her pride and joy. My stepmother Deedy Abernethy came into our lives many years ago, and she loved my brother and me with reckless abandon. She also taught us how to love spicy food (this is very

important) and dominoes.

I have to thank my advisor, Dana Anderson, for believing in me and supporting me. I came in your lab little more than an overgrown child, and I will leave it an adult. You taught me grace and maturity, as well as some science along the way. Thank you for letting me take this idea and run with it. I've valued the independence that you allowed me to have and the leadership that came along with it. What you and I have built here is pretty darn cool, and I'm thankful that I stumbled into your lab one December afternoon and asked for a job.

I owe a great debt to my undergraduate advisor Heather Lewandowski. She was crazy enough to hire me as a freshman and see through the many obvious flaws to find a budding young scientist underneath the cracks. Because of her, I had a safe place to learn and grow as an undergraduate. She then became a mentor, asking the hard questions when necessary and picking me up when I had failed and fallen down. I hope that with this, I have been able to take a step towards being a professional scientist in my own right, and I could not have done it without her. I also can't think of anyone I'd rather have in an infield with me as I pitch (except for maybe Raina Gough).

I have a number of friends to thank, and I fear that the pages of this acknowledgement will not do them justice. First, I want to thank my labmates (and hockey teammates) Seth Caliga and Cam Straatsma. I miss skating with you guys and will value your friendship throughout our lives. My other labmates through the years deserve recognition: Kai Hudek, Hoon Yu, Taylor Froelich, and Ava Ashby (who, after the apocalypse, will be found working quietly at her desk, unaware of any disturbance). To my current labmates, Brad Dinardo, Curtis Rau, and Aaron Mutchler: you three have been so fun to work with. I have to thank the JILA machine shop staff (especially Hans, Todd, Kim, and Kyle: I've managed to keep all of my fingers through grad school), the electronics shop staff (especially Terry, for letting me ask him an infinite number of ridiculous questions), the computing staff (Jim, JR, and Corey: your help kept my experiment running and the angry theorists off my back), and everyone in between. JILA is a fantastic place to learn, work, and grow. I would not be half of the scientist I am today without the support structure it provided.

I must acknowledge everyone who was ever a JILA Monster or a Moosecat, you all provided

me with a valuable hour of social interaction, even if I had to go back to the lab afterwards. Marika Meertens deserves a special shout out for being a bastion of love and support and mountain biking through these years. I want to thank my oldest friend Nicki Seminara for her friendship and tolerance of my asocial tendencies. This goes out to the crew who lived at 1220 East Ridge (especially Sean Sammons (and Ily Reiling!) and Madison Kittleson), to Enrique Gonzales-Moore, Brian Lynch, Jordan McConnell and John Palen. A special thank you to Barb Brown and Ed, Jonathan, and Kathryn Tollerud: you adopted my family into your family and provided me with such love (and food) during these last few years. (A special thanks to Ed Tollerud for proofreading this entire thesis.) Further thanks to Jan Hartley and Tim Cunningham for always thinking to invite us to dinner, take me biking, and give me a hard time. Much love to my graduate student buddies Courtney Peck, Mathis Habich, Will Lewis, John Bartolotta, and Jennifer Ellis. Likewise to my high school buddies: Amy Reed (and Sally Stermer) and the Triumvirate. There's no way I got everyone in here, and to those I missed (and you know who you are), please feel free to berate me at your earliest convenience.

Finally, I have to thank the people that first fostered my academic spark: Joann Taylor, Jeanne Westmoreland, and Dee Willars. You took an angry and frustrated teenager and helped build a (reasonably well-functioning) human being. This thesis is for you.

Go Cubs!

Contents

Chapter

1	Introduction	1
1.1	Motivation	1
1.2	The story of the shaken lattice interferometer	5
1.3	Outline of this dissertation	7
2	The physics of ultracold atoms in optical lattices	8
2.1	Cold and ultracold atoms, briefly	8
2.2	Optical dipole traps	10
2.3	Lattices	11
2.3.1	The Bloch basis	12
2.3.2	The Bose-Hubbard Hamiltonian, briefly	15
2.4	Atoms in a shaken optical lattice	17
3	Simulation and optimization methods used in this thesis	20
3.1	The split-step method	20
3.1.1	How the split-step method works	20
3.1.2	Errors in the split-step method	21
3.2	Motivating the use of learning algorithms	23
3.3	The genetic algorithm	25
3.4	The Krotov method	30

3.4.1	A general formulation of the Krotov method	31
3.4.2	The Krotov method applied to the linear TDSE	33
3.4.3	Numerical implementation of the Krotov method	34
3.4.4	The Krotov method applied to the shaken lattice	35
3.5	The CRAB method	37
3.5.1	Formulation of the CRAB method	37
3.5.2	The Nelder-Mead algorithm	38
3.5.3	Avoiding local minima in the CRAB algorithm	40
3.6	What’s next	41
4	Simulating the shaken lattice interferometer	42
4.1	The simulation and optimization of the shaken lattice interferometer	42
4.1.1	The “ingredients” of interferometry	43
4.1.2	Fitness and error functions	47
4.1.3	Learning curves	48
4.2	Adding a signal to the shaken lattice interferometer	49
4.3	Proof-of-principle results using the Krotov method	51
4.3.1	Initial interferometer optimization	51
4.3.2	Adding a signal	53
4.4	Determining the interferometer sensitivity	55
4.4.1	Initial optimization with the genetic algorithm	55
4.4.2	Robustness and stability of the shaken lattice interferometer	60
4.4.3	Transfer function optimization	65
4.4.4	Calculating the interferometer sensitivity	72
4.5	Next steps	76
5	Building a shaken lattice interferometer	78
5.1	From dispenser to BEC on a chip	78

5.1.1	The laser system	79
5.1.2	The physics package and vacuum cell	80
5.1.3	BEC production in the experimental system	86
5.2	Building an optical lattice around a BEC system	89
5.3	Shaking the lattice and measuring the results	98
5.3.1	Electro-optic shaking	100
5.3.2	Shaking via frequency modulation	102
5.3.3	Advantages and disadvantages of the two shaking methods	106
5.4	What's next	107
6	Optimization and characterization of a shaken lattice interferometer	108
6.1	Closed-loop optimization in the experimental system	108
6.2	Learning in the closed-loop system	110
6.2.1	Narrowing the frequency band used in the optimization	110
6.2.2	An experimental learning curve	113
6.2.3	The effect of a finite momentum width	115
6.3	Experimental shaken lattice interferometry	115
6.3.1	Optimizing the interferometers	116
6.3.2	Determining the interferometer sensitivity	119
6.3.3	Experimental confirmation of other interesting aspects of shaken lattice in- terferometry	124
6.4	Next steps	126
7	An analysis of splitting in the shaken lattice interferometer	129
7.1	Introduction	129
7.2	Band decomposition of the lattice wavefunction	131
7.3	The analytics of phase modulation of an optical lattice	136
7.4	Single-frequency splitting	137

7.4.1	Theory and simulation of single-frequency splitting	137
7.4.2	Experimental results	141
7.4.3	How much must we split?	143
7.5	Optimization of higher-order splitting using band-to-band transitions	145
7.6	Conclusion	152
8	Shaken lattice interferometry in the future	153
8.1	Future experimental improvements	154
8.2	Future simulation improvements	157
8.2.1	A shaken lattice gyroscope	158
8.2.2	A semiclassical Hamiltonian	158
8.2.3	Towards simulations of a shaken lattice gyroscope	158
8.2.4	Experimental rotation sensing	160
8.3	Final thoughts	161
	Bibliography	162

Tables

Table

4.1	Genetic algorithm results, best of 5 optimization runs.	59
7.1	Band transition frequencies $f_{r,r'}$, $V_0 = 10E_R$. The single (double) asterisk marks frequencies with matrix element overlaps $M_c(M_s) > 0.1$ (see Fig. 7.6 and Sec. 7.3). .	135

Figures

Figure

- 1.1 A diagram of an optical Mach-Zehnder-style Sagnac interferometer. The light from a laser is split and recombined with two beamsplitters (BS). Mirrors (M) direct the split light around an enclosed area $|\vec{A}| = L_1 L_2$. If the system is rotating with rate $\vec{\Omega}$ (about an axis directed out of the page and parallel to the area vector \vec{A}), then there will be a phase shift $\Delta\Phi$ between the two recombined beams given by Eq. (1.1). This rotation shift can be detected by counting the fringes that pass one of the two detectors (D) at the two output ports of the recombining beamsplitter. 2
- 2.1 The ground (green) and first four excited Bloch states (red, blue, magenta, and cyan, respectively) for lattice depths of (a) $\alpha = 10$ and (b) $\alpha = 20$. The lattice depth is shown via a black, dotted line. 15
- 2.2 The ground ($r = 0$) Bloch state for (a) $\alpha = 10$ and (b) $\alpha = 20$ at quasimomentum $q = 0$ (red) and $q = \hbar k_L$ (green). The plots show the real part of the wavefunction and are normalized such that the maximum amplitude of the state is one; the imaginary part is zero. For a spatial reference, the function $\cos(2k_L x)$ is plotted (blue), representing the lattice. 16

- 2.3 The (a) position- and (b) momentum-space representations of the ground (red) and first (green) excited Bloch states of an optical lattice with $\alpha = 10$. As in Fig. 2.2 the plots in (a) are of the real part of the wavefunction and are normalized such that the maximum amplitude of the state is one; the imaginary part of the wavefunctions are zero. For a spatial reference, the function $\cos(2k_L x)$ is plotted (blue), representing the lattice. In plot (b) the relative population in each of the $2n\hbar k_L$ states is shown by the height of the bars. More plots of the Bloch state momentum populations are shown in Fig. 7.1. 16
- 3.1 The phase error in radians from Eq. (3.3) plotted vs. timestep δt in scaled units. For the typical timestep used in the simulations, $\delta t = 1/1000$, limiting the phase error to tens of microradians. 22
- 3.2 The error between the final state and the finer-time reference state in Eq. (3.4) plotted vs. timestep δt in scaled units for sinusoidal shaking with various frequencies ω . The case where no shaking is applied is shown in red, which demonstrates the propagation of the initial eigenstate in a stationary lattice. For the typical timestep used in the simulations, $\delta t = 1/1000$, and for all shaking frequencies the percent error is $< E_{lim} = 10^{-3} \%$ 24
- 3.3 A block diagram illustrating the steps taken in the GA. Given the initial and desired states, the first generation $G = 1$ of A individuals is randomly generated with unique shaking functions $\phi_\alpha(t)$. We then solve the TDSE for each individual in the generation. These results are then used to produce the next generation of individuals, denoted $G = 2$. After the j^{th} run of the simulation a generation $G = j + 1$ results from the mixing of the previous generation's individuals (see Fig. 3.4). Once the convergence criterion is met, the GA stops. 27

- 3.4 Each iteration, the genetic algorithm mixes the best individuals from the previous generation to create “children” that populate the next generation. Each individual is a vector of values where red corresponds to a minimum value and purple corresponds to a maximum value. In each case the indices of crossover or mutation are chosen at random. The numbers label the index and their color is chosen for legibility. The methods presented here are adapted from reference [92]. (a) One-point crossover at index 6. (b) Two-point crossover at indices 5 and 9. (c) Random mutation of the value at index 2. (d) Change of the value at index 5 through a small “creep” of the value. 29
- 4.1 A cartoon of (a) an optical Mach-Zehnder interferometer and (b) a Michelson interferometer, with the basic “ingredients” of the interferometry sequence labeled. . . . 44
- 4.2 A cartoon of the full interferometer sequence, with atom position plotted along the y -axis versus time on the x -axis. Blue clouds represent atom wavepackets interacting with the shaken lattice. Atoms begin in the ground Bloch state of the lattice (Fig. 2.3) and are split into two oppositely propagating wavepackets. The atoms are then reflected, reverse-propagated, and recombined back into their initial state (in the absence of a signal), thus completing the interferometer sequence. 46
- 4.3 Learning curves for first-order splitting using the genetic algorithm showing (a) the fitness of the best individual (black) and the average fitness of all individuals (red) and (b) the percent error of the final state relative to the desired split state. 48
- 4.4 Learning curve for first-order splitting using the Krotov method showing (left axis, blue) the percent error of the final state relative to the desired split state and (right axis, red) the λ_a parameter as a function of iteration number. 50

4.10	Variation of the final state with varying noise amplitudes added to the optimized splitting shaking function. The results shown here are the average of 5 runs with random white Gaussian noise added to the shaking function $\phi(t)$. Error bars give the standard deviation of the variations. Noise amplitude is given as a fraction of the maximum of the $\phi(t)$ shown in Fig. 4.8. The red line marks a variation of 1%.	64
4.11	Variation of the final split state after the addition of spurious lattice potentials with varying phase due to unwanted reflections with reflection amplitudes of $\epsilon = 0.1\%$ (red), 1% (blue), and 4% (cyan).	66
4.12	The response of an interferometer optimized in the presence of a DC bias acceleration $a_{\text{DC}} = 0.76 \text{ m/s}^2$. The interferometer response is clearly minimized in the vicinity of a_{DC} and increases away from this bias. This shows that shaken lattice interferometry can be used to reject a DC bias of a given magnitude or measure perturbations around this bias.	68
4.13	The reciprocal interferometer sequence. The reciprocal interferometer modifies the standard Michelson interferometer sequence shown in Fig. 4.2 so that the atoms travel a fully symmetric path. This configuration is designed to be sensitive to AC accelerations and immune to DC accelerations.	70
4.14	Response of the reciprocal (red, solid) and non-reciprocal (blue, dashed) interferometers to a sinusoidal signal as in Eq. (4.7). The response is given in terms of the variation $D_{\text{f,i}}$ between the optimized final state of the interferometer and the final state after shaking with the applied signal as in Eq. (3.4). The total interrogation time of each interferometer is 2.008 ms and the amplitude of the applied acceleration signal is $a_0 = 0.115 \text{ m/s}^2$. The reciprocal interferometer is 20 times more sensitive than the non-reciprocal interferometer to a signal with $f \approx 7 \text{ kHz}$, showing that shaken lattice interferometry may be modified to tailor the interferometer response to an AC signal.	71

- 4.15 The minimum detectable acceleration δa scaled by $\sqrt{N_a}$ plotted on a log scale versus the interrogation time T_I . The black points are simulation results, and the red line is a fit of the form CT_I^{-n} . Here, lower values of δa correspond to a more sensitive interferometer. 74
- 4.16 The acceleration sensitivity δa (relative to g) for varying atom numbers N_a at $T_I = 10$ ms (red), 100 ms (blue), and 1 s (black). 75
- 5.1 Schematic of the lattice laser system with servo. The angle of the first-order AOM output is not shown for ease of drawing. The servo works by comparing a control voltage from LabVIEW to the voltage output of the photodetector. The error signal is sent to a mixer that attenuates the output of a voltage-controlled oscillator (VCO) that controls the AOM frequency. The attenuated signal is then amplified by a 2 W RF amplifier and sent to the AOM. The symbols represent the following: lens (L), mirror (M), shutter (S), optical Faraday isolator (ISO), half-wave plate (HWP), and quarter-wave plate (QWP). Other symbols are defined in the main text. 81
- 5.2 The two-chamber vacuum system used in the experiments described in this thesis. In our system the cell is oriented so that the atom chip is on top. The 2D MOT chamber contains a rubidium dispenser and an NEG. The atoms are pushed up through a pinhole between the 2D cell and the spherical cube (yes, that's actually what it's called) and into the 3D cell. A 0.5 L/s ion pump pumps on the 3D cell, which is topped by the atom chip (the pinhole also serves as a differential pumping stage). Copper wires patterned on the vacuum side of the atom chip are barely visible through the 3D cell. While the 2D cell is not AR-coated, the 3D cell is. The copper pinch-off is also labeled; this separates the vacuum chamber from the larger bake-out station. 83

- 5.3 A view from the underside of the 2D cell, up through the aperture on the cell bottom and focused on the pinhole. A light is being shone through the pinhole to backlight it for clarity. The 0.5 L/s ion pump can be seen off to the right, and the copper pinch-off is to the left. The push beam that moves the atoms from the 2D cell/MOT to the 3D cell/MOT travels through the aperture in the bottom of the cell and through the pinhole. The metallized pads on the underside of the cell (surrounding the aperture) make electrical connections to the rubidium dispenser and NEG. Connections are made to the pads using conductive tape and colloidal silver. 84
- 5.4 (a) CAD and (b) a photograph of the physics package provided by ColdQuanta. The physics package is roughly 28x28x32 cm in size, and a cell similar to that shown in Fig. 5.2 is bolted down in the center of the package. The lower two tiers contain the optics needed to produce a 2D MOT and the push beam that pushes atoms up through the pinhole to the 3D cell. The upper two tiers contain the optics for 3D MOT production, optical pumping, and probing. In the CAD drawing, the cooling beams are shown in red, and the imaging beam is gray. In the photograph, the system is rotated clockwise 90° from the CAD view. The coils surrounding the cell in the CAD drawing are not installed in the photograph, so the photo shows the 3D cell poking above the top tier of the physics package in the center. The cable connecting the 0.5 L/s ion pump to its controller is shown running off to the right. 85
- 5.5 Picture of the atom chip used in the experiments described in this work. The chip is made of high-conductivity silicon coated in an insulating dielectric and measures roughly 2.5 cm by 2.5 cm by 420 μm thick. All wires are patterned in gold on the ambient side of the chip. The only two wires used here are the thicker main wire (labeled by the current I_0) and the thinner dimple wire (I_d); these wires intersect in the center of the chip and the atoms are trapped below this intersection. 88

- 5.6 Density plots of the magnetic trap potential used to hold the atoms in (a) the initial chip trapping configuration, (b) the compressed evaporation trap, and (c) the waveguide-like decompressed pre-lattice trap. The color indicates the magnetic field in Gauss. The calculations were made using the Biot-Savart law and assume infinitesimally thin wires. Note the difference in scale in plot (c) relative to (a) and (b). 88
- 5.7 The relevant currents run through the big-Z coil (black, scaled down by 50%), X-bias coil (red) Y-bias coil (blue), Z-bias coil (gray), main chip wire (magenta), and dimple wire (cyan) during the following experimental stages: (a) optical pumping, (b) big-Z capture and transport, (c) initial chip loading, (d) chip compression and first-stage evaporation, and (e) the rest of the RF evaporation. The optical pumping stage usually lasts for only 0.5 ms and is extended in time in this plot for clarity. . . 90
- 5.8 Changes in currents (left axis) and lattice light (magenta, right axis) as we ramp from the BEC trap to lattice. The X-bias and Y-bias coil currents are shown in black and red, respectively, and the main and dimple wire currents are shown in blue and gray, respectively. The amplitude of the lattice beam is given in arbitrary units related to the voltage sent to the intensity servo. Due to the fact that the servo has trouble servoing around zero intensity, the lattice beam is held at an arbitrary high value (with the mechanical shutter closed) as the atoms are evaporated to BEC (end of evaporation shown in a), then ramped down to zero as the atoms are loaded into the waveguide trap (b). The lattice light intensity is nearly zero for about 1 ms as the mechanical shutter is opened (c), then it ramps on after the chip wires are ramped off (d). The X-bias coil is not ramped to zero, but rather a constant value that allows the atoms to maintain their spin polarization in the lattice. The reason for this will be explained in more detail in Chapter 6. 91

- 5.9 Schematic of the optics used in the lattice and imaging systems. The cell is shown outlined in purple, with the atoms (blue) in the center. The probe light for the two axes of imaging are shown in light red, and the lattice light is shown in dark red. Elements surrounded by the black dashed box are contained on the physics package, elements surrounded by the dark gray dashed lines are contained on the secondary lattice board, and the lattice optics surrounded by the light gray box are on a three-axis translation stage. All abbreviations are explained in the text. 93
- 5.10 The entire experiment, including the physics package (black, center) and the auxiliary lattice board surrounding it. (a) View of the experiment showing the lattice launcher (right) and the EOM and lattice retro-reflecting mirror (left) on the auxiliary lattice board. (b) View of the experiment along the direction of lattice propagation showing the lattice launcher (left) and the second axis of imaging launcher (right) on the auxiliary lattice board. The coils used to apply the magnetic field gradient to simulate a signal are shown surrounding the coil assembly (red coils surrounding black coil structure). (inset) A close-up view of the cell surrounded by the coil assembly (coils in orange). The red gradient coil can also be seen to the right of the cell. 94
- 5.11 Stability of (a) the tweaker plate (TP) and (b) the angular mirror directing the lattice beam into the cell (see Fig. 5.9). The TP displacement is shown for plate angles of 0° (red), 10° (blue), and 20° (black) relative to the plane orthogonal to the incoming laser beam. For maximum resolution (in terms of screw turns per displacement), the TP is mounted near 0° . The number of screw turns is calculated using the different mounts for these two optics: (a) a Newfocus 9873-K mount and (b) a Thorlabs Polaris K1T mount. The angular mirror is extremely sensitive and thus a high-quality mirror mount is used to house this optic. In both plots the experimental image resolution is shown with black dotted lines. 96

5.12	Drift δz of the z -position of the pre-lattice waveguide magnetic trap as a function of drift δI of the current in the x-bias field (red), y-bias field (blue), main wire (black), and dimple wire current (blue). Calculations show that the drift in the x - and y -directions is negligible, as confirmed by the fact that only the z -position of the lattice beam needs adjustment over the course of a day.	99
5.13	Schematic of lattice shaking with an EOM. The lattice is shown in red, while the atoms are blue and the cell is gray. Amplifier gains are given by the numbers on the amplifier elements. The details of shaking are explained in the text.	101
5.14	Double-pass phase change of the EOM as a function of the voltage output by the AWG (the applied voltage to the crystals is this voltage amplified by $40\times$). The error bars are high for the low-voltage points due to the difficulty in measuring the phase at low applied voltage. A linear fit $\phi_{\text{lin}}(V) = aV$ to the data is shown in red, and a quadratic fit $\phi_{\text{quad}}(V) = aV^2 + bV$ is shown in blue. The linear fit gives $a = 0.746(6)$ rad/V and the quadratic fit gives $a = 0.12(2)$ rad/V ² and $b = 0.64(1)$ rad/V. The linear fit is sufficient for our purposes.	103
5.15	The unshifted lattice potential (gray) and the lattice potential when a frequency shift of $\delta\nu = 300$ MHz is applied (black), showing the predicted π phase shift. The plot extends over $x = d_0 \pm 3\lambda_L$	104
5.16	Phase error (as a fraction of π phase shift) as a function of distance (in wavelengths λ_L) for the experimentally relevant case where $d_0 = 0.25$ m and $\lambda_L = 85$ nm. The error is less than 4% over a range of 10^4 wavelengths.	105
6.1	A block diagram of the closed-loop optimization protocol used in the experiments shown in this chapter. The control code runs the experiment, analyzes the images produced by the camera, and updates the lattice shaking function based on the optimization algorithm.	111

- 6.2 Results of scanning the shaking frequency from (a) 10 – 18.25 kHz, (b) 18 – 28 kHz and (c) 26 – 42 kHz and measuring the percent error of the final state with respect to the split state. The different colored curves simply denote different experimental runs. The points (error bars) denote the mean (standard deviation) of the error. The data was taken over the course of three days (each separate plot corresponds to one day of data), and this explains small “jumps” in the error as one moves from plot to plot. 112
- 6.3 Sample learning curves for the experiment. The points (error bars) show the average (standard deviation) of three points. The optimization resulting in the learning curve in (a) was done first, then the results were improved upon in a second optimization run that took the best results of optimization (a) and produced the learning curve in (b). 114
- 6.4 Two example shaking functions. (a) Splitting and recombination shaking protocols, totaling $T_I = 0.4$ ms. (b) Splitting and recombination protocol from (a), but with 4 propagation steps added in between, resulting in an interrogation time $T_I = 2$ ms. We optimize the interferometer so that the atoms remain split at the end of each propagation step. In both cases the second half of the shaking protocol is simply the reflection of the first half. 118
- 6.5 Percent error in the (a) splitting and (b) recombination protocols as a function of (a) the splitting time $T_I/2$ and (b) the total interrogation time T_I . The insets show example experimental images from each protocol, where the color represents optical density (OD). (a, inset) An image from optimized splitting of the atoms into equal population in the $\pm 2\hbar k_L$ momentum states and (b, inset) recombining the atoms into the ground Bloch state. 120

- 6.6 Data showing the atoms' acceleration as a function of magnetic field gradient G applied via two coils placed on the side of the cell. The coils (visible in Fig. 5.10(b)) are oriented so that the current through them goes in opposite directions, resulting in a magnetic field gradient felt by the atoms. 121
- 6.7 Minimum detectable effective acceleration δa plotted as a function of interrogation time for each of the five interferometers optimized for this work (black) and fit (red) to $f(T_I) = aT_I^b + c$. The scaling value b is consistent with the expected T_I^2 scaling, and the offset c arises due to imaging noise and is measured experimentally. (inset) Data taken with no atoms present (blue) and no shaking applied to the atoms (red) showing no signal other than imaging noise. Blue data is scaled by the ratio of the relative atom number as explained in the text. 123
- 6.8 Momentum population of the atoms after the $T_I = 2$ ms interferometer sequence as a function of the applied acceleration signal. Atoms in the $2n\hbar k_L$ state are denoted by open blue circles ($n = -2$), blue crosses ($n = -1$), black dots ($n = 0$), red plusses ($n = 1$) and red asterisks ($n = 2$). As the applied signal is varied away from zero, we can distinguish positive and a negative signals. The dotted lines are cubic spline fits to guide the eye. 125
- 6.9 As in Fig. 6.8, plot of the momentum state variation as a function of applied acceleration with the biased interferometer, showing variation of the final state as the acceleration is varied from the bias value of $a_{\text{bias}} = -0.71$ m/s² (black dashed line). Data points and splines colored as in Fig. 6.8. (inset) An experimental image of the optimized split state in the biased interferometer. OD is indicated by the colorbar on the right. 127

- 7.1 (a-h, in order of increasing band number from $r = 0$ to 7) Momentum state populations and energies for the first 8 Bloch states $|r\rangle$ for atoms trapped in a lattice with $V_0 = 10E_R$. Note that for $r > 0$ the states begin to resemble split states of higher and higher orders n 133
- 7.2 Value of the projection integral in Eq. (7.1) versus the state $|r\rangle$ corresponding to the band r for splitting order $n = 1$ (red circles), $n = 2$ (black diamonds), and $n = 3$ (blue squares) for (a) $\theta = 0$ and (b) $\theta = \pi$. As the band number increases, the Bloch wavefunctions look more and more like the split states, with the relative phase θ between the $\pm 2n\hbar k_L$ momentum states equal to $\theta = 0$ (π) for even (odd) band numbers. Alternating bands have relative phases θ of 0 and π , depending on the band parity. 134
- 7.3 Band-to-band transition frequencies (in units of $\omega_R = E_R/\hbar$) versus lattice depth V_0 (in units of E_R) for the transition from band $r = 0$ to $r' = 1$ (red, solid), 2 (blue, dashed), 3 (gray, dotted), 4 (magenta, dash-dotted), and 5 (black). The band 6 transitions are almost degenerate with the band 5 transitions for the entire range of depths considered here, so the $r' = 6$ curve would completely overlap the $r' = 5$ curve. 134
- 7.4 The band transition frequencies ω_{01} (red, solid) and $\omega_{02}/2$ (blue, dashed) as a function of the lattice depth. 139
- 7.5 Results of shaking simulations for a lattice depth of $V_0 = 10E_R$. The percent error relative to the split state (red, see Eq. (3.4)), projection of the current state $|\psi\rangle$ onto the split state with $\theta = 0$ (blue, dashed, see Eq. (7.1)) and $\theta = \pi$ (black, dot-dashed), plotted versus shaking time for a shaking frequency of (a) ω_{01} and (b) $\omega_{02}/2$ and a shaking amplitude of $\alpha = 0.3$ rad. 140
- 7.6 (a) The matrix elements $|M_c|^2$ (red) and $|M_s|^2$ (black) for band-to-band transitions from bands $r = 0$ to r' plotted versus final band r' . (b-c) Plots of the matrix elements (b) $|M_c|^2$ and (c) $|M_s|^2$ for band-to-band transitions from bands r (x-axis) to r' (y axis). The colorbar on the right gives the magnitude of the transition matrix element. 142

- 7.7 Experimental results showing the percent overlap with the split state versus shaking amplitude A for frequencies of 17 kHz $\approx \omega_{02}/2$ (red), 21.5 kHz (blue), and 22 kHz (black), where $\omega_{01} \approx 21.7$ kHz for a lattice depth $V_0 \approx 15.3E_R$ (see Fig 7.4). The points show the mean of three separate experimental runs with the same shaking function and the error bars show the standard deviation of these points. 144
- 7.8 The percent error relative to the (a) third- ($n = 3$) and (b) fourth-order ($n = 4$) split state (red, see Eq. (3.4)), projection of the current state $|\Phi(t)\rangle$ onto the split state with $\theta = 0$ (blue, dashed, see Eq. (7.1)) and $\theta = \pi$ (black, dot-dashed), plotted versus shaking time for a shaking amplitude $\alpha = 1$ rad. In each case, there are two counterpropagating lattices moving at velocities $v = \pm 2n\hbar k_L/m$ 146
- 7.9 (a) Percent error vs. splitting order for $n = 1, 2$ and 3 for the five simulations considered in this work. The frequency ranges considered here are: band transition frequencies (cyan diamonds), half-band transitions (black stars), both band and half-band transitions (blue squares), frequencies with non-negligible matrix element overlap (magenta triangles), and all frequencies in the band (red dots). (b) The best optimized shaking functions for the select frequency case corresponding to the magenta triangles in (a) for $n = 1$ (black, solid), $n = 2$ (blue, dot-dashed), and $n = 3$ (red, dashed). 148
- 7.10 Phase θ between the split arms (blue asterisks, left axis) and error (red dots, right axis) for (a) $n = 1$, (b) $n = 2$, and (c) $n = 3$. The best of ten runs for each of the five simulation classes is shown. The classes are labeled on the x-axis as follows: band transition frequencies (b), half-band transitions (hb), both band and half-band transitions (b+hb), select frequencies with non-negligible matrix element overlap (s), and all frequencies in the band (all) 150

7.11	Results of genetic optimization showing the mean (open markers) and best (closed markers) error for splitting with (a) $n = 1$, (b) $n = 2$ and (c) $n = 3$ versus number of iterations of the optimization algorithm. The black points indicate optimizations with all frequencies in the band from DC to 121 kHz, and the red points indicate optimizations using only the truncated frequency space corresponding to the select band-to-band transitions indicated in Table 7.1. The error improvement is much faster with the truncated subspace.	151
8.1	Location below the chip of the atom chip trap as the pre-lattice waveguide trap is ramped off over an arbitrary ramp time (on the x -axis). As the chip trap is ramped off, the trap location moves rapidly towards the chip.	155
8.2	A cartoon of an atom-based Sagnac gyroscope in a two-dimensional lattice. The atoms (depicted as red and green circles) enclose an area as they undergo the conventional Michelson atom interferometry sequence in one dimension while moving out and back in the orthogonal dimension.	159

Chapter 1

Introduction

1.1 Motivation

Since Michelson and Morley’s experiments that refuted the existence of the luminiferous aether, light-based interferometry has been a useful tool for making scientific measurements. Technology based on the interference of light is able to measure the characteristics of stellar and other astronomical systems [79], the material flatness and deformation [80], rotational motion [11, 37, 85], and more recently, gravitational waves [1, 2, 3].

In the realm of inertial sensing, light-based Sagnac interferometry currently provides the rotation sensing capabilities in commercial inertial measurement units like Northrop Grumman’s LN-200S.¹ An optical Sagnac interferometer works by splitting light into two paths, then recombining them after they traverse an enclosed area \vec{A} (Fig. 1.1). If the system is rotating at a rate $\vec{\Omega}$, then the phase shift $\Delta\Phi$ between the two arms can be calculated to be [103]

$$\Delta\Phi_{\text{light}} = \frac{8\pi\vec{\Omega} \cdot \vec{A}}{\lambda c} \quad (1.1)$$

where for light, c represents the speed of light and λ its wavelength.

This thesis, however, is not about light. Instead, we work with ultracold atoms, or “matter waves,” specifically in the context of using atom-based interferometers as inertial sensors of acceleration and rotation. The utility of atoms as inertial sensors is quite easily seen if we take Eq. (1.1) and substitute matter wave properties for optical properties. That is, the wavelength of light

¹ <http://www.northropgrumman.com/Capabilities/LN200FOG>

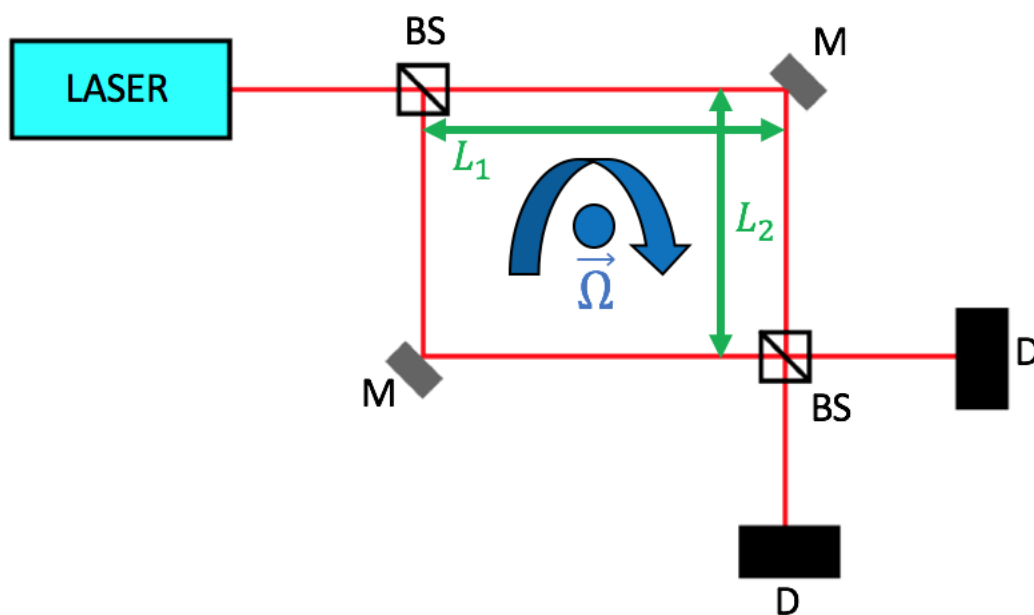


Figure 1.1: A diagram of an optical Mach-Zehnder-style Sagnac interferometer. The light from a laser is split and recombined with two beamsplitters (BS). Mirrors (M) direct the split light around an enclosed area $|\vec{A}| = L_1 L_2$. If the system is rotating with rate $\vec{\Omega}$ (about an axis directed out of the page and parallel to the area vector \vec{A}), then there will be a phase shift $\Delta\Phi$ between the two recombined beams given by Eq. (1.1). This rotation shift can be detected by counting the fringes that pass one of the two detectors (D) at the two output ports of the recombining beamsplitter.

becomes the deBroglie wavelength $\lambda_{DB} = h/p = 2\pi\hbar/mv$ for an atom of mass m moving at a speed v (such that the atom's momentum is $p = mv$).² This allows us to write

$$\Delta\Phi_{\text{atom}} = \frac{4m\vec{\Omega} \cdot \vec{A}}{\hbar}. \quad (1.2)$$

Assuming that $\vec{\Omega}$ and \vec{A} are the same for both the light- and matter-based interferometers, this results in a ratio of the measured phase shifts $\Delta\Phi$ of

$$\frac{\Delta\Phi_{\text{atom}}}{\Delta\Phi_{\text{light}}} = \frac{\lambda mc}{\hbar}. \quad (1.3)$$

Assuming we work with 852 nm light and ⁸⁷Rb atoms, this ratio is on the order of 10^{10} . Note that if the atoms are moving at a constant velocity v around a circular enclosed area $A = \pi r^2$, then the phase shift scales as the square of the interrogation time T_I .

Regardless of the type of wave used to perform interferometry, there is generally a pre-set list of “ingredients” that are used to perform interferometry. We will revisit these ingredients as we move forward in this thesis (see Chapter 4 and beyond), but it is useful to make note of them here. Given a source of waves, most interferometers implement a means of splitting the waves, allowing the split wavepackets to propagate, reflect, reverse-propagate, then finally recombine once the wavepackets have overlapped. All of these operations can be seen in Figs. 1.1 and 4.1, and they are present in the majority of the interferometry works cited in this thesis. These are the ingredients with which we build the interferometer presented in this work.

There are, of course, caveats to the number quoted in the ratio of Eq. (1.3). For example, it is relatively simple to put light into a fiber and allow it to propagate over kilometers, thus greatly increasing the effective area of the device. Getting atoms to do this is much more difficult. Furthermore, photons are rather cheap and easy to produce by the quadrillions/second,³ but atoms are less readily plentiful; a typical BEC produced in our lab contains a few 10^4 atoms. This lower flux contributes to higher shot noise. Shot noise in interferometers typically goes as $1/\sqrt{N}$ for N atoms or photons measured per shot, unless one is using squeezed states to lower this noise

² In this thesis we will use the reduced Planck's constant \hbar preferentially over the non-reduced version $h = 2\pi\hbar$. It's just how we were raised.

³ Assuming a 1 mW laser at 852 nm.

below the standard quantum limit [35, 51]. Continuous measurements are also easier to achieve in laser-based systems, whereas atom-based systems are typically plagued with dead-time issues.

On the flip side, there is no practical laser-based accelerometer, as most commercial systems use micro-electrical-mechanical-system (or MEMS) based devices for accelerometry. Atom-based accelerometers are relatively straightforward to implement, and their sensitivity, like the Sagnac gyroscope sensitivity, is proportional to T_I^2 [74].

This simple back-of-the envelope calculation is enough to motivate the pursuit of atom interferometry, and since the first interferometers were demonstrated in 1991 [21, 58, 59, 97] great advances have been made in the field. This has given rise to some very interesting physics and the development of incredible technology. What follows in the next paragraphs is not meant to be an exhaustive review of the literature on atom interferometry. Rather, we want to highlight some interesting papers that demonstrate some of the more fun aspects of atom interferometry. We have most certainly neglected some of the most interesting developments, but if we included everything, this thesis would double in size.⁴ Recent review papers on the subject of atom-based sensing can be found in Refs. [8, 26, 61, 78].

Typically, the atom-based beamsplitting operation occurs in momentum space using nanofabricated gratings [59, 113], two-photon Raman transitions [75], and/or light pulses in the Bragg [116] or Raman-Nath [117] regime. The Kasevich group at Stanford was able to use Raman and Bragg pulses to incorporate momentum splitting of $102\hbar k_L$ [116] and more recently demonstrated coherence between two split wavepackets that had obtained half-meter spatial separation [63]. Work has been done to optimize the light-pulse sequences used in these systems to increase pulse efficiency [10, 109]. Another method for achieving higher momentum separations (and faster sensitivity scaling of T_I^3) between two wavepackets is using Bloch oscillations in an optical lattice to accelerate one wavepacket relative to the other [22, 74, 91, 123].

To reduce the dead time between shots, an experiment has been done that relies on simultaneous cold atom preparation and detection [33]. Atom-based measurements have been made of

⁴ To my committee: you're welcome.

rotation rates via shifts in the so-called “tune-out” wavelength (where the atomic vector polarizability changes sign) [113]. Six-axis inertial sensors have been demonstrated using light-pulse atom interferometry [20], and atom interferometry has been done with Bose-Einstein condensates in microgravity [84].

Most of the interferometers mentioned above are free-space atom interferometers, where the atoms propagate freely between pulses. While such interferometers have been used to make some of the most sensitive measurements to date, e.g. in 10 m high atom fountains [29, 63], they have the disadvantage that the sensitivity of the device is tied to its size, even though compact devices have been demonstrated [110]. Trapped-atom interferometers (as in Ref. [117]) can get around this limitation but they are sensitive to other deleterious effects like phase diffusion [55]. A trapped atom interferometer in an optical cavity has been implemented by Holger Müller’s group in Berkeley [49]. Our unique approach to atom interferometry is a trapped-atom scheme where we interrogate atoms trapped in an optical lattice via phase modulation (shaking) of the optical lattice. This system, which is amenable to compactification and has some interesting and unique advantages, is described in more detail in the next section (and in the rest of this work).

1.2 The story of the shaken lattice interferometer

As stated in the last section, our approach to interferometry was to use atoms trapped in a phase-modulated (shaken) optical lattice. We wished to implement transformations of the atoms’ momentum state in order to perform an interferometry sequence like those described above. For example, an atom beamsplitter would split the atom population equally into momentum states with equal magnitude and opposite sign.

When we were first exploring this problem we simulated the effects of sinusoidal shaking on atoms trapped in the ground Bloch state of the lattice (see Sec. 2.3.1) using the time-dependent Schrödinger equation. We found that as we shook the lattice, the atoms’ momentum state populations would change only by quanta of $2\hbar k_L$, where $k_L = 2\pi/\lambda_L$ was the wavenumber of the lattice light. However, an analytical solution that found the shaking function for a desired momentum

state transformation eluded us.

As a result, we decided to think about the shaken lattice problem as a boundary value problem. That is, we had an initial state, a desired final state, and some time T that we would shake the lattice. The unknown piece was the shaking function $\phi(t)$ that would transform the initial state into the desired final state. To figure this out, we built on the pioneering work in Ref. [92] and implemented a set of learning and optimal control algorithms that would optimize the shaking function to carry out the desired transformations.

Our optimization simulations were successful, and we soon had a set of state-to-state transformations that described a simulated one-dimensional interferometer. Furthermore, we found that when we added a simulated acceleration signal, the atoms' response varied based on the magnitude and direction of the signal. We not only had a trapped-atom interferometer but one that could determine the direction of an applied signal; the symmetry of the system is broken by the shaking of the lattice. We extended this to optimize our interferometer to be sensitive to perturbations on a bias signal and/or AC signals with a given frequency. Finally, we were able to determine that the sensitivity of our theoretical interferometer scaled as the square of the interrogation time (like the light-pulse interferometers described above).

The challenge then was to implement this system experimentally. We built a 1D lattice around an existing experiment that could produce Bose-Einstein condensates of ^{87}Rb on an atom chip and started modulating the lattice. Every successful graduate student can pinpoint a time (or three) when they realized that the system that they worked so hard on for so long would work. Here, that time occurred shortly before Christmas in 2015 when we performed the first shaking experiments. When the results showed distinctly separate clouds of diffracted atoms separated by $2\hbar k_L$, we knew that the system would work.

The next years were spent optimizing the interferometer in the same way that we theorized and calibrating its response to an applied signal. We were able to show that the interferometer sensitivity scaling was consistent with the T_I^2 scaling that was theoretically predicted by our sim-

ulations.⁵ We were able to show measurements of perturbations on a bias acceleration. At this point we started writing this thesis.

Interspersed throughout the experimental efforts were additional theoretical efforts to understand the analytics of the shaken lattice interferometer. Shortly after the experiments presented here were completed we were able to put together a detailed theoretical analysis of the shaken lattice interferometer; these results are presented in a later part of this work.

The overarching goal of this dissertation is to provide a (reasonably) self-contained document that describes shaken lattice interferometry and our efforts to realize a shaken lattice interferometer. The focus, however, is on the next student to work on this project. The idea is this: when armed with this dissertation and the base level of intelligence and diligence that one requires to succeed in graduate school, the next student should be able to pick up where we have left off without too much issue.⁶ Potential pathways for future students are discussed at the end of this thesis.

1.3 Outline of this dissertation

The rest of this dissertation is organized as follows: Chapter 2 deals with the theoretical aspects of the system, outlining the physics of dipole traps, optical lattices, and shaken lattices. Chapter 3 discusses the various learning and optimal control algorithms used to obtain the computational and experimental results presented in this work. Chapter 4 describes the shaken lattice interferometer in detail and presents the simulation results. Chapter 5 details the experimental setup, and Chapter 6 summarizes experimental results. Chapter 7 outlines some advanced analytics that allow us to better understand why the shaking works. This chapter also describes methods to improve the optimization of atom beamsplitting and reduce deleterious effects that cause atom heating and decoherence. Chapter 8 provides some insight into the future as well as a hypothetical plan for future work.

⁵ This data was taken around midnight on a normal weekday in September 2017, and it marked the second significant point where we realized that we may one day graduate.

⁶ Just don't break the vacuum cell...

Chapter 2

The physics of ultracold atoms in optical lattices

This chapter will lay out the basic theoretical physics background necessary to understand the rest of the thesis, focusing mainly on the pieces required to understand the shaken lattice. Therefore, we will outline the physics of dipole trapping and optical lattices in some detail, but we will largely gloss over the theoretical aspects concerning the production of cold and ultracold atoms. The physics of shaking will be briefly treated but an in-depth treatment will be relegated to Chapters 4 and 7.

2.1 Cold and ultracold atoms, briefly

Our lab focuses on the application of atomic and optical physics techniques developed here at JILA and around the world, so we tend towards using well-developed techniques with the simplest atoms to explore new applications. Thus, in this work we use the relatively simple (but useful!) rubidium-87 atom. This isotope of rubidium was among the first condensed [6] (with sodium being the other condensed atom [28]). More specifically, we follow the well-established method of cooling and trapping using the $F = 2 \rightarrow F' = 3$ D2 transition of ^{87}Rb , and we magnetically trap and cool the atoms to degeneracy in the $|F, m_F\rangle = |2, 2\rangle$ ground state. (Useful numbers and energy level diagrams can be found in Ref. [106], a useful reference for any experimental physicist working with ^{87}Rb to have in their toolbox.)

As usual we start by trapping cold atoms in a magneto-optical trap (MOT) before proceeding to a magnetic trap and cooling to condensation using forced RF evaporation. Magneto-optical

trapping of atoms is a well-established technique first demonstrated in 1987 [94]. An excellent general reference for laser cooling and trapping is the book by Metcalf and van der Straten [77].

Magneto-optical trapping can only cool atoms to the Doppler limit

$$T_D = \frac{\hbar\Gamma}{2k_B} \quad (2.1)$$

where Γ is the lifetime of the $F' = 3$ state, and k_B is Boltzmann's constant. For ^{87}Rb T_D is on the order of 100 μK whereas we make BECs with temperatures around hundreds of nanokelvin. Using sub-Doppler cooling techniques (namely polarization gradient cooling, or PGC [27]) the atoms' temperature may be lowered to about 4 μK . We find the original paper [27] to be rather dense from a basic experimental point of view. A simpler explanation¹ is given by S. Du in his thesis [32]; this has proven sufficient for our practical understanding.

Once the atoms (in the $F = 2$ hyperfine sublevel of the $5^2S_{1/2}$ ground state of ^{87}Rb) are cooled to the sub-Doppler limit set by PGC, they must be magnetically trapped and transported to the atom chip for condensation. However, the magnetic sublevels of the $F = 2$ state are degenerate in the absence of a magnetic field, and thus the atoms populate the sublevels equally when a field is applied. Earnshaw's theorem states that we cannot make a local magnetic field maximum with static magnets. Therefore, we must trap the atoms in magnetic field minima. In order to trap the atoms in a static magnetic field, we apply a magnetic field to break the hyperfine states' degeneracy and optically pump the atoms into the $|F, m_F\rangle = |2, 2\rangle$ state using circularly-polarized light. After optical pumping we can trap the atoms in a magnetic trap produced by external coils and transport them to our atom chip where they are cooled to degeneracy. We will defer practical details of the chip trap to Chapter 5. For the interested reader curious about the details of BEC, the book by Pethick and Smith [90] is a good starting point to learn the theory. Further resources focusing on experimental implementation may be found in Refs. [60, 67].

¹ Read: for experimentalists.

2.2 Optical dipole traps

The physics behind optical dipole traps is discussed thoroughly in Ref. [46] and only the relevant details will be noted here. Optical dipole traps are formed when atoms interact with light that is detuned from an atomic resonance. When the atom interacts with an electric field \vec{E} , a dipole moment $\vec{p} = \alpha\vec{E}$ is induced in the atom. Here α is the complex polarizability of the atom. The potential takes the form

$$U_{\text{dip}} = -\frac{1}{2\epsilon_0 c} \text{Re}(a)I \quad (2.2)$$

where $\epsilon_0 = 8.85 \times 10^{-12}$ F/m is the electric permittivity in vacuum and $c = 2.998 \times 10^8$ m/s. Thus, the force that arises from this potential is proportional to the gradient of the intensity, $\vec{F} \propto \nabla I(\vec{r})$. For light that is red-detuned from resonance the sign of α is such that the potential is attractive, and blue-detuned light creates repulsive potentials. In this work we use only red-detuned light for our optical traps.

If we trap an alkali metal atom in a linearly polarized Gaussian beam, we can represent the dipole potential as [46]

$$U_{\text{dip}}(\vec{r}) = \frac{\pi c^2 \Gamma}{2\omega_0^3} \left(\frac{2}{\Delta_{2,F}} + \frac{1}{\Delta_{1,F}} \right) I(\vec{r}) \quad (2.3)$$

where $\Gamma = \Gamma_{\text{D1}} + \Gamma_{\text{D2}}$ is the sum of the natural linewidths of the Rubidium D1 and D2 lines, $\omega_0 = (\omega_{\text{D1}} + \omega_{\text{D2}})/2$ is the average of the two lines' transition frequencies, and $\Delta_{i,F}$ is the detuning of the trapping laser from each line.² The intensity profile of a Gaussian beam with beam power P is given by the expression

$$I(x, r) = \frac{2P}{\pi w(x)^2} e^{-r^2/w(x)^2}. \quad (2.4)$$

In Eq. (2.4) the beam propagates in the x -direction, and the Gaussian radial confinement in the other two dimensions is represented by the radial variable $r = \sqrt{y^2 + z^2}$. The beam waist along x is given by $w(x) = w_0 \sqrt{1 + \left(\frac{x}{x_{\text{R}}}\right)^2}$ where the Rayleigh range $x_{\text{R}} = \pi w_0^2/\lambda_{\text{L}}$ and w_0 is the beam waist at the focus. In practice, we use the mathematics presented here when designing an

² While we trap and cool using the D2 line of Rubidium-87, we need to take both the D1 and D2 lines into account when analyzing dipole trapping. Details on the frequencies and linewidths of each line are found in Ref. [106].

experiment. For example, if we design a system to trap atoms in a Gaussian beam with wavelength λ_L and waist w_0 at the beam focus, Eqs. (2.3) and (2.4) can be used to estimate the laser powers required to trap atoms and hold them against gravity (which is defined to be along the z -direction).

We must also take photon scattering into account when designing experiments. Even though the above expression (2.3) is derived in a far-detuned limit, there is still a nonzero probability for the atoms to scatter photons. This scattering rate Γ_{sc} often sets the fundamental limit for the lifetime of atoms in optical trap. For alkali metal atoms we can write

$$\Gamma_{\text{sc}}(\vec{r}) = \frac{\pi c^2 \Gamma^2}{2\hbar\omega_0^3} \left(\frac{2}{\Delta_{2,F}^2} + \frac{1}{\Delta_{1,F}^2} \right) I(\vec{r}). \quad (2.5)$$

If we approximate $\Delta_{D1} \approx \Delta_{D2} \approx \Delta$ (that is, the dipole trapping is very far detuned from both resonances), for a given detuning Δ dipole potential depth scales with I/Δ . The scattering rate depends on I/Δ^2 . Therefore, detuning the laser farther and farther from resonance decreases the photon scattering rate, but more power is required to trap the atoms.

In our experiments we use $\lambda_L = 852$ nm out of convenience, as we had an available laser from an old cesium experiment. This is far enough detuned from the rubidium-87 resonance at $\lambda_0 = 780$ nm that we can trap atoms long enough to interrogate them but close enough that we have enough laser power available to provide a sufficiently deep trap.

2.3 Lattices

If we have sufficient light for a dipole trap we can retro-reflect the light back onto itself and create a standing wave with node spacing $d = \lambda_L/2$ and wavevector $k_L = 2\pi/\lambda_L$. This standing wave is known as an optical lattice. When discussing the theoretical basis behind optical lattices we will largely follow what is done in Ref. [42], slightly modifying the notation to avoid confusion later in this thesis.

There are generally two regimes considered when one analyzes the physics of atoms in optical lattices. The first of these is the Bloch model which allows us to derive the band structure of the periodic system. If we restrict ourselves to the lowest band we can also describe the physics of

the system in terms of the Bose-Hubbard Hamiltonian. In this thesis we work primarily in the Bloch basis and describe the Bose-Hubbard basis at the end of this section for completeness and to facilitate a brief discussion of localized vs. delocalized behavior.

In the experiments presented in this thesis, we work with a one-dimensional Gaussian beam retro-reflected onto itself. The radial confinement of the atoms is thus provided by the radial Gaussian potential, and the transverse confinement is due to the lattice potential. We work with lattices shallow enough for the atoms' to be "delocalized" in the lattice potential. That is, the atoms' position wavefunction extends over many lattice sites so the physics is best described in the Bloch basis.

2.3.1 The Bloch basis

If we ignore atom-atom interactions and the Gaussian extent of the beam along x ,³ we can write the Bloch Hamiltonian for an optical lattice in one dimension as

$$H = -\frac{\hat{p}^2}{2m} - \frac{V_0}{2} \cos(2k_L x) \quad (2.6)$$

where \hat{p} is the momentum operator and we usually quantify the prefactor V_0 in terms of the recoil energy of the photons in the lattice, $V_0 = \alpha E_R$ where the recoil energy is $E_R = \hbar^2 k_L^2 / 2m$. This can be converted into a recoil frequency in the usual way: $\omega_R = E_R / \hbar = 2\pi \times 3.169$ kHz (in our system). The recoil frequency, lattice wavelength, and recoil energy set the natural time, distance, and energy scales for the system.

As with any Hamiltonian containing a periodic potential, we can apply Bloch's theorem to solve the system. That is, we assume that the functions $\psi(x)$ that solve the time-independent Schrödinger equation (TISE, $H\psi = E\psi$) can be written in the form

$$\psi_q^{(r)}(x) = e^{iqx} u_q^{(r)}(x). \quad (2.7)$$

The wavefunctions $\psi_q^{(r)}$ can be written as the product of a plane wave and a function $u_q^{(r)}(x)$ that we define to have the same periodicity as the potential. We will see later that the index q labels

³ This is a reasonable approximation, since our lattice beams have a Rayleigh range of x_R on the order of a few millimeters, which can be considered roughly infinite relative to the other length scales in the experiment.

the quasimomentum of an atom in Bloch band (or energy level) r . If we insert this into the TISE, we obtain an analogue to the TISE for the $u_q^{(r)}(x)$ as

$$H_B u_q^{(r)}(x) = E_q^{(r)} u_q^{(r)}(x) \quad (2.8)$$

where

$$H_B = \frac{1}{2m}(\hat{p} + q)^2 - \frac{V_0}{2} \cos(2k_L x). \quad (2.9)$$

Here we see the origin of the term ‘‘quasimomentum,’’ as the q factor is lumped in with the kinetic energy term in the expression for H_B . If we decompose both the $u_q^{(r)}(x)$ and the lattice potential in terms of a Fourier series, we obtain

$$u_q^{(r)}(x) = \sum_l c_{l,q} e^{2ilk_L x} \quad (2.10)$$

and

$$V(x) = \sum_j V_j e^{2ijk_L x} \quad (2.11)$$

where both l and j are indices. For the cosinusoidal lattice potential, we have $V_{\pm 1} = -1/4$ with all other terms being zero.⁴ This allows us to write Eq. (2.9) as

$$\frac{(2\hbar k_L l + q)^2}{2m} c_{l,q}^{(r)}(x) e^{2ilk_L x} - \frac{1}{4} V_0 \left(e^{2ik_L x} + e^{-2ik_L x} \right) c_{l,q}^{(r)} e^{2ilk_L x} = E_q^{(r)} c_{l,q}^{(r)}. \quad (2.12)$$

In matrix form, we write Eq. (2.12) as

$$\sum_{l,l'} H_{B,l,l'} c_{l,q}^{(r)} = E_q^{(r)} c_{l,q}^{(r)} \quad (2.13)$$

where H_B is the tridiagonal matrix with entries

$$H_{B,l,l'} = (2l + q/\hbar k_L)^2 E_R \delta_{l,l'} - (1/4)\alpha E_R \delta_{l,l'-1} - (1/4)\alpha E_R \delta_{l,l'+1}. \quad (2.14)$$

In Eq. 2.14, δ_{ij} represents the Kronecker delta.

We find that there is a eigenvector $c_q^{(r)}$ with energy eigenvalue $E_q^{(r)}$ for each band r and quasimomentum q . The quasimomentum can vary from $-\hbar k_L$ to $+\hbar k_L$ within each band r ; this is

⁴ Note that we can set V_0 to be whatever we want, because it's simply a constant offset.

known as the “first” Brillouin zone, as the band structure repeats with periodicity $2\hbar k_L$ (with the first Brillouin zone centered around $q = 0$).

To calculate this band structure we solve Eq. (2.13) using the matrix in Eq. (2.14) for values of $E_q^{(r)}$ within the first Brillouin zone. The Bloch wavefunction for a given quasimomentum q and band r is found by finding the eigenvectors $c_q^{(r)}$, using Eq. (2.10) to find the corresponding $u_q^{(r)}(x)$, then multiplying by the plane wave factor in Eq. (2.7) to obtain $\psi_q^{(r)}(x)$. To carry out these calculations we must limit l to a finite value. In Ref. [42], truncation at $|l| \leq 5$ is recommended. Because we are not limited in computational power and sometimes consider higher bands than considered in Ref. [42], we generally truncate at $|l| \leq 10$.

The band energies in the first Brillouin zone as a function of the quasimomentum q for $\alpha = 10$ and $\alpha = 20$ are shown in Fig. 2.1, and the wavefunctions corresponding to Bloch states in lowest band $r = 0$ for $q = 0$ and $q = \hbar k_L$ are shown in Fig. 2.2. Bloch wavefunctions for the lowest and first excited band at $q = 0$ and $\alpha = 10$ are shown in Fig. 2.3.

Two important limiting cases should be noted. As the lattice depth $V_0 \rightarrow 0$ the gaps between the bands (the “bandgaps”) vanish. At this point the atoms’ energy structure becomes continuous and resembles that of a free particle, as expected given the absence of a potential for $\alpha = 0$. In the other limit, as $V_0 \rightarrow \infty$, the atoms become more and more confined to a single lattice site. In this limit the potential looks harmonic, and as expected the bands become flatter and more evenly spaced. In Fig. 2.1, one can see that the lower bands are flatter and lattice bandgap is increased for a depth of $20E_R$ (b) relative to the $10E_R$ case (a). In both cases, as one moves to the higher bands, the bandgap decreases and eventually vanishes.

When considering the Bloch states (as shown in Figs. 2.2 and 2.3), some important features can be noted. First of all, at higher lattice depths the bands are more confined to the “classically allowed” regions. This makes intuitive sense, as the deeper the lattice depth, the more likely that the atom is localized to a given site. (Generally, in this regime the physics is described in terms of the Bose-Hubbard Hamiltonian discussed briefly in the next section.) For a given band the parity of the Bloch state at $q = 0$ is opposite that at $q = \hbar k_L$. Similarly, at a given quasimomentum,




Figure 2.1: The ground (green) and first four excited Bloch states (red, blue, magenta, and cyan, respectively) for lattice depths of (a) $\alpha = 10$ and (b) $\alpha = 20$. The lattice depth is shown via a black, dotted line.

q_0 , the parity of the state corresponding to band r will be opposite that of band $r + 1$. This will be discussed further in Chapter 7. Furthermore, because the position-space Bloch wavefunctions are periodic, the momentum space representation is quantized in units of $2\hbar k_L$ as shown in Fig. 2.3(b) [50]. The quantization of this momentum is a key result that we will use extensively when we design and optimize our interferometer.

2.3.2 The Bose-Hubbard Hamiltonian, briefly

If we limit the system to a single Bloch band the atom dynamics are described by the Bose-Hubbard (B-H) Hamiltonian. In this regime we assume atoms are localized to a single site, and the relevant quantities are the atom-atom interaction and the atom tunneling parameters, usually denoted U and J , respectively.⁵ If we ignore the effects of external confinement due to the Gaussian beam, we can write the 1-D B-H Hamiltonian as

$$H_{\text{B-H}} = -J \sum_{i,j} \hat{a}_i^\dagger \hat{a}_j + \sum_i \frac{1}{2} U \hat{n}_i (\hat{n}_i - 1). \quad (2.15)$$

In Eq. (2.15) the \hat{a}_i^\dagger and \hat{a}_i are the creation and annihilation operators for a boson on lattice site i , and $\hat{n}_i = \hat{a}_i^\dagger \hat{a}_i$ is the number of bosons on site i . We will not be using the B-H

⁵ One can also use the Bloch bands to build a basis of states localized to lattice sites. These Wannier functions will not be considered here, but they are presented and used to derive the B-H Hamiltonian and associated factors U and J in Ref. [42].




Figure 2.2: The ground ($r = 0$) Bloch state for (a) $\alpha = 10$ and (b) $\alpha = 20$ at quasimomentum $q = 0$ (red) and $q = \hbar k_L$ (green). The plots show the real part of the wavefunction and are normalized such that the maximum amplitude of the state is one; the imaginary part is zero. For a spatial reference, the function $\cos(2k_L x)$ is plotted (blue), representing the lattice.

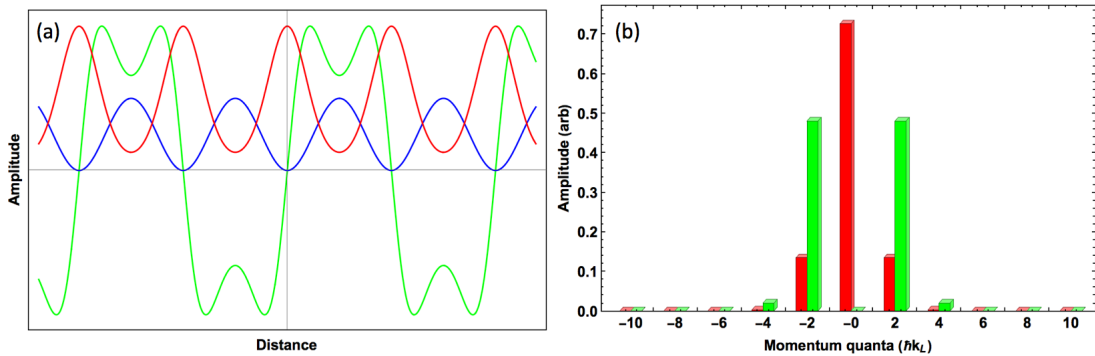


Figure 2.3: The (a) position- and (b) momentum-space representations of the ground (red) and first (green) excited Bloch states of an optical lattice with $\alpha = 10$. As in Fig. 2.2 the plots in (a) are of the real part of the wavefunction and are normalized such that the maximum amplitude of the state is one; the imaginary part of the wavefunctions are zero. For a spatial reference, the function $\cos(2k_L x)$ is plotted (blue), representing the lattice. In plot (b) the relative population in each of the $2n\hbar k_L$ states is shown by the height of the bars. More plots of the Bloch state momentum populations are shown in Fig. 7.1.

Hamiltonian in this work and present it here only for completeness. In particular, work has been done to identify the regime where a BEC in an optical lattice transitions from the delocalized (or superfluid) state to the localized (or Mott-insulating) state as the values of J and U change relative to one another [42, 44].

The superfluid state is characterized by the ability to resolve distinct $2\hbar k_L$ -quantized momentum states in far-field imaging, whereas in the Mott-insulating state the atom localization in position leads to a delocalization in momentum space and the distinct momentum states are indistinguishable. In a 1-D lattice this transition has been observed to occur at potential depths $V_0 \approx 31E_R$ [101]. This is far from our operating regime where $\alpha = 10 - 20$, so our atoms remain delocalized in our optical lattice, although we do have to take into account the effects of atom localization in Chapters 6 and 7. Furthermore, when we make transitions between different momentum states, we involve multiple Bloch bands. Therefore, the single-band Bose-Hubbard Hamiltonian is insufficient for our analysis, and we will use the Bloch basis to describe the physics of the atoms in the shaken lattice.

2.4 Atoms in a shaken optical lattice

In this section we will detail some of the physics of the shaken lattice. More detailed analytical treatment is deferred to Chapter 7, and only the basics needed to understand the shaken lattice interferometer are presented here. We consider atoms trapped in a shallow optical lattice potential with an arbitrary phase-modulation (shaking) function $\phi(t)$ described by the (now time-dependent) Hamiltonian

$$H(x, t) = -\frac{\hat{p}^2}{2m} - \frac{V_0}{2} \cos(2k_L x + \phi(t)). \quad (2.16)$$

The evolution of a state $\psi(x, t)$ is given by the time-dependent Schrödinger equation (TDSE). The TDSE is written in position space as

$$i\hbar \frac{\partial \psi}{\partial t} = H(x, t) \psi(x, t). \quad (2.17)$$

In shaken lattice interferometry, we are working with atoms trapped in a shallow lattice, so

we work in the delocalized Bloch basis. Due to the fact that the Bloch states are quantized in momentum space, it is convenient to examine how this Hamiltonian behaves in this space. We define the momentum-space wavefunction $\tilde{\psi}(k, t)$ as

$$\tilde{\psi}(k, t) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} dx \psi(x, t) e^{ikx}. \quad (2.18)$$

Inserting Eq. (2.18) into the TISE, we obtain [92]

$$i\hbar \frac{\partial \tilde{\psi}}{\partial t} = \frac{\hbar^2 k^2}{2m} \tilde{\psi} - \frac{V_0}{2} \left[\tilde{\psi}(k - 2k_L) e^{-i\phi(t)} + \tilde{\psi}(k + 2k_L) e^{i\phi(t)} \right]. \quad (2.19)$$

This equation tells us one very important thing: the shaking couples states of different momentum, but these states must be separated by $2\hbar k_L$. This is verified by our simulations in Chapters 4 and 7. Thus, we will make only direct transitions between states of the same quasimomentum q . This is important because it allows us to restrict ourselves to a single value of $q = q_0$ and work in the Bloch basis defined by the lattice depth $V_0 = \alpha E_R$ and this chosen quasimomentum q_0 .

For the simulation work presented in this thesis we assume that the atoms are loaded into the lattice at rest in the lowest Bloch band ($r = 0$). This allows us to set $q_0 = 0$, corresponding to the energy minimum at the center of the first Brillouin zone (Fig. 2.1). The atoms have a momentum state population similar to that shown in Fig. 2.3(b), depending on the value of α . As expected from the limiting behaviors discussed in Sec. 2.3 as α decreases, more atoms populate the $0\hbar k_L$ state. Conversely, more atoms populate the higher-order momentum states as α increases. As discussed in Sec. 2.3, we will work in the regime where $\alpha \approx 10 - 20$.

The theory discussed here is enough to begin exploring the physics of the shaken lattice. However, it is difficult to solve the TDSE with the Hamiltonian in Eq. (2.16) analytically except in the simplest of cases.⁶ Therefore, as described in Chapter 1, in order to control the atoms' behavior in both simulation and experiment, we used learning and optimal control algorithms to circumvent these analytical difficulties. To facilitate understanding of the underlying mechanisms governing our

⁶ These cases will be discussed in Chapter 7.

simulations, before discussing simulation results in Chapter 4 and experimental results in Chapter 5, we will outline the algorithms used for optimization in Chapter 3.

Chapter 3

Simulation and optimization methods used in this thesis

This chapter will discuss the simulation and optimization methods used in this thesis. First, we will briefly describe the symmetric split-step algorithm used to solve the time-dependent Schrödinger equation (TDSE) and discuss errors in the method. This is the method used to perform the simulations presented in this thesis. Then, we will discuss the motivation for using an optimal control or learning algorithm to solve the shaken lattice problem. This will lead into a discussion of the algorithms that were used in this work. We will discuss the genetic algorithm, then turn to two separate optimal control methods: the Krotov method and the CRAB method. While these algorithms are discussed extensively in the literature, we will describe both, with an emphasis on how they were implemented in simulation or experiment. As in the rest of this dissertation the information contained here is presented with a focus on informing the next student to take on this project of what was done in my time in the lab.

3.1 The split-step method

3.1.1 How the split-step method works

When using the split-step method of propagation [40, 65] the TDSE is written in terms of a unitary propagation operator. For a generic Hamiltonian of the form $\hat{H} = \hat{K} + \hat{V}$, the propagator for a state $\psi(x, t)$ takes the form

$$\psi(x, t + \delta t) = \hat{U}\psi(x, t) = e^{-\frac{i\delta t}{\hbar}\hat{H}}\psi(x, t). \quad (3.1)$$

The split-step method is implemented by separating the kinetic energy operator \hat{K} and potential energy operator \hat{V} and propagating them separately. Because the kinetic energy operator is diagonal in momentum space and the potential energy operator is diagonal in position space, the kinetic energy propagation occurs in Fourier (momentum) space and the potential energy propagation occurs in real (position) space. The fast Fourier transform allows for quick computation via this method, and the norm of ψ is preserved due to the fact that the operator \hat{U} is unitary. When using this method the boundary conditions are periodic, making it an ideal scheme for propagation of the TDSE with the periodic Hamiltonian in Eq. (2.16). However, there are issues that arise when we attempt to simulate the addition of an acceleration signal, and these will be addressed in Sec. 4.2.

3.1.2 Errors in the split-step method

Due to the fact that the $[\hat{K}, \hat{V}] \neq 0$, this method incurs errors in phase. For the symmetrized split-step method currently used in SLI computation the propagation for a timestep δt is approximated as [40, 65].

$$\psi(x, t + \delta t) = e^{\frac{i\delta t}{\hbar}\hat{K}} e^{\frac{i\delta t}{\hbar}\hat{V}} e^{\frac{i\delta t}{\hbar}\hat{K}} \psi(x, t). \quad (3.2)$$

If the propagator is approximated as in Eq. (3.2) then the error is on the order of δt^3 .¹ For the shaken lattice the Hamiltonian is given by Eq. (2.16). For N timesteps, where the total propagation time $T = N\delta t$, the phase error increases linearly. In this case, the phase error takes the form

$$E(\delta t) = e^{i\alpha N(\omega_R \delta t)^3}. \quad (3.3)$$

In Eq. (3.3), $\omega_R = \hbar k^2/2m$ is the recoil frequency of the atoms in the lattice and α is the lattice depth in recoils, as defined in Chapter 2.

¹ This may be explicitly derived by splitting the unitary evolution into two separate but identical operators with timestep $\delta t/2$, then applying the Baker-Hausdorff identity. The second-order correction cancels, and the lowest-order error is on the order of δt^3 .

In simulation we scale units such that $\hbar = m = E_R = 1$.² Therefore, $\omega_R = 1$ and the scaled time is such that $T = 1 \approx 50 \mu\text{s}$.³ In simulations we generally set $T = 10$, and the timestep is typically set to $\delta t = 1/1000$ in scaled units so that the total number of steps is $N = 10^4$. Using these numbers and $\alpha = 20$ for a worst-case scenario, the phase error is on the order of 40×10^{-6} radians. The total phase error after a total simulation time $T = 10$ is plotted versus δt in Figure 3.1.




Figure 3.1: The phase error in radians from Eq. (3.3) plotted vs. timestep δt in scaled units. For the typical timestep used in the simulations, $\delta t = 1/1000$, limiting the phase error to tens of microradians.

To verify that this phase error is negligible simple simulations were run wherein the wavefunction begins in the ground Bloch state of a lattice of $10E_R$ and the lattice was shaken sinusoidally with various frequencies. The percent error was calculated from the final momentum state population vector \vec{P}_f obtained after shaking; this final state is dependent on the timestep δt .⁴ The error

² This is done to avoid extremely large and extremely small numbers, both of which cause problems in numerical computation.

³ This value is dependent on the lattice wavelength used, and the numbers given here assume $\lambda_L = 852 \text{ nm}$.

⁴ Here and in (most of) the rest of this work, only amplitudes are considered and phases are ignored due to the fact that they are not accessible experimentally.

is represented as

$$E(\delta t) = \left(1 - \frac{\vec{P}_f(\delta t) \cdot \vec{P}_f(\delta t')}{|\vec{P}_f(\delta t)| |\vec{P}_f(\delta t')|} \right) \times 100\%. \quad (3.4)$$

In equation (3.4) the reference vector $\vec{P}_f(\delta t')$ is taken as the final momentum state population vector resulting from simulation with a timestep $\delta t'$ of 1/51200, corresponding to a phase error of 1.53×10^{-8} radians in Fig. 3.1.⁵ An absolute reference is not obtainable due to the fact that this equation is not analytically solvable, so the finest-timestep final state is taken instead as the most accurate representation of the true final state due to the fact that the phase error is negligibly small in this case. The results are shown in Fig. 3.2. For typical simulation parameters the relative error stays above $E_{lim} = 10^{-3}$ %. Generally, the convergence of the optimization methods give errors between 10^{-3} % and 0.1 %. Therefore, convergence to an error $E < E_{lim}$ is limited by the simulation accuracy.

3.2 Motivating the use of learning algorithms

Solving the Hamiltonian in Eq. (2.16) is not straightforward except in the simplest of cases. These cases will be discussed more in Chapter 7. Beyond single-frequency shaking (which can be analyzed using the Jacobi-Anger expansion or Floquet theory) the analytics become quickly very complex. We have not been able to find analytic solutions to multiple-frequency shaking that are useful for our work. Thus, as stated in Chapter 1 it is useful to formulate the shaken lattice problem as a boundary-value problem. That is, we have an initial state, our desired final state, the shaking time T , and an equation that governs the time evolution of the system. This system is rather straightforward to set up as an optimization problem, and we solved the shaken lattice problem using the optimization methods described in the rest of this chapter.

The idea of using learning algorithms to solve the shaken lattice problem was not our own. Our work on the shaken lattice interferometer builds on the pioneering work done in Pierre Meystre's group back in 2001 [92, 93]. In this work, they simulated the dynamics of a BEC in a shaken

⁵ This is the smallest timestep that would run in a reasonable amount of time on a desktop computer.




Figure 3.2: The error between the final state and the finer-time reference state in Eq. (3.4) plotted vs. timestep δt in scaled units for sinusoidal shaking with various frequencies ω . The case where no shaking is applied is shown in red, which demonstrates the propagation of the initial eigenstate in a stationary lattice. For the typical timestep used in the simulations, $\delta t = 1/1000$, and for all shaking frequencies the percent error is $< E_{lim} = 10^{-3} \%$.

optical lattice and used a genetic algorithm (GA) to control the dynamics, specifically related to the problem of coherent acceleration of a BEC. Our GA (described in detail in Sec. 3.3) is largely inspired by the algorithms described in Refs. [57,92]. While further exploring optimization methods, we simulated and optimized the shaken lattice system using the Krotov method [88,89,104]. The Krotov method is derived in detail and applied to the shaken lattice problem in Sec. 3.4. Finally, we used the Chopped RAndom Basis (CRAB) method [19,31,95] to optimize our experimental shaking. This method has been successfully implemented in experimental systems [41,69,114,115] and is detailed in Sec. 3.5.

3.3 The genetic algorithm

The text and images in this section have been adapted from the work published in Ref. [119]. Details have been added where needed, and other details have been deferred to Chapter 4.

The genetic algorithm was the first method used to optimize lattice shaking, and it remains a useful tool when exploring lattice dynamics. A block diagram of the GA procedure is shown in Fig. 3.3. First, the algorithm produces an initial population of A individuals; each individual corresponds to a different shaking function. This is the first generation, denoted $G = 1$. Each individual α is produced by generating a random vector $\vec{\alpha}(\omega)$ of l Fourier amplitudes at uniformly-spaced frequencies ω_i from DC up to a certain frequency limit set by the user; this sets the bandwidth of the shaking. For the work presented in this thesis we set $A = 20$ and the Fourier amplitudes are randomly chosen from a normal distribution with standard deviation $\sigma = 100$. We pick $l \approx 100$ and limit the bandwidth to about 35 kHz unless otherwise stated. The value of σ is chosen to limit $\phi(t)$ to within approximately $\pm 2\pi$ radians and the bandwidth is generally chosen to limit the search space.⁶ We choose an l that gives good convergence of the GA within a reasonable amount of computation time.⁷

⁶ We have in the past tried to limit the force on the atoms due to shaking via bandwidth control, as discussed in Ref. [119].

⁷ Practically, the limit on the value of l is given by the timestep dt used in the problem and the bandwidth of the shaking. That is, we set dt and the bandwidth, and l is determined based on those parameters. This arises because we use a discrete Fourier transform to go between the temporal and frequency domain.

Once the vector $\vec{\alpha}$ is chosen its discrete Fourier transform is taken to produce a time-varying function. To maintain smooth turn-on and turn-off, this function is multiplied by an envelope function

$$f_{\text{env}}(t) = \sin^2(\pi t/T) \quad (3.5)$$

that goes to zero at its endpoints. This results in a shaking function $\phi_\alpha(t)$ for each individual. The shaking time T is a parameter chosen by the user and, as stated in Sec. 3.1.2 this is usually set to approximately 0.5 ms. We choose the sine envelope shape over e.g, a Gaussian, because it possesses a quicker “turn-on” time, and in practice, the discontinuities at the endpoints do not cause undesired behavior in the shaking function.

Once the initial individuals are generated, the initial state ψ_0 is propagated forward in time by solving the TDSE with the split-step method with $\phi_\alpha(t)$, producing a final state ψ_α . Because experimental measurements can provide momentum state populations via time-of-flight imaging, we represent a simulated result for the Fourier transform of the final wavefunction $\Psi_\alpha(k) = \mathcal{F}\{\psi_\alpha(x)\}$ as a vector $\vec{P}_\alpha(|\Psi_\alpha|^2)$. As in Sec. 3.1.2, this vector has components $P_{\alpha,n}$ representing the relative atom population in each momentum state $2n\hbar k_L$ for $n = [-\infty, \infty]$. A similar vector $\vec{P}_d(|\Psi_d|^2)$ is constructed for the desired final state. In this work n is truncated at $n = [-5, 5]$ since higher-order momentum states are negligibly populated. Due to the presence of the applied potential, only transitions between momentum states separated by $2\hbar k_L$ are allowed (see Eq. (2.19) and surrounding text for more information). As stated in Sec. 2.4, we work with atoms with quasimomentum $q = 0$ and thus couple only states with momentum $2n\hbar k_L$ for integer n .

For a given shaking protocol $\phi_\alpha(t)$, the GA assesses the quality of the final momentum-space population via a fitness function $f(\vec{P}_\alpha)$ which quantifies the difference between the final and desired states. The fitness function used in this work will be discussed in detail in Chapter 4. Once the fitness of the initial population is evaluated the individuals are ranked in terms of their fitness (because we are minimizing the fitness function, lower fitness values are better) and the genetic algorithm uses this ranking to produce the next generation of individuals. In general the procedure

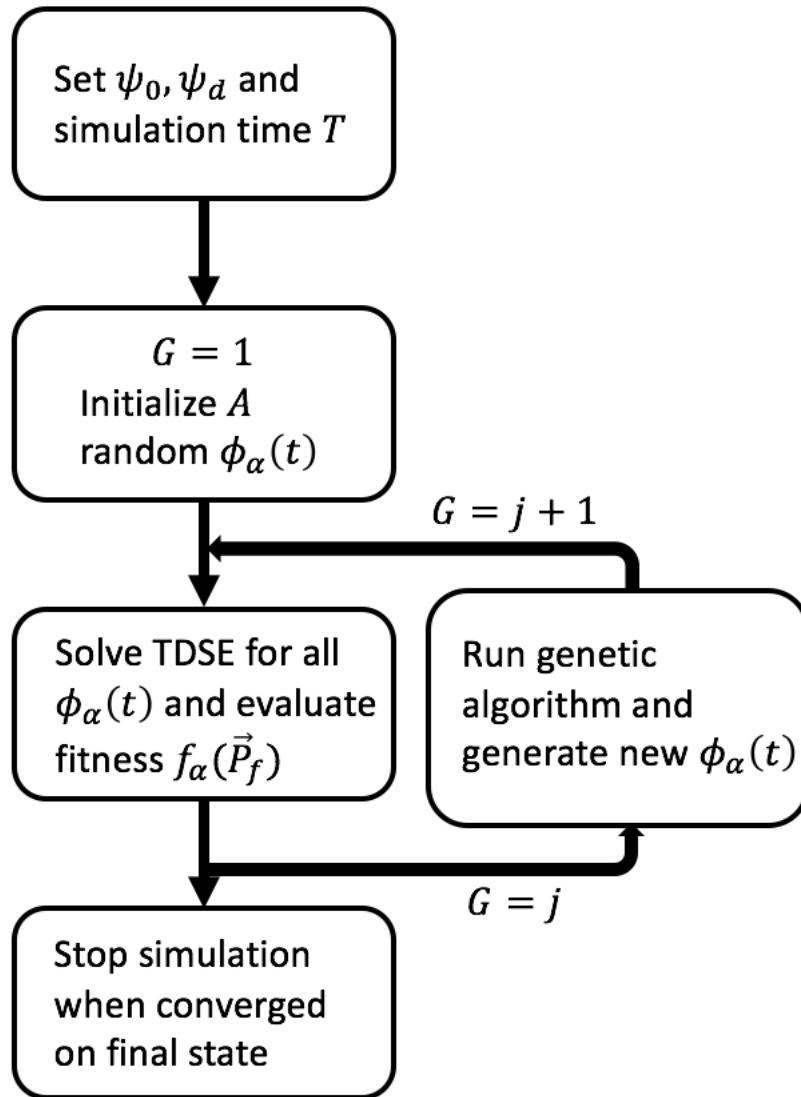


Figure 3.3: A block diagram illustrating the steps taken in the GA. Given the initial and desired states, the first generation $G = 1$ of A individuals is randomly generated with unique shaking functions $\phi_\alpha(t)$. We then solve the TDSE for each individual in the generation. These results are then used to produce the next generation of individuals, denoted $G = 2$. After the j^{th} run of the simulation a generation $G = j + 1$ results from the mixing of the previous generation's individuals (see Fig. 3.4). Once the convergence criterion is met, the GA stops.

is identical from one generation $G = j$ to the next, $G = j + 1$. The A_{live} best individuals are allowed to proceed unaltered to the next generation, a process known as “elitism” that ensures that the best (lowest) fitness value will never increase. The A_{die} worst individuals are deleted entirely. In this work we choose $A_{\text{live}} = 2$ and $A_{\text{die}} = 4$. The remaining $A - A_{\text{live}}$ new individuals are generated by randomly picking “parents” from the surviving population for each new individual.

The new generation is created from the previous generation using methods adapted from Ref. [92]. These methods are discussed here and illustrated in Fig. 3.4. Given two parent vectors of Fourier amplitudes $\vec{\alpha}$ and $\vec{\alpha}'$, one-point crossover picks a random index c and creates two new vectors $\vec{\alpha}_{\text{a}}$ and $\vec{\alpha}_{\text{b}}$ by swapping the values of the parent vectors starting at index c such that $\vec{\alpha}_{\text{a}} = \{\alpha_1, \alpha_2, \dots, \alpha_c, \alpha'_{c+1}, \dots, \alpha'_l\}$ and $\vec{\alpha}_{\text{b}} = \{\alpha'_1, \alpha'_2, \dots, \alpha'_c, \alpha_{c+1}, \dots, \alpha_l\}$. Two-point crossover performs the same swap, but two indices c_1 and $c_2 > c_1$ are randomly chosen. This results in two children: $\vec{\alpha}_{\text{a}} = \{\alpha_1, \alpha_2, \dots, \alpha_{c_1}, \alpha'_{c_1+1}, \dots, \alpha'_{c_2}, \alpha_{c_2+1}, \dots, \alpha_l\}$ and $\vec{\alpha}_{\text{b}} = \{\alpha'_1, \alpha'_2, \dots, \alpha'_{c_1}, \alpha_{c_1+1}, \dots, \alpha_{c_2}, \alpha'_{c_2+1}, \dots, \alpha'_l\}$. Mutation takes a vector $\vec{\alpha}$ and produces a child $\vec{\alpha}_{\text{a}}$. The child $\vec{\alpha}_{\text{a}}$ is identical to $\vec{\alpha}$ except at a random index c where $\alpha_{\text{a},c} = \tilde{m}$ and \tilde{m} is a random number such that $-M \leq \tilde{m} \leq M$ for some mutation limit M . Creep is similar to mutation, but $\alpha_{\text{a},c} = \alpha_c + (0.5 - r) * C$, where r is a random number between 0 and 1 and C is defined as the “creep” rate. In most cases we set both M and C to be 1000. In general one should choose these values to be large enough to be effective at changing the population but small enough to keep the final shaking function phase to within approximately $\pm 2\pi$ radians.

The idea behind the mixing in this and other GAs is to use the components of the $A - A_{\text{die}}$ best individuals to produce a new generation with better fitness. When generation $j + 1$ has been generated, the simulation runs again and evaluates the new fitness values. This procedure iterates until a preset number of iterations has been performed or the fitness reaches a suitable level.

The goal of the GA is to minimize the fitness function $f(\vec{P}_{\alpha})$. The fitness function is constructed so that the final momentum state converges to the desired state, and the lattice “learns” to control the atoms. In general, any fitness function may be tailored to the individual requirements of each protocol in order to converge more quickly on a desired state. This is one of the distinct

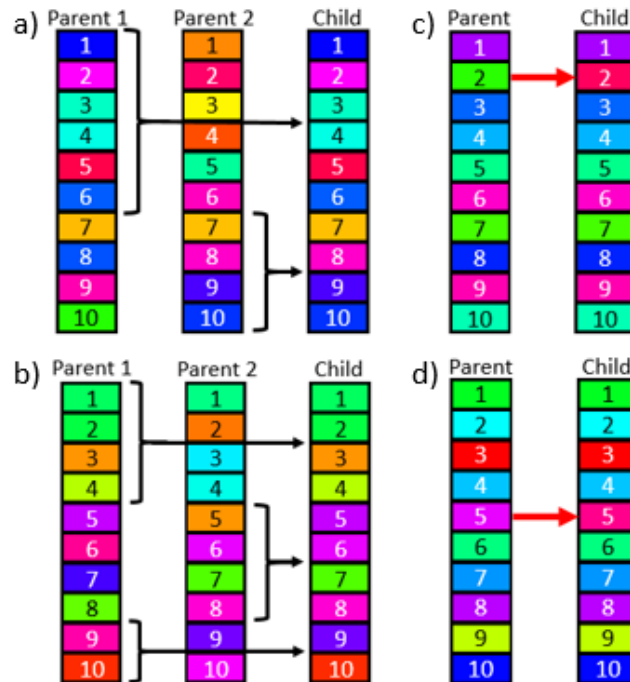


Figure 3.4: Each iteration, the genetic algorithm mixes the best individuals from the previous generation to create “children” that populate the next generation. Each individual is a vector of values where red corresponds to a minimum value and purple corresponds to a maximum value. In each case the indices of crossover or mutation are chosen at random. The numbers label the index and their color is chosen for legibility. The methods presented here are adapted from reference [92]. (a) One-point crossover at index 6. (b) Two-point crossover at indices 5 and 9. (c) Random mutation of the value at index 2. (d) Change of the value at index 5 through a small “creep” of the value.

advantages of the GA.

The GA has proven to be very easy to use in simulation, and it generally converges nicely on a solution. However, there is always an inherent degree of randomness in the GA, which is addressed by the optimal control methods described in the next sections. Furthermore, a prohibitively large number of individuals per generation are needed for convergence in our experimental systems. For example, we are currently limited by an experimental cycle time of roughly one shot per minute in order to avoid damaging our atom chip.⁸ Thus, if we take 3 experimental shots per individual and use $A = 20$, one iteration takes an hour. This is prohibitively slow, and the CRAB method described in Sec. 3.5 reduces this time to 15 minutes because in practice it requires only 5 distinct individuals to proceed.⁹

3.4 The Krotov method

In this section the basics of the Krotov method for the optimization of quantum mechanical systems described by the linear Schrödinger equation will be introduced. The derivation follows the one by Sklarz and Tannor in Ref. [104], and its generalization to quantum mechanical system is given in Refs. [88,89]. This section is admittedly the most initially opaque of the thesis, and the interested reader should study it and the references provided above until a thorough understanding is attained. Those wanting a general understanding of the method can skip to Sec. 3.4.2.

Given an initial state $|\psi(t=0)\rangle$ denoted $|\psi_0\rangle$, the time evolution of the state $|\psi(t)\rangle$ given by the a general evolution equation (here, the TDSE)

$$|\dot{\psi}(t)\rangle = f(\epsilon, t)|\psi(t)\rangle = -\frac{i}{\hbar}\hat{H}(\epsilon, t)|\psi(t)\rangle. \quad (3.6)$$

In equation (3.6), the control field $\epsilon(t)$ is the variable that is controlled when the optimization occurs (for the shaken lattice problem this will be the shaking function $\phi(t)$). After a time T we

⁸ Higher rep-rate systems have been demonstrated. In similar chip-based systems repetition rates of 0.3 Hz [38] were demonstrated in 2010 and rates in excess of 1 Hz were reported in 2013 [39].

⁹ We could likely reduce A in the GA by selecting part of the search space and reducing the number of Fourier frequencies l used in the algorithm, but the CRAB method has proven to be a useful and capable method in our experiments. As such, modifications to the GA to make it more amenable to experiment were not necessary.

denote the final state as $|\psi_T\rangle$. Given this, we must formulate the Krotov method as it specifically pertains to this problem. This will be done by first deriving the general Krotov method, then applying it to a quantum mechanical system governed by the TDSE.

3.4.1 A general formulation of the Krotov method

As a matter of terminology we define a process w to be the set that includes the state $|\psi(t)\rangle$ and the control field $\epsilon(t)$. One can define a general functional $J[w]$ that depends on the process w and is minimized via the Krotov method. This functional may take the form

$$J[w] = F(\psi_f) + \int_0^T dt f^{(0)}(t, \psi(t), \epsilon(t)). \quad (3.7)$$

In $J[w]$, the general functions F and $f^{(0)}$ depend on the final and intermediate states of the system, respectively, and the idea behind the Krotov method is to find the process w that minimizes J . Note that the functional can also include terms that penalize based on the projection of $|\psi\rangle$ into a forbidden state $|\psi_{\text{forbid}}\rangle$ denoted by a projection operator $\hat{P} = |\psi_{\text{forbid}}\rangle\langle\psi_{\text{forbid}}|$ [89]. In this work, however, we set $\hat{P} = 0$.

The next step is to define the functional $L[w, \Phi]$, which includes a general function $\Phi(t, \psi)$. This functional is written as

$$L[w, \Phi] = G(\psi_T) - \int_0^T dt R(t, \psi(t), \epsilon(t)) - \Phi(0, \psi_0) \quad (3.8)$$

where

$$G(\psi_T) = F(\psi_T) + \Phi(T, \psi_T) \quad (3.9)$$

and

$$R(t, \psi(t), \psi^\dagger(t), \epsilon(t)) = -g_a(\epsilon) + \frac{\partial\Phi}{\partial\psi} \frac{d\psi}{dt} + \frac{d\psi^\dagger}{dt} \frac{\partial\Phi}{\partial\psi^\dagger} + \frac{\partial\Phi}{\partial t} = -g_a(\epsilon) + \frac{d\Phi}{dt} \quad (3.10)$$

for some function $g_a(\epsilon)$. The functions R and G serve the same purpose in L as the functions $f^{(0)}$ and F did for J . That is, they separate dependence on intermediate and final time states. For any function Φ , one can show that $L[w, \Phi] = J[w]$ [104]. For any Φ finding an optimum to L will ensure an optimum in J .

The idea behind the Krotov method is that given complete freedom in choosing Φ , we can choose the Φ that maximizes L . We begin with the state $|\psi_0\rangle$ and an initial guess for the control field ϵ_0 that will then define an initial process w_0 . Then, the algorithm proceeds as outlined in Refs. [89, 104]

- Find Φ such that G is maximized and R is minimized (with respect to ψ); this maximizes L . Note that only ψ_T matters in the maximization of G .
- Given the Φ found in the previous step minimize R with respect to ϵ . Denote this minimizing value of epsilon as $\tilde{\epsilon}(t, \psi)$, which satisfies

$$\frac{\partial R}{\partial \tilde{\epsilon}} = 0 \quad (3.11)$$

- To remove the dependence of $\tilde{\epsilon}$ on ψ , note that $\tilde{\epsilon}$ and ψ must be consistent with each other through equations (3.6) and (3.11). This gives two equations that may be solved for the two unknowns, $\epsilon(t)$, $\psi(t, \epsilon)$. In practice, for the linear problem this is done by constructing Φ to first order in ψ , defining a function χ such that

$$\chi(t) = \frac{\partial \Phi}{\partial \psi^\dagger}. \quad (3.12)$$

The Hermitian conjugate of equation (3.12) defines an equation for χ^\dagger as well. Furthermore, the “initial” condition on χ is given by the extremum condition $\frac{\partial G}{\partial \Phi_T} = 0$ as

$$\chi(T) = \frac{\partial F}{\partial \psi_T^\dagger}. \quad (3.13)$$

Therefore, to first order, write [89]

$$\Phi = \chi^\dagger \psi + \psi^\dagger \chi. \quad (3.14)$$

Using this Φ , the minimization condition on R may be used to find the new field ϵ .

- The new process w defined by the new ψ , ψ^\dagger and ϵ now satisfies $J[w] < J[w_0]$, and so the current iteration has been successful. Therefore, the process may now be repeated until the functional J is sufficiently minimized.

Criteria for monotonic convergence may be established [88,89,104], but they will not be considered here. The mathematics presented here will be put to use in the next section.

3.4.2 The Krotov method applied to the linear TDSE

In this section we apply the method described in Sec. 3.4.1 to formulate the problem in terms of concrete equations that may be used in a general quantum mechanical system governed by the linear TDSE corresponding to a Hamiltonian $\hat{H}(t, \epsilon(t))$. As with the other optimization methods described in this chapter, for a given problem we know the initial state $|\psi^{(0)}(t=0)\rangle$, the desired final state after a final time T , $|\psi_f(t=T)\rangle$, and an initial guess for the controlling field $\epsilon^{(0)}(t)$. We denote the projection operator onto the final state as

$$\hat{D} = |\psi_f\rangle\langle\psi_f|. \quad (3.15)$$

Following [89], the equations for $|\chi(t)\rangle$ may be written

$$|\dot{\chi}(t)\rangle = -\frac{i}{\hbar}\hat{H}(\epsilon^{(0)})|\chi(t)\rangle \quad (3.16)$$

with “initial” condition.

$$|\chi(T)\rangle = -\lambda_0\hat{D}|\psi^{(0)}(T)\rangle \quad (3.17)$$

where λ_0 is a constant parameter set by the user. Propagating backwards in time using Eq. (3.16) and our final time condition Eq. (3.17), we can find $|\chi(t)\rangle$ for all times $t = 0$ to T .

Utilizing this and the minimization conditions (3.11) outlined in the previous section and choosing [88]

$$g_a(\epsilon) = \lambda_a(t)[\epsilon^{(1)}(t) - \epsilon^{(0)}(t)]^2, \quad (3.18)$$

it is possible to determine an equation for the correction to the field $\Delta\epsilon = \epsilon^{(1)} - \epsilon^{(0)}$, which leads to [89]

$$\epsilon^{(1)}(t) = \epsilon^{(0)}(t) - \frac{1}{\hbar\lambda_a(t)} \text{Im} \left[\left\langle \chi(t) \left| \frac{\partial \hat{H}}{\partial \epsilon} \right| \psi^{(1)}(t) \right\rangle \right]. \quad (3.19)$$

In equations (3.18) and (3.19), $\lambda_a(t)$ denotes a time-dependent shaping function, e.g. a slow turn-on and turn-off of an electromagnetic field as in Eq. (3.5). Using equations (3.16), (3.17), and (3.19), a new field ϵ may be found, and this process may be repeated until $\Delta\epsilon$ is sufficiently small.

This methodology can be extended to the simultaneous optimization of two states with the same control field. For example, if we begin our optimization with a state $|\psi_{1,0}\rangle$ and a state $|\psi_{2,0}\rangle$ subject to the same control field, we can optimize both simultaneously. Going through the same mathematics as above, one finds that

$$\epsilon^{(1)} = \epsilon^{(0)} - \frac{1}{\hbar\lambda_a(t)} \sum_j \text{Im} \left\langle \chi_j \left| \frac{\partial \hat{H}}{\partial \epsilon} \right| \psi_j^{(1)} \right\rangle. \quad (3.20)$$

3.4.3 Numerical implementation of the Krotov method

The numerical implementation of the Krotov method here is described as in Ref. [88]. Numerical implementation can be tricky due to the fact that the field $\epsilon^{(1)}(t)$ appears on both sides of equation (3.19), since $|\psi^{(1)}(t)\rangle$ also depends on $\epsilon^{(1)}(t)$. Therefore we use an interleaved grid method to solve the numerical problem. The first grid of $N_t + 1$ time points $t = 0, \Delta t, \dots, T$ propagates the states $|\psi(t)\rangle$ and $|\chi(t)\rangle$. The second grid of N_t time points $t = \Delta t/2, 3\Delta t/2, \dots, T - \Delta t/2$ propagates the correction to $\epsilon(t)$. The process proceeds in the following way:

Given the initial state $|\psi(0)\rangle$ and the initial guess for the field $\epsilon^{(0)}(t)$, the state $|\psi^{(0)}(t)\rangle$ is found by propagating the Schrödinger equation forward in time to $t = T$ on the initial grid, which results in a state $|\psi_T\rangle$. We then use $|\psi_T\rangle$ to find the “initial” condition $|\chi(T)\rangle$ for the time evolution of $|\chi(t)\rangle$ by applying Eq. (3.17). Equation (3.16) is propagated backwards in time on the initial grid from $t = T$ to $t = 0$ to find $|\chi(t)\rangle$ for all times t .

After we find $\chi(t)$ on the original grid, the new field is determined in the interleaved grid. For example, at $t = \Delta t/2$ the new field is given by

$$\epsilon^{(1)}(\Delta t/2) = \epsilon^{(0)}(\Delta t/2) - \frac{1}{\hbar\lambda_a(\Delta t/2)} \text{Im} \left[\left\langle \chi(0) \left| \frac{\partial \hat{H}}{\partial \epsilon} \right| \psi^{(1)}(0) \right\rangle \right]. \quad (3.21)$$

Using the result of Eq. (3.21), we can then propagate $|\psi(t)\rangle$ one step forward in time on the initial grid, giving us $|\psi(\Delta t)\rangle$. The field may then be propagated by adjusting equation (3.21) accordingly

to

$$\epsilon^{(1)}(\Delta t/2) = \epsilon^{(0)}(3\Delta t/2) - \frac{1}{\hbar\lambda_a(3\Delta t/2)} \text{Im} \left[\left\langle \chi(\Delta t) \left| \frac{\partial \hat{H}}{\partial \epsilon} \right| \psi^{(1)}(\Delta t) \right\rangle \right]. \quad (3.22)$$

In this way, the new field $\epsilon^{(1)}(t)$ may be found on the interleaved grid, and the new state $|\psi^{(1)}(t)\rangle$ may be found on the initial grid. This completes one iteration, and this interleaved grid method may be repeated until acceptable convergence has been reached.

In what has been presented here there is no limitation on the shaking bandwidth, which is problematic when we are trying to compare our simulation results to those of a real system. We can limit the bandwidth of the Krotov method to a frequency f_0 in the following way. We define a function γ as a subset of Eq. (3.19):

$$\gamma = \text{Im} \left[\left\langle \chi(\Delta t) \left| \frac{\partial \hat{H}}{\partial \epsilon} \right| \psi^{(1)}(\Delta t) \right\rangle \right]. \quad (3.23)$$

We can then define a band-limited γ_{HP} by applying a high-pass filter to γ with corner frequency f_0 . Then, we modify the projection in Eq. (3.19) to

$$\epsilon^{(1)}(t) = \epsilon^{(0)}(t) - \frac{1}{\hbar\lambda_a(t)} (\gamma - \gamma_{\text{HP}}). \quad (3.24)$$

This ensures that the high frequency components of γ and thus $\epsilon(t)$ are always filtered out. This works as long as our initial guess $\epsilon^0(t)$ does not have frequency components outside the bandwidth of $\gamma - \gamma_{\text{HP}}$.

3.4.4 The Krotov method applied to the shaken lattice

It is straightforward to adapt the mathematics presented here to the shaken lattice problem. Here, we denote the initial state as $|\psi_0\rangle$ and the desired final state as $|\psi_d\rangle$. Using Eq. (3.19) or (3.24), it is possible to find the correction to the field $\phi(t)$. The difference here is that the derivative $\frac{\partial \hat{H}}{\partial \epsilon}$ becomes $\frac{\partial \hat{H}}{\partial \phi} = \frac{\partial \hat{V}}{\partial \phi}$ for the shaking function $\phi(t)$.

Evaluating this for the lattice potential in Eq. (2.16) gives

$$\frac{\partial \hat{H}}{\partial \Phi} = -\frac{V_0}{2} \sin(2k_L x + \Phi(t)). \quad (3.25)$$

From here Eq. (3.19) may be modified by using Eq. (3.25) to obtain

$$\phi^{(1)}(t) = \phi^{(0)}(t) + \frac{1}{\hbar\lambda_a(t)} \frac{V_0}{2} \text{Im}[\langle \chi(t) | \sin(2k_L x + \phi(t)) | \psi^{(1)}(t) \rangle]. \quad (3.26)$$

Using Eq. (3.26), the equations (3.21) and (3.22) may be appropriately modified, and the interleaved grid method may be used. Furthermore, this method may also be modified to band-limited shaking using Eq. (3.24). Using this we are now able to use the Krotov method to optimize the lattice shaking function $\phi(t)$. For the work presented in Chapter 4 we have simply used $\phi^{(0)}(t) = 0$ as our initial guess to initialize the optimization algorithm.

When implementing the Krotov method one has freedom in how they choose the parameters λ_0 and $\lambda_a(t)$. As stated before, $\lambda_a(t)$ defines the envelope of the shaking function and can be used, for example, to ensure smooth turn-on and turn-off. In practice we use the envelope function Eq. (3.5) to ensure this smooth behavior, and we scale the envelope function by the constant λ_a . Written this way, we can see that λ_a sets a threshold for how much $\phi(t)$ can change from iteration to iteration. The parameter λ_0 is a weight on the final state. It can be shown [89] that in order to perform minimization, λ_a must be greater than zero, and λ_0 must be less than zero.

In our implementation of the Krotov method, we set $\lambda_0 = -1$ and initialize $\lambda_a = 0.1$. The parameter λ_0 remains constant through optimization. However, λ_a is dynamic, as described in the next paragraph. After each iteration we check the error in the final state ψ_f relative to the desired state ψ_d by checking the error in the momentum population vector \vec{P}_f relative to the desired state \vec{P}_d ¹⁰.

If the error in iteration $j + 1$ is greater than the error in iteration j , we adjust λ_a , doubling its value so that $\lambda_{a,\text{new}} = 2\lambda_{a,\text{old}}$. The reasoning behind this is that if the error increased, we have modified $\phi^{(j+1)}(t)$ too dramatically. With this, the stopping condition for the Krotov method is such that the loop breaks after the error has reached a lower threshold or a certain number of iterations have taken place. One could also implement a break condition that stops once λ_a becomes too large (as the modifications to $\phi(t)$ are negligible for large enough values of λ_a). In

¹⁰ See Sec. 3.1.2 for a description of the momentum population vectors.

practice a value of $\lambda_{a,\text{break}} = 10^8$ is sufficient to achieve reasonable convergence.

One of the drawbacks of the Krotov method is that it is not suitable for use in an experimental system. This is an unfortunate consequence of the fact that backwards time propagation is necessary for the method to work. However, the Krotov method does not have the inherent randomness of the genetic algorithm; if the system is set up in the same way for two separate optimization runs, the two results will be identical. The next section will introduce the CRAB method as an optimal control method that *is* suitable for experimental implementation.

3.5 The CRAB method

3.5.1 Formulation of the CRAB method

This section will describe our implementation of the CRAB method in some detail. Excellent references to learn about this method are given in Refs. [19, 31, 95], and the information in this section is distilled from these references.

Like the Krotov method, this method can be used to determine an optimal path from an initial quantum state $|\psi_0\rangle$ to a final state $|\psi_F\rangle$ in a time T given a control field $\epsilon(t)$ in the Hamiltonian describing the system. As before, in our case, the control field is the shaking function $\phi(t)$. As with the GA the error function is reasonably flexible [31], but we determine error as in Eq. (3.4).

To implement the CRAB method we solve the TDSE with a guess for the control field $\phi^{(0)}(t)$.¹¹ The CRAB method attempts to correct this guess to a field $\phi^{(\text{opt})}(t) = \phi^{(0)}(t)g(t)$. In general, the function $g(t)$ is expanded in a basis set of functions; in this work, we use the Fourier basis so that [31]

$$g(t) = 1 + \frac{\sum_n A_n \cos(\omega_n t) + B_n \sin(\omega_n t)}{\lambda(t)} \quad (3.27)$$

where A_n and B_n are the Fourier cosine and sine coefficients, respectively and ω_n are the frequencies used in the basis set. The function $\lambda(t)$ allows us to set the envelope of our shaking function the same way as in Sections 3.3 and 3.4. We can then formulate our problem in terms of the control

¹¹ We generally set $\phi^{(0)}(t) = 1$ when using this method.

parameters A_n , B_n and ω_n . While these parameters could be continuous, in practice n is limited to a finite number to narrow the search space. Further narrowing of the search space can be done by choosing the frequencies of the control field ω_n intelligently based on knowledge of the system. We will explore the narrowing of the optimization search space (or “landscape”) in more detail in Chapter 7. Unless otherwise stated we choose the values of ω_n randomly in our band of interest. Thus, for a given number of frequencies N_ω our control space has $2N_\omega$ parameters corresponding to a Fourier sine and cosine amplitude for each ω_n .

Our problem has been cast into a simpler one that can be solved in more traditional methods, e.g. gradient-descent methods. Gradients are expensive to compute both in simulation and experiment, so we avoid gradient-based methods. Instead, we work mainly with the gradient-free Nelder-Mead method, described in the next subsection.

3.5.2 The Nelder-Mead algorithm

The Nelder-Mead (NM) algorithm was developed in 1965 by J. A. Nelder and R. Mead [87].¹² The method described here is more or less identical to the method described in the initial paper. What follows is a brief description of the method with an emphasis on application to the CRAB method and the shaken lattice problem.¹³

The NM algorithm proceeds as follows: we pick a number of points N_c . These points are represented as vectors in the $2N_\omega$ -dimensional search space, and they define the vertices of a “simplex” in the search space. The point of the NM algorithm is to warp the simplex around the geometrical landscape until it converges on a local minimum. The fact that the NM scheme only finds a local minimum is an important point that leads to the modification of the CRAB algorithm discussed in Sec. 3.5.3.

After the simplex vertices are chosen we run the simulation (or experiment) and evaluate the error E_n for the n^{th} vertex \vec{x}_n using Eq. (3.4). The vertices are then ranked from smallest error

¹² At the time, Nelder and Mead were at the National Vegetable Research Station in Warwick, UK.

¹³ As unprofessional as it is to cite Wikipedia, the Wikipedia article on the NM algorithm is far and away the best resource that we have found in terms of adaptable pseudocode. Of course, given the malleability of Wikipedia (and the transient nature of Internet-based resources in general), this may change in the future.

to largest (in the case of minimization, which is all that is considered in this work). Excluding the worst-ranked vertex \vec{x}_{\max} with maximum error E_{\max} , we calculate the centroid of the remaining vertices. For a set of discrete points the centroid of the set is just the mean of the points. $\vec{x}_{\text{ct}} = \sum_n \vec{x}_n / N_c$.

After we have found the centroid \vec{x}_{ct} , we reflect the worst vertex about the centroid in an attempt to improve the simplex. This reflected vertex is found via

$$\vec{x}_{\text{refl}} = \vec{x}_{\text{ct}} + \alpha_{\text{NM}}(\vec{x}_{\text{ct}} - \vec{x}_{\max}) \quad (3.28)$$

where α_{NM} is a ‘‘reflection coefficient’’ controllable by the user. We use the standard value $\alpha_{\text{NM}} = 1$ in this work. The error E_{refl} of the reflected vertex is found, and the algorithm proceeds in the following ways based on this error:

- If $E_{\text{refl}} > E_{\text{best}}$, where \vec{x}_{best} and E_{best} are the best vertex and its error, respectively, then the simplex has reflected in a better direction. We then expand the simplex in that direction by performing an ‘‘expansion’’. This expanded vertex \vec{x}_{exp} is found by

$$\vec{x}_{\text{exp}} = \vec{x}_{\text{ct}} + \gamma_{\text{NM}}(\vec{x}_{\text{ct}} - \vec{x}_{\max}). \quad (3.29)$$

As with reflection, γ_{NM} is the ‘‘expansion coefficient’’ and is generally (and in this work) set to $\gamma_{\text{NM}} = 2$. If the error in the expanded vertex E_{exp} is less than E_{refl} , then we replace \vec{x}_{\max} with \vec{x}_{exp} and go back to step 1 (recalculating the centroid and reflecting about it). Otherwise, we replace \vec{x}_{\max} with \vec{x}_{refl} and go back to step 1.

- If $E_{\text{refl}} > E_{\max}$, we have moved in the wrong direction or jumped over a minimum.¹⁴ . The reflected vertex then undergoes a contraction of the form

$$\vec{x}_{\text{cont}} = \vec{x}_{\text{ct}} + \rho_{\text{NM}}(\vec{x}_{\text{ct}} - \vec{x}_{\max}) \quad (3.30)$$

where $\rho_{\text{NM}} = 1/2$ is the contraction coefficient. The same replacement rules apply here as with expansion.

¹⁴ Some methods contract if E_{refl} is greater than $E_{\max-1}$, that is if the reflected vertex has error greater than the second-worst vertex. In hindsight this is probably the better form of the algorithm and an error in the code used to obtain the results presented in this thesis. This code still works, but an adjustment may provide faster convergence.

- In all other cases, we reduce the size of the simplex around the best vertex \vec{x}_{best} . This is because we are likely surrounding a minimum. The reduction of the simplex replaces all vertices but the best with

$$\vec{x}_{n,\text{red}} = \vec{x}_{\text{best}} + \sigma_{\text{NM}}(\vec{x}_n - \vec{x}_{\text{best}}) \quad (3.31)$$

where we set the reduction coefficient $\sigma_{\text{NM}} = 1/2$. The algorithm then returns to step 1.

An iteration of the NM algorithm has occurred each time one of the itemized steps is taken and returns to a calculation of the centroid.

In simulation we have some flexibility when choosing the dimensionality $2N_\omega$ of the search space and the number of vertices N_c that we choose to span that space. We will choose values between 10 – 30 for these parameters. However, in experiment we must be more selective. In the experimental implementations of the CRAB method described in Chapter 6 we set $N_\omega = N_c = 5$ and take three “shots” (different experimental runs with the same experimental parameters) for each point considered in the NM algorithm. We find that this method is the fastest way to produce results (we can run 15 – 18 iterations of the algorithm in roughly 4 – 6 hours).

3.5.3 Avoiding local minima in the CRAB algorithm

The CRAB algorithm is (like the GA) limited in that it can only find *local* minima and hence there is no guarantee that one has found a *global* minimum when the CRAB algorithm converges on a point in the search space. Thus, a method was developed in Ref. [95] that attempts to circumvent this. This new method, called dressed CRAB or dCRAB, is the method used in our experiments.

The difference between CRAB and dCRAB is that when the simplex is stuck in a potentially suboptimal local minimum, the search space is changed. In many ways one can think of it as “re-rolling” to get better performance. The potentially suboptimal control field $\phi_{\text{sub}}(t)$ is modified by a new control pulse $\phi_{\text{new}}(t)$ such that the control field is given by $\phi(t) = \phi_{\text{old}}(t) + \phi_{\text{new}}(t)$. The old field has the form derived from Eq. (3.27)

$$\phi_{\text{old}}(t) = \phi_0(t) \left(1 + \frac{\sum_n A_n \cos(\omega_n t) + B_n \sin(\omega_n t)}{\lambda(t)} \right) \quad (3.32)$$

with (suboptimal) coefficients A_n and B_n corresponding to the original set of frequencies ω_n . The new field has the form

$$\phi(t) = \phi_{\text{old}}(t) + \phi_0(t) \left(1 + \frac{\sum_n A'_n \cos(\omega'_n t) + B'_n \sin(\omega'_n t)}{\lambda(t)} \right) \quad (3.33)$$

where the new optimization parameters are A'_n and B'_n and a new set of frequencies ω'_n has been chosen in the bandwidth of interest. The NM optimization then proceeds with the new simplex defined by A'_n and B'_n . This process can then be repeated as many times as one wishes, defining a new set of coefficients and frequencies A''_n , B''_n and ω''_n , and relegating the suboptimal $\phi(t)$ to $\phi_{\text{old}}(t)$ and optimizing a new $\phi(t)$ on the new set of parameters.

In practice this “re-rolling” is done when the optimization error stays stagnant (or increases due to experimental noise or drift) for more than two iterations. The downside to this is that the complexity of the shaking function (in terms of the number of frequency components it contains) increases, but the method works for experimental optimization shaking. More insight into the downsides of the presence of unnecessary shaking frequencies in $\phi(t)$ will be given in Chapter 7.

3.6 What’s next

With the basis for the simulation methods and the optimization algorithms laid out in this chapter, we will dive into simulation results using the GA and Krotov methods in Chapter 4. We will explain the basics of the experiment and describe experimental optimization and results using the CRAB algorithm later in this thesis.

Chapter 4

Simulating the shaken lattice interferometer

This chapter will describe the implementation of simulation methods detailed in the last chapter for the shaken lattice interferometer. We will also present results of the simulations of the shaken lattice interferometer. In particular, we will describe how the interferometer is simulated, optimized, and tested. Proof-of-principle results with the Krotov method will be presented. Then, we will show more realistic (read: lower shaking bandwidth) results using the genetic algorithm. From these results we will determine that the interferometer sensitivity scales with the square of the interrogation time T_1 . We will also present results that show that we are able to tune the interferometer transfer function. For example, we are able to optimize the interferometer to measure perturbations on a DC bias acceleration or an AC acceleration signal. To our knowledge, this is the first atom interferometer capable of tuning its sensitivity in this way. Finally, we will describe some next steps; these will be elaborated on further in Chapter 7.

4.1 The simulation and optimization of the shaken lattice interferometer

In this section we will outline the basics of optimization of the shaken lattice interferometer. Note that much of this section and Sec. 4.4 is adapted from our work in Ref. [119]. We will begin by reviewing the main “ingredients” of interferometry that were introduced in Chapter 1. This will allow us to define the language that we use throughout the rest of this thesis. Then we will discuss the fitness and error functions used in this chapter and Chapters 6 and 7. To illustrate the functionality of the optimization algorithms used in this work, we will present exemplary learning

curves for the genetic algorithm and Krotov method.

It is important to note that in these proof-of-principle simulations we do not consider any nonlinearities due to atom-atom interactions. To first order when considering atom-atom interactions in BEC simulations, the effective potential due to the condensate mean-field is modified such that [82]

$$V_{\text{eff}} = \frac{V_0}{1 + 4C} \quad (4.1)$$

where $C = g/E_r$ and g is the typical atom-atom interaction parameter in the Gross-Pitaevskii equation [90]

$$g = \frac{4\pi n_0 \hbar^2 a_s}{m}. \quad (4.2)$$

In Eq. (4.2), n_0 is the peak condensate density and a_s is the s-wave scattering length, which for Rb-87 is $a_s \approx 95a_0$ where $a_0 = 0.529 \times 10^{-11}$ m is the Bohr radius [36, 76].

In neglecting atom-atom interactions we assume that the lattice is sparsely populated, e.g. through the use of a 3-D lattice where in the non-shaken dimensions the lattice is so deep that few atoms populate each site, but in the shaken dimension, the lattice is shallow enough that the atoms remain delocalized. The infinite lattice approximation used here is reasonable if the lattice beam has a Rayleigh range much larger than the lattice wavelength. This is true for lattice beams with near-infrared wavelengths and waists on the order of tens of microns; we use such a scheme in our experiments. We find that although our simple simulation results cannot be applied directly to experimental results, the basic ideas do translate to experiment. This is explored in detail in Chapter 6.

4.1.1 The “ingredients” of interferometry

As stated in Chapter 1 the vast majority of interferometers contain many of the same basic “ingredients:” a source of waves, a beamsplitter, a means for the waves to propagate, reflect, and propagate further until they are recombined. For example, see the optical Mach-Zehnder and Michelson interferometers in Fig. 4.1. In a typical Mach-Zehnder interferometer a source of light

is split by a beamsplitter and propagates until it is reflected by the mirrors (M). After reflection the light continues to propagate until it is recombined by another beamsplitter. Detectors at each output measure the intensity of the light; when the interferometer is perfectly aligned one detector will measure a high (low) intensity when the other detector measures a low (high) intensity. By measuring how many light-to-dark fringes pass one of the detectors, one can make a measurement of the path length change between two arms. This path length difference could be due to a rotation or a change in the refractive index through which the light in one arm travels relative to the other. The Michelson interferometer operates similarly, but the splitting and recombining operations occur in the same optic and there is only one accessible detector¹. The Michelson scheme can measure motion, e.g. of one of the reflecting mirrors.




Figure 4.1: A cartoon of (a) an optical Mach-Zehnder interferometer and (b) a Michelson interferometer, with the basic “ingredients” of the interferometry sequence labeled.

The labels given in Fig. 4.1 are those that we will use to denote our interferometer operations in this chapter and Chapter 6. More specifically, we will begin with atoms in the ground Bloch state of the optical lattice (see Fig. 2.3) [50] and phase modulate (shake) the lattice to change the atoms’ momentum state, as in Ref. [92]. In this way we cause the atoms to undergo splitting, propagation, reflection, reverse propagation, and recombination while trapped in the shaking optical lattice.

As a reminder, the Hamiltonian of the system is described by Eq. (2.16), and the “knob” that we have to turn is the shaking function $\phi(t)$. The main goal of shaken lattice interferometry is to

¹ The other beam path goes back into the laser and is usually not measured.

find an optimal shaking function $\phi(t)$ that transforms an initial state ψ_0 to a desired final state $|\psi_d\rangle$. To accomplish this we run our optimization algorithms to find a shaking function that accomplishes this task, resulting in a final state $|\psi_f\rangle$ that is similar to $|\psi_d\rangle$ within a specified precision. In this way we construct a sequence of shaking functions that cause the atoms to undergo a conventional interferometer sequence. We describe the simulations of the interferometer sequence in some detail below.

When splitting, we wish to split the atoms' wavefunction so that half of the atoms are moving in the $+x$ direction with momentum $2n\hbar k_L$ and the other half are moving in the $-x$ direction with momentum $-2n\hbar k_L$. That is, the final split state is a pair of equal-amplitude waves moving with equal and opposite momenta. In this chapter we will only consider the first-order splitting case where $n = 1$. Higher-order splitting will be treated in Chapter 7.

As the split states are not eigenstates of the lattice we find another shaking protocol that maintains the atoms' momentum states as they propagate. Three additional shaking protocols are needed to reflect, reverse-propagate, and recombine the atoms. Once the different phase-modulation protocols are learned, interferometric measurements may be performed by “stitching” together the protocols, resulting in the atoms undergoing the full interferometer sequence as illustrated in Fig. 4.2. The sinusoidal envelope function Eq. (3.5) allows us to stitch the protocols together with minimal discontinuity.

In simulation, during propagation and reflection it is important to control each momentum component separately because the initial and final momentum populations are the same. In these cases the split-step TDSE solver is run twice. The first run begins with an initial state where all atoms are in the $-2n\hbar k_L$ state and modulates the lattice with a function $\phi_{pr}(t)$. The second run begins with atoms in the $2n\hbar k_L$ state and applies the same phase modulation function $\phi_{pr}(t)$. When propagating the atoms, the desired state is identical to the initial state for each momentum component. The optimization algorithm then sums the error of both final states and optimizes $\phi_{pr}(t)$. The optimal $\phi_{pr}(t)$ will thus propagate the atoms without “crosstalk” between the two momentum states when applied to the linear combination of the two states. For reflection this




Figure 4.2: A cartoon of the full interferometer sequence, with atom position plotted along the y -axis versus time on the x -axis. Blue clouds represent atom wavepackets interacting with the shaken lattice. Atoms begin in the ground Bloch state of the lattice (Fig. 2.3) and are split into two oppositely propagating wavepackets. The atoms are then reflected, reverse-propagated, and recombined back into their initial state (in the absence of a signal), thus completing the interferometer sequence.

simultaneous two-state optimization proceeds in the same way, but for an initial state with $\pm 2\hbar k_L$, the final state has momentum $\mp 2\hbar k_L$. After reflection the atoms are again propagated to maintain their (now reversed) momentum state.

The final step of the interferometry sequence recombines the two split waves. The recombination scheme is as follows: the initial and desired states considered in splitting are swapped, such that the desired state is now the initial state and vice versa. The optimization algorithm is run to find a modulation sequence that returns all of the atoms in the two split matter waves to the ground Bloch state. It is this final state of the atoms that changes when a signal is applied.

4.1.2 Fitness and error functions

In this work, each of the optimization algorithms needs a “fitness” or an “error” function around which the optimization is centered. The terminology “fitness” is generally used when one is referring to a genetic algorithm (in line with the rest of the biologically-inspired terminology used when describing a GA), and when using optimal control algorithms (like the Krotov or CRAB methods), one usually refers to an “error.” This work will follow the conventions laid out in this paragraph. For the Krotov and CRAB methods we generally optimize using the percent error defined by Eq. (3.4). More details will be given in Chapter 6.

In all cases, the ultimate goal is to minimize Eq. (3.4). With the genetic algorithm, however, we have implemented a bit more control over the fitness function.² For example, in Ref. [119] when splitting (Sec. 4.4.1) we minimize a fitness function of the form

$$f(\vec{P}_\alpha) = |\vec{P}_d - \vec{P}_\alpha| + \sum_{|n| \neq 1} |P_{d,n} - P_{\alpha,n}| + \left| \frac{P_{\alpha,1} - P_{\alpha,-1}}{P_{\alpha,1} + P_{\alpha,-1}} \right|. \quad (4.3)$$

The first term in Eq. (4.3) quantifies the difference between the two momentum state populations for the final state and the desired state. The second term penalizes for atoms not in the desired $\pm 2\hbar k_L$ momentum states. The last term penalizes for asymmetry in the $\pm 2\hbar k_L$ momentum states. The penalizing terms are added to improve the rate of convergence to the desired state. Similar

² There is nothing keeping us from modifying the optimization functional with the Krotov or CRAB methods to implement similar control, but we have not found that this is necessary.

fitness functions are used for all shaking protocols optimized with the GA, and in general any fitness function may be tailored to the individual requirements of each protocol in order to converge more quickly on a desired state.

4.1.3 Learning curves

An exemplary learning curve for first-order splitting using the genetic algorithm is shown in Fig. 4.3. This learning curve plots the average fitness of all individuals as well as the fitness of the best individual from Eq. (4.3) versus the number of generations (denoted here as iterations). We see that the majority of the learning occurs within the first 100 generations with refinement occurring in the last iterations. The best fitness after 1000 iterations is 0.0011, corresponding to an error of $9.9 \times 10^{-6}\%$, which is the same order of magnitude as in Table. 4.1.³ Note that even though the average fitness hovers around the same value, the best fit individual can only decrease due to elitism. Furthermore, because the fitness function and the percent overlap defined by Eq. (3.4) are not the same, a lower fitness can correspond to a higher percent overlap, but this generally only occurs near the minimum (i.e. best fit) that the GA has found.




Figure 4.3: Learning curves for first-order splitting using the genetic algorithm showing (a) the fitness of the best individual (black) and the average fitness of all individuals (red) and (b) the percent error of the final state relative to the desired split state.

³ The values given here and in Table 4.1 do not take into account the limitations due to phase error in the split step method (discussed in Sec. 3.1.2).

A sample learning curve for the Krotov method is shown in Fig. 4.4. With this method, we use Eq. (3.4) to determine the error. In this case the learning curve is not monotonically decreasing, but whenever the error increases the value of λ_a is adjusted (see Sec. 3.4). Thus, when the optimization algorithm overshoots the correction to the shaking function $\phi(t)$, the algorithm is adjusted to make smaller corrections. The final error is $8.0 \times 10^{-5}\%$. We find that when the simulation is converged, λ_a increases dramatically. Thus, in practice we can break the loop early, e.g. when $\lambda_a > 10^8$. This allows us to speed up our simulations and conserve computing resources.

4.2 Adding a signal to the shaken lattice interferometer

After optimization it is of interest to simulate how the interferometer would respond to an applied inertial signal. This can be done by modifying the shaking Hamiltonian in Eq. (2.16) to

$$H(x, t) = -\frac{\hat{p}^2}{2m} - \frac{V_0}{2} \cos(2k_L x + \phi(t)) + ma_x x \quad (4.4)$$

where a_x is the component of the acceleration \vec{a} in the lattice propagation direction \vec{x} . However, this “tilt” of the lattice ruins the periodic boundary conditions outlined in Sec. 3.1.1. Thus, we must set up our wavefunction so that it is no longer infinite in position space. To do this, we change the momentum space wavefunction so that the momentum components are no longer delta functions but rather Gaussian functions of a width δp centered around $p = 2n\hbar k_L$. Then, when we Fourier transform the momentum state wavefunction it is now finite in extent and the “tilt” of the lattice due to the acceleration is no longer as problematic. Specifically, if we use a value of $\delta p = 0.01\hbar k_L$ and look at small applied accelerations a_x , we are able to reasonably accurately estimate the behavior of the atoms as they propagate in the lattice. For larger values of a_x , we need to use $\delta p \approx 0.05\hbar k_L$ for more accurate results.⁴

⁴ We have verified that our sensitivity results do not change dramatically when the Gaussian width of the atom wavefunction is changed from $0.01\hbar k_L$ to $0.05\hbar k_L$ and the scaling is still consistent with the expected T^2 scaling.




Figure 4.4: Learning curve for first-order splitting using the Krotov method showing (left axis, blue) the percent error of the final state relative to the desired split state and (right axis, red) the λ_a parameter as a function of iteration number.

4.3 Proof-of-principle results using the Krotov method

In this section we will describe proof-of-principle results using the Krotov method. These results were obtained with few restrictions on the bandwidth and amplitude of the shaking functions. That is, the shaking bandwidth was as high as a few 100 kHz and the amplitude of the shaking was as high as tens of radians, in contrast to the shaking functions produced by the GA (see Fig. 4.8). Thus, the optimizations using the GA in Sec. 4.4 produce results that are more suitable for comparison to a realistic experimental scenario. Despite these drawbacks, these simulations are presented here as they are useful illustrations of some important concepts in shaken lattice interferometry. As in Sec. 4.4, there is nothing inherently keeping us from doing these simulations with the GA, but these simulations were done with the Krotov method as a means of illustrating its utility.

4.3.1 Initial interferometer optimization

As in the GA simulations detailed in Sec. 4.4.1 the initial optimization of the shaken lattice interferometer using the Krotov method implemented splitting, reflection, and recombination with a number of propagation steps in between. By varying the number of propagation steps, we vary the interrogation time and thus sensitivity of the interferometer. Two examples of optimized interferometry sequences are shown in Fig. 4.5 where we plot the momentum state population vector $\vec{P}(t)$ versus time for the optimized interferometer sequence. This figure shows the atom dynamics during the splitting, reflection, and recombination protocols, and the high-bandwidth simulations maintain the atoms' split momentum state during propagation. The simulated shaking functions that produced these plots will be used in the next section when we add a simulated acceleration signal and glean some basic insight into some important characteristics of the shaken lattice interferometer.

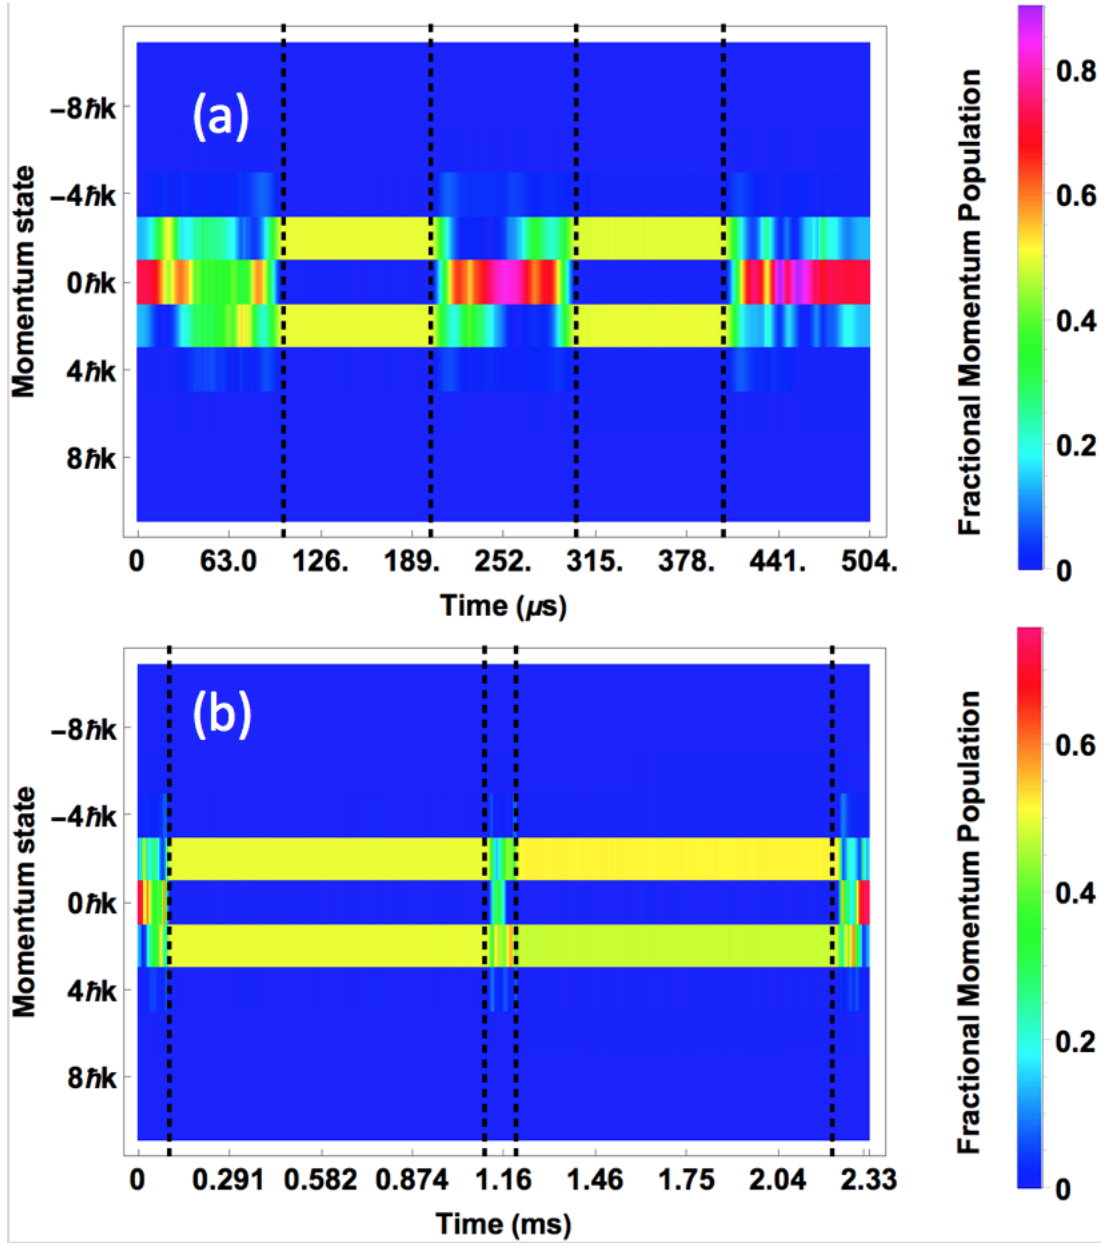


Figure 4.5: High-bandwidth Krotov simulations of the shaken lattice interferometer for two different interrogation times T_I (note the different units in the y-axes). These plots show the atom momentum state population (quantized in units of $2\hbar k_L$) versus the shaking time. The color shows the relative population in each of the separate momentum states, and the black dashed lines separate (from left to right) the splitting, propagation, reflection, reverse propagation, and recombination regions. At a given time t , the colors representing the fractional momentum state population are a visual representation of the vector $\vec{P}(t)$ described in 3.3.

4.3.2 Adding a signal

Once the high-bandwidth Krotov interferometers are optimized, we add a signal as in Sec. 4.2. The results for both interferometers shown in Fig. 4.5 are shown in Fig. 4.6. Here we are plotting the final momentum state \vec{P}_a after the full interferometry sequence as a function of the applied signal a . That is, we are observing how the interferometer state changes as a signal is applied; it is the measurement of this final state that allows us to measure a meaningful quantity (here, the acceleration a).

We notice three things immediately. First, Fig. 4.6 shows that the sensitivity of the interferometer increases as the interrogation time increases, which makes intuitive sense. We will explore the sensitivity scaling of the shaken lattice interferometer in Sec. 4.4.

Less intuitively, we see in Fig. 4.6 that the final state of the interferometer depends on the direction of the applied signal; that is, we can measure both the magnitude and direction of acceleration with this system. This is a rather unique characteristic of this interferometer relative to other atom interferometry schemes. The directionality arises because when the lattice is shaken, the symmetry of the system is broken (for example, if we shook with $\phi(\tilde{t}) = -\phi(t)$, we would get different results).

Finally, we note that in contrast to conventional light-pulse atom interferometry (e.g. [20, 29, 47, 58]) there are more than two states interfering. We see momentum state population in states up to $\pm 4n\hbar k_L$. Therefore, we have effectively built a multi-path interferometer, and the dynamics of the system are more complex than in the typical two-state atom interferometer.

The high-bandwidth simulations with the Krotov method can tell us one more interesting thing. In most interferometers one can easily define an “arm” of the interferometer. As such, one usually refers to the phase difference between one arm relative to another. However, in shaken lattice interferometry, when we add a signal the definition of an “arm” of the interferometer becomes ambiguous and difficult to define. To show this we have plotted the result of adding an acceleration signal to the shaken lattice interferometer in Figure 4.7. We see that even though our initial




Figure 4.6: The final state after an acceleration signal is applied for each of the optimized high-bandwidth Krotov shaken lattice interferometers shown in Fig. 4.5. Each plot shows the final atom momentum state population (quantized in units of $2\hbar k_L$) after shaking is plotted versus the applied acceleration. The color shows the relative population in each of the separate momentum states. This can be thought of as a visual representation of the vector \vec{P}_a described in e.g. Sec. 4.4.4.

optimization seems to allow us to define two arms of the interferometer, once a signal is applied, the momentum state population starts changing, although it remains quantized [112]. Thus we must get more creative when defining the interferometer sensitivity, and this leads us to the use of the classical Fisher information and Cramer-Rao bound, described in more detail in the next section.

4.4 Determining the interferometer sensitivity

The work in this section will largely follow what is published in Ref. [119] using the GA, and most of the figures in this section come directly from the publication. Note that there is nothing preventing one from using the Krotov method to obtain results similar to those presented here, but we used the GA because at the time this work was done, the framework for doing lower-bandwidth simulations using the GA was more developed. As in Ref. [119] we will begin by describing the initial optimization of the interferometer using the genetic algorithm, followed by an analysis of the robustness of the optimized interferometer. Results showing how we can tune the transfer function of the shaken lattice interferometer will be presented. We will then discuss how we determine the interferometer sensitivity using the classical Fisher information and present some speculative calculations of the ultimate interferometer sensitivity based on reasonable atom numbers and interrogation times.

4.4.1 Initial optimization with the genetic algorithm

To implement the Michelson interferometer sequence we begin with atoms in the ground Bloch state of the lattice. The momentum population of the ground state is shown in Fig. 2.3(b) [50]. The GA then uses Eq. (4.3) as the fitness function to find the optimal shaking protocol that splits the atom wavefunction. As stated in the last section the final split state is a pair of equal-amplitude waves moving with equal and opposite momenta. An optimized shaking protocol for splitting is shown in Fig. 4.8. The optimized splitting protocol shows the slow turn-on/off and the band-limited nature of the shaking function.




Figure 4.7: Plot similar to those shown in Fig. 4.5, except with an acceleration signal added. We see that the definition of an interferometer “arm” quickly breaks down due to the interference between different momentum states in the presence of a signal.




Figure 4.8: Shaking protocol for the optimal splitting shown in Table 4.1. The envelope function ensures slow turn-on and turn-off at the endpoints. (inset) Fourier transform of the bandwidth-limited shaking function.

For each shaking protocol Table 4.1 shows the initial, desired, and final momentum states, as well as the difference between the optimized final state and the desired state as defined in Eq. (3.4). The best result of 5 runs is shown, but in all cases, the variation is below 0.1%. Variation of the optimized results with respect to variations in parameters will be discussed in the Sec. 4.4.2. Section 4.4.4 will discuss the interferometer sensitivity and scaling with the interrogation time.

Table 4.1: Genetic algorithm results, best of 5 optimization runs.

Protocol	State	Momentum population ¹					% Diff.	Bandwidth (kHz) ²
		$-4\hbar k_L$	$-2\hbar k_L$	$0\hbar k_L$	$2\hbar k_L$	$4\hbar k_L$		
Split	Init.	0.0026	0.1345	0.7259	0.1345	0.0026		
	Des.	0	0.5	0	0.5	0		
	Final	0	0.4999	0.0001	0.4998	0	3.6×10^{-6}	57.8
Prop.	Init.	0	0.5	0	0.5	0		
	Des.	0	0.5	0	0.5	0		
	Final	0.0006	0.4992	0.0010	0.4980	0.0006	2.5×10^{-4}	53.8
Refl.	Init.	0	0.5	0	0.5	0		
	Des.	0	0.5	0	0.5	0		
	Final	0.0012	0.4958	0.0037	0.4980	0.0008	1.8×10^{-3}	55.8
Recomb.	Init.	0	0.5	0	0.5	0		
	Des.	0.0026	0.1345	0.7259	0.1345	0.0026		
	Final	0.0026	0.1345	0.7258	0.1343	0.0026	5.4×10^{-6}	55.8

4.4.2 Robustness and stability of the shaken lattice interferometer

This section will deal with the effects of errors and instability in shaken lattice interferometry. The analysis here has allowed us to build the shaken lattice interferometer described in Chapters 5 and 6. First, we will discuss the effects of errors in the system in Sec. 4.4.2.1, motivating the use of a closed-loop system in experiment. These sources of error are considered in other conventional atom interferometry schemes, and potential solutions for the specific case of shaken lattice interferometry are given. In many cases if the lattice power does not drift and the optics are clean and stable, most errors can be corrected for in a closed-loop system. Then we will move on to describing simulations that verify the robustness of the shaken lattice interferometer in Sec. 4.4.2.2. Here, analysis is done with respect to the optimal splitting function shown in Fig. 4.8, but the results presented here should hold for all stages of the interferometer due to the similarities in shaking time and bandwidth. As the interrogation time increases, however, we expect the stability requirements to become more stringent, and we present the results here only as a general example.

4.4.2.1 Errors in shaken lattice interferometry

Ideally, in practice an optimized shaking function can be found computationally and adjustments may be made by running an experiment with the learning algorithm in a closed-loop scheme, which has been done before in cold and ultracold atom experiments [98, 115, 121]. In practice, we simply optimize the whole experimental system in a closed-loop manner without reference to our simulations, which exist for proof-of-principle verifications of our interferometric scheme.

Our closed-loop system described in detail in Chapter 6 may be (and is) used to optimize in the presence of inevitable systematics due to nonlinearities in shaking, laser wavefront errors, atom-atom interactions, and finite lattice effects. For example, a closed-loop system can correct deviations from optimal fitness due to the parasitic reflections discussed in the next section as long as the deleterious effects are constant from shot-to-shot. Uncertainties in lattice parameters such as the lattice depth or wavelength due to imperfect lattice alignment or the Gouy phase [22, 49]

may also be corrected for in a closed-loop system.

Other errors can be corrected by the use of techniques common to other atom interferometry schemes. As in light-pulse atom interferometry the effects of unwanted inertial signals, e.g. spurious rotations, can be subtracted out with the use of two interferometers operating in differential mode [83]. The common-mode signal can then be recovered by subtracting the two interferometer measurements.

Trapped-atom interferometers generally suffer from the deleterious effects of phase diffusion due to atom-atom interactions [55]. Heating due to atom-atom interactions is also explored in detail in Ref. [96] In the shaken lattice interferometer one may lower these interactions through sparse population of the lattice, e.g. via the use of the three-dimensional scheme introduced in Sec. 4.1. That is, a deep two-dimensional lattice with low single-site occupation may be used. This results in an array of 1-D low atom number interferometers largely immune to the effects of phase diffusion. Shaking then takes place along the third dimension. If each of the 1-D interferometers is shaken in the same way, a collective measurement of their responses can be made. To counter the effects of shot noise, a few hundred of these 1-D “tubes” can be loaded with about 100 atoms each. Then total atom numbers can reach 10^6 , lowering shot noise to levels comparable with state-of-the-art light-pulse atom interferometers based on Bose-condensed atoms [63].

The use of a shaken lattice to perform atom interferometry can cause issues due to atom decoherence. The effects of decoherence from lattice shaking have been studied extensively elsewhere [4, 54, 68, 122]. The effects of decoherence occur for a certain range of shaking amplitudes and frequencies, depending on the lattice depth and applied acceleration. If the desired dynamic range is known then undesired shaking frequencies and amplitudes can be filtered out accordingly in the learning algorithm, limiting the effects of this decoherence.

4.4.2.2 Robustness of shaken lattice interferometry

In simulation the shaking function is optimized in an ideal situation where the lattice wavelength and depth are known exactly. However, in an experimental setting there is uncertainty

in these values. Thus, the robustness of the shaking function to these errors is of interest. In what follows the variation $D_{f,i}$ between the final state after perturbation and the optimized state is calculated using Eq. (3.4).

The results of variation of lattice depth and wavelength are shown in Fig. 4.9. Changes in the lattice depth of 5% and variations in the wavelength of 0.68% maintain the variation between the perturbed final state and the optimized state to within 1%.

A common limit to the lifetime of lattice experiments is defined by photon scattering as detailed in Sec. 2.2. In other optical lattice experiments long atom lifetimes in the lattice are enabled by servo systems that lower laser noise [12]. If the lattice light is sufficiently far off of resonance the limit due to photon scattering is reduced (although more power is needed). Therefore, interrogation times on the order of tens of seconds are possible in shaken lattice systems. In our retro-reflecting lattice scheme laser phase noise is irrelevant, but unwanted motion of the retro-reflecting mirror will cause unwanted shaking and give rise to spurious signals. Thus any noise in the retro mirror motion must be stabilized via a servo system.

For a mechanically stable system lattice depth variations can be controlled by servoing the laser intensity, and as described above, such a servo is desirable to limit heating due to laser intensity noise. In our experiment wavelength drift is controlled by locking the laser to an atomic hyperfine transition, providing frequency stability to better than 1 MHz. Therefore the wavelength stability provided by laser locking is sufficient, as this is well within the limits shown in Fig. 4.9(b).

Even if the lattice parameters are known exactly, in an experimental system the shaking will have some associated noise. For example, noise in the electronics used to shake will be added to the desired shaking function. Undesired shaking due to mechanical instability will also add noise to the optimal shaking function. To simulate this, Gaussian white noise of varying amplitudes was added to the optimized splitting function. The resulting variation in the perturbed final state relative to the optimized state was recorded and is shown in Fig. 4.10. At a threshold of about 10%, the variation rises rapidly as more noise is added to the shaking function. Thus, in a practical situation the noise in the shaking function should be kept below this threshold.




Figure 4.9: Percent variation of the optimized splitting shaking function shown in Fig. 4.8 after variations of (a) the simulated lattice depth and (b) the simulated lattice wavelength. A variation of 1% is marked in each plot with a blue dashed line.




Figure 4.10: Variation of the final state with varying noise amplitudes added to the optimized splitting shaking function. The results shown here are the average of 5 runs with random white Gaussian noise added to the shaking function $\phi(t)$. Error bars give the standard deviation of the variations. Noise amplitude is given as a fraction of the maximum of the $\phi(t)$ shown in Fig. 4.8. The red line marks a variation of 1%.

Finally, we calculate the robustness of the shaking function to undesired stray reflections in the lattice system. In Ref. [22] the forces due to these reflections caused undesired phase shifts and imposed a limit on the contrast of the interferometer fringes. In our case such undesired reflections set up secondary lattice potentials that can cause the final state to deviate from the optimized result. We simulate this by shaking the lattice and adding a parasitic potential of the form

$$V_P(\epsilon, \delta) = -\frac{\epsilon V_0}{2} \cos(2k_L x + \delta). \quad (4.5)$$

In Eq. (4.5), ϵ is the reflection amplitude and δ is the phase of the parasitic lattice relative to the initial unshaken main lattice potential. The value of δ is generally unknown in a real experiment. The simulation results are shown in Fig. 4.11 for varying values of ϵ from 0.1% (i.e. reflection from an anti-reflection coated window) to 4% (reflection from uncoated glass). A single 4% reflection can change the variation between the final state and optimal state by as much as 1%, and multiple such reflections should be managed carefully. Therefore, in an experimental realization, it is important to use AR coatings on the windows of the science chamber and align optics so that parasitic lattices do not interfere with the main lattice. Our science chamber, discussed in Chapter 5, is AR-coated for our cooling/trapping and lattice wavelengths. It is important to note, however, that as long as the spurious reflections are the same from shot-to-shot, the closed-loop system will adjust its shaking function to achieve the desired final state.

4.4.3 Transfer function optimization

One of the advantages of shaken lattice interferometry is the ability to control the transfer function of the interferometer. That is, we can change the shaking protocols to control the range and type of signals to which shaken lattice interferometry is sensitive, such as an AC acceleration signal. Conversely, we can also design the interferometer to reject certain signals. In this section we give two examples of transfer function control with shaken lattice interferometry. Section 4.4.3.1 discusses optimizing the interferometer in the presence of a DC bias signal in order to measure perturbations on this bias, and Section 4.4.3.2 discusses building a reciprocal interferometer to




Figure 4.11: Variation of the final split state after the addition of spurious lattice potentials with varying phase due to unwanted reflections with reflection amplitudes of $\epsilon = 0.1\%$ (red), 1% (blue), and 4% (cyan).

measure sinusoidal signals. With sufficient interrogation time and the combination of these two methods, one can fully specify the interferometer transfer function and optimize it to the signal of interest. For example, a seismic signal could be detected while rejecting a DC signal due to gravity. An interferometer optimized in the presence of gravity at one spatial location could also be used to perform gravity gradiometry.

4.4.3.1 Optimizing in the presence of a DC bias

To show that we can optimize the interferometer to measure perturbations on a DC signal, the interferometer may be optimized with a DC bias so that it rejects a DC signal of a given magnitude. To obtain the results shown in Fig. 4.12 the standard non-reciprocal interferometer was optimized in the presence of a DC acceleration $a_{\text{DC}} = 0.76 \text{ m/s}^2$. Then the response of the interferometer was simulated for accelerations around a_{DC} . Figure 4.12 shows that the interferometer may be optimized and operated around a DC bias point.

It is straightforward to improve these results. A longer interrogation time will narrow the “dip” shown in Fig. 4.12, improving the interferometer response to signals around the bias. Currently, the only limits to total interrogation time are due to computation time and practical limits are due only to the experimental limitations discussed in Sec. 4.4.2.

4.4.3.2 Optimizing the reciprocal interferometer for AC sensitivity

To be sensitive to AC signals, the interferometer may be set up in a reciprocal configuration as shown in Fig. 4.13. The difference between this and the interferometer configuration shown in Fig. 4.2 is that the atoms take a fully symmetric path. Because the space-time area of the interferometer is zero [73] the interferometer should be immune to DC accelerations but maximally sensitive to a sinusoidal acceleration

$$\vec{a}(t) = a_x \sin(\omega t) \hat{x} \quad (4.6)$$




Figure 4.12: The response of an interferometer optimized in the presence of a DC bias acceleration $a_{\text{DC}} = 0.76 \text{ m/s}^2$. The interferometer response is clearly minimized in the vicinity of a_{DC} and increases away from this bias. This shows that shaken lattice interferometry can be used to reject a DC bias of a given magnitude or measure perturbations around this bias.

at $\omega = 2\pi/T_1$. Conversely, the non-reciprocal Michelson interferometer should be less sensitive to this acceleration but responsive to DC accelerations.

To show this the interferometer was simulated in the following ways: First, the standard Michelson interferometer was optimized so that the atom wavefunction was split, propagated for a time $2T_P$, reflected, reverse propagated for a time $2T_P$, then recombined into the ground state, as in Fig. 4.2. The second simulation of the reciprocal interferometer (Fig. 4.13) split the wavefunction, propagated for a time T_P , reflected the atoms, reverse propagated them for a time $2T_P$, reflected the atoms again, propagated them for a final time T_P , then recombined them back into the ground state. In both cases the total propagation time was $4T_P$. Once the simulations were completed, the shaking function $\phi_{\text{opt}}(t)$ was used to simulate propagation of the TDSE with the potential

$$V(x, t) = -\frac{V_0}{2} \cos \{2k_L x + \phi_{\text{opt}}(t)\} + ma_x x \sin(\omega t). \quad (4.7)$$

Note that any potential term linear in x that is added to a lattice potential may be unitarily transformed into phase factor modifying the lattice potential (and vice versa) [112]. Thus, from Eq. (2.19), we expect that the momenta remain quantized even in the presence of an applied force, and this is verified by our simulations.

The phase of the acceleration during interrogation matters, and here, we analyze the effects of sinusoidal accelerations with no added phase. For the proof-of-principle simulations done here we set $T_P = 0.502$ ms, $a_x = 0.115$ m/s², and scan the acceleration frequency from DC to $f = \omega/2\pi \approx 10/T_P$ as shown in Fig. 4.14. We expect that the sensitivity of the reciprocal interferometer is maximal around $f = 1/4T_P = 0.5$ kHz. Figure 4.14 shows a maximum sensitivity for the reciprocal interferometer around $f_{\text{max}} \approx 7$ kHz. Moreover the reciprocal interferometer shows a factor of 20 enhancement in sensitivity over the non-reciprocal interferometer.

The discrepancy between simulation and theory in the value of f_{max} is due to the fact that the splitting, reflection, and recombination times are on the order of the propagation time, not negligibly small as assumed. Thus, the atoms interact with the shaken lattice for much longer than assumed, increasing the frequency at which the interferometer is maximally sensitive. Better




Figure 4.13: The reciprocal interferometer sequence. The reciprocal interferometer modifies the standard Michelson interferometer sequence shown in Fig. 4.2 so that the atoms travel a fully symmetric path. This configuration is designed to be sensitive to AC accelerations and immune to DC accelerations.




Figure 4.14: Response of the reciprocal (red, solid) and non-reciprocal (blue, dashed) interferometers to a sinusoidal signal as in Eq. (4.7). The response is given in terms of the variation $D_{f,i}$ between the optimized final state of the interferometer and the final state after shaking with the applied signal as in Eq. (3.4). The total interrogation time of each interferometer is 2.008 ms and the amplitude of the applied acceleration signal is $a_0 = 0.115 \text{ m/s}^2$. The reciprocal interferometer is 20 times more sensitive than the non-reciprocal interferometer to a signal with $f \approx 7 \text{ kHz}$, showing that shaken lattice interferometry may be modified to tailor the interferometer response to an AC signal.

agreement can be reached by increasing the propagation time relative to the splitting, reflection, and recombination times. Furthermore, as in the DC-biased interferometer, longer interrogation times will narrow the “spike” around f_{\max} . In the case shown here, having the atoms traverse multiple periods of the reciprocal interferometer (i.e. repeating the reciprocal sequence twice so that the net spacetime area remains zero but the atoms are reflected four times instead of two) will narrow this spike and increase the interferometer sensitivity to signals at frequencies around f_{\max} . The interferometer sensitivity is calculated in Sec. 4.4.4.

4.4.4 Calculating the interferometer sensitivity

In this section we calculate the interferometer sensitivity using the Fisher information metric [48]. Because the envelope function allows for smooth turn-on and turn-off of the shaking function, we can “stitch” together successive propagation steps to increase the interrogation time. Therefore the dynamic range of the interferometer can be controlled by changing the total propagation time. The GA is used to force the atom wavefunction to maintain its state over longer propagation times, correcting for any errors that arise as the propagation protocol is repeatedly applied.

In other atom interferometry schemes a phase is measured between two atomic wavepackets. However, in shaken lattice interferometry an applied acceleration will change the atoms’ momenta as they interact with the shaken lattice, as shown in Sec. 4.3.2. Therefore, because the atoms’ momentum population changes under the influence of an applied signal, the definition of a phase difference between two arms becomes ambiguous because of the ambiguity in defining the arms of the interferometer.

It is due to this ambiguity that we use the classical Fisher information (CFI) to quantify the interferometer sensitivity, given an optimized state \vec{P}_f and the final state under acceleration \vec{P}_a . The CFI for a measured parameter θ given a probability distribution $f(x, \theta)$ can generally be written as

$$F_C(\theta) = \int dx \left(\frac{\partial}{\partial \theta} \ln(f(x, \theta)) \right)^2 f(x, \theta). \quad (4.8)$$

We may then use the Cramer-Rao bound to find the smallest resolvable change in the variable θ .

The Cramer-Rao bound is written

$$\delta\theta = \frac{1}{\sqrt{F_C}}. \quad (4.9)$$

In simulation we have access to the full atom wavefunction. Thus, we could calculate the quantum Fisher information [48] for the acceleration parameter a . However, experimentally we only have access to amplitude information and lose information about relative phases between the different interfering momentum states. Therefore, in this work we only use the CFI. To simplify the problem, we will re-write Eq. (4.8) using the momentum population vector \vec{P}_a . In this case, we can write

$$F_{C,P}(a) = N_{\text{at}} \sum_{n=-N}^N \frac{(\partial P_{a,n}/\partial a)^2}{P_{a,n}} = N_{\text{at}}(\vec{A} \cdot \vec{B}) \quad (4.10)$$

where \vec{A} has components $A_n = 1/P_{a,n}$ and \vec{B} has components $B_n = (\partial P_{a,n}/\partial a)^2$ and N_{at} is the total number of atoms. The sum extends from $-N$ to N where $N = 5$ is where we truncate the number of momentum states considered, as higher-order momentum states are negligibly populated. The factor N_{at} arises because if the atoms are non-interacting, each atom counted is a separate measurement of the probability distribution \vec{P}_a . This will give us a factor of $\sqrt{N_{\text{at}}}$ in the denominator of the expression for δa , as expected from conventional interferometry. The CFI $F_{C,P}$ increases as $\partial P_{a,n}/\partial a$ increases. This can be intuitively understood because operating in a regime with high $\partial P_{a,n}/\partial a$ is analogous to operating on the edge of a fringe in conventional interferometry. From Eq. (4.10), we can immediately write down the smallest resolvable acceleration δa using Eq. (4.9)

$$\delta a = \frac{1}{\sqrt{N_{\text{at}}}} \frac{1}{\sqrt{\vec{A} \cdot \vec{B}}}. \quad (4.11)$$

To study how δa changes with interrogation time T_1 , we optimized the full interferometer for varying interrogation times from approximately 1 to 20 ms. Once the optimization was complete, we took the optimized shaking function $\phi_{\text{opt}}(t)$ and added an acceleration signal, altering the potential term in Eq. (2.16) to

$$V(x, t) = -\frac{V_0}{2} \cos \{2k_L x + \phi_{\text{opt}}(t)\} + ma_x x. \quad (4.12)$$

We then solved the TDSE with this potential and recorded the final momentum state vector \vec{P}_a

for various values of the acceleration a_x . From this we use the Cramer-Rao bound to get $\delta a(T_I)$, where the derivatives were taken around $a = 0$.

The value of δa should decrease as T_I increases such that $\delta a \propto T_I^{-n}$ where n is the interferometer scaling, since a lower value of δa corresponds to an interferometer that is more sensitive to acceleration. To obtain n , we use the Levenberg-Marquardt algorithm to fit the resulting data to a curve of the form $f(T_I) = CT_I^{-n}$ for some constant C . The results are plotted in Figure 4.15. Note that the total interrogation time also includes the nonzero splitting, reversal, and recombination time (about 1.5 ms of total shaking time). The fit gives $n = 2.21 \pm 0.31$. This is consistent with the $n = 2$ scaling achieved in most other atom interferometers (e.g. light-pulse atom interferometers [20, 29, 47, 58, 75]).




Figure 4.15: The minimum detectable acceleration δa scaled by $\sqrt{N_a}$ plotted on a log scale versus the interrogation time T_I . The black points are simulation results, and the red line is a fit of the form CT_I^{-n} . Here, lower values of δa correspond to a more sensitive interferometer.

Using the results above for the sensitivity scaling with T_I we can plot the acceleration sensitivity of the interferometer versus the atom number for various interrogation times. This is shown in Fig. 4.16. From this plot, we can expect a sensitivity better than $10^{-11}g$ for interrogation times of 1 s and 10^6 atoms.




Figure 4.16: The acceleration sensitivity δa (relative to g) for varying atom numbers N_a at $T_I = 10$ ms (red), 100 ms (blue), and 1 s (black).

It is possible that other shaken lattice interferometer configurations could scale with higher powers in T_I . For example, one could optimize the interferometer to accelerate the atoms in the lattice as they propagate, much like the continuous-acceleration-Bloch interferometers that scale as T_I^3 [74]. One could also alter the fitness function so that the GA now minimizes the Cramer-Rao bound in Eq. (4.11) and maximizes the interferometer sensitivity.

4.5 Next steps

There are a wide variety of things that one can explore further in simulation. Simplified shaking protocols for atom wavefunction splitting are worth exploring, and we address this in Chapter 7. Some aspects of shaken lattice interferometry that are not treated here are listed below.

We consider the Michelson and reciprocal interferometer schemes in this work as they are natural and familiar starting points. As stated in Sec. 4.4.4, in general it is possible to optimize the interferometer to obtain higher sensitivities by changing the shaking protocol, and optimization using the Fisher information in Eq. (4.11) as a fitness or error metric is an interesting prospect for future work. This could allow for optimization of a shaken lattice interferometer with sensitivity scaling that is faster than the current quadratic scaling.

One possible next step would be to include atom-atom interactions in our simulations. To first-order, we could include the changes to the effective lattice depth described in Sec. 4.1 quite straightforwardly. Full inclusion of the nonlinear term in the Gross-Pitaevskii equation [90] would be reasonably straightforward in the genetic algorithm, but nonlinear terms in the Krotov method must be treated differently [104].

Another possibility would be to improve the transfer function optimization presented here. First, one can work on increasing the DC bias so that gravity gradiometry can be performed with shaken lattice interferometry. Similarly the AC frequency f_{\max} of maximum reciprocal interferometer sensitivity may also be decreased to move into the regime of seismic signals or other signals of interest. This will require an increase in the interferometer interrogation time. As stated in Sec. 4.4.3 an increase in the interferometer interrogation time will increase the interferometer sensitivity

in all cases, narrowing the “spikes” in sensitivity shown in Figs. 4.12 and 4.14.

A multi-dimensional lattice opens up a myriad of possibilities. Interrogating the atoms by shaking along each of the three Cartesian dimensions will allow us to build a three-axis accelerometer, and ideally we could measure along each of the three axes simultaneously. Furthermore, it should be possible to make a shaken lattice gyroscope. In order to do this we can use a two-dimensional Sagnac scheme where the atoms move around an enclosed area as the lattice is shaken along two dimensions. This setup should be sensitive to rotations. In this way it is potentially possible to build a six-axis inertial sensor (capable of measuring acceleration and rotation along each of the three Cartesian axes) using shaken lattice interferometry.

Chapter 5

Building a shaken lattice interferometer

This chapter will describe the basic experimental setup used to produce the results presented in this thesis. The layout of this chapter is as follows: First, we will describe briefly the experimental sequence, listing applicable references to the interested reader who wants to learn more. Some of the experimental processes, e.g. magnetic trapping and Bose-Einstein condensate (BEC) production are described thoroughly and sufficiently in other references and will be only briefly treated here. The focus of this chapter will be on the practical implementation of the shaken optical lattice.

5.1 From dispenser to BEC on a chip

The purpose of this section is not to describe every step of the experiment in excruciating theoretical detail, but rather to describe the relevant pieces critical to shaken lattice interferometry. Citations in the section will lead the reader to more in-depth discussion of the experimental sequence. In particular we will try to highlight the citations most helpful in learning the systems in use in the Anderson lab. As such, some of the nitty-gritty experimental details regarding lattice loading will be provided here as a practical reference, but trapping and cooling of atoms will be explained only cursorily. Specifically, to obtain BEC we follow generally the experiments explained in detail in Refs. [53, 103, 105].

5.1.1 The laser system

Laser systems of the type used in this work are detailed exhaustively in previous work, and we will mention only relevant details of this specific experiment. The master lasers used for trapping, cooling, pumping, and imaging are the JILA-style external-cavity diode lasers (ECDLs) developed originally by Carl Wieman’s group [70]. The cooling light is fed into a tapered amplifier (New Focus TA7613, free-space output) then split into two fibers containing roughly 80 mW of power apiece. The light from the low-power repump master laser is injected into a higher-powered slave diode. The light from the slave laser is then cross-polarized with the cooling light and coupled into the same fibers. A third ECDL provides light for optical pumping (the “pump” light) and absorption imaging (the “probe” light). In order to image along two orthogonal axes, we couple our probe light into two separate fibers for the two- and three-dimensional MOTs. In this experiment, all fibers are single-mode polarization maintaining fibers, ensuring that light with a nice Gaussian mode and a well-defined linear polarization comes out of the fiber.

A schematic of the lattice system is shown in Fig. 5.1 and described here. Part of the light that exits the laser is picked off by a polarizing beamsplitter (PBS) and used in a saturated absorption spectroscopy setup to lock the laser to the $|F = 4\rangle \rightarrow |F' = 4/5\rangle$ crossover transition in cesium.¹ The rest of the light is coupled into a fiber going into a TA (New Focus TA7616, fiber output). The output of the TA is isolated by a Faraday isolator and passes through a polarizer. The polarizer stabilizes the polarization of the light coming into the AOM, and fluctuations in the laser polarization are converted into intensity fluctuations. The resulting intensity fluctuations are then suppressed by the intensity servo. The AOM acts as a high-speed switch for the light as well as allowing us to servo the laser intensity. After the AOM one of the first-order beams is used for the lattice light.² The remaining light in the zero-order beam is blocked. A wedged plate beamsplitter (BS) picks off about 4% of the shifted light and sends it to a photodetector (PD,

¹ The actual locking transition does not matter, as long as it is consistent from day-to-day.

² It doesn’t matter which of the two first-order beams is used, as the AOM is used as a fast shutter and an intensity servo. The AOM frequency does not matter, but day-to-day consistency in the frequency shift induced by the AOM does matter. That is, pick a frequency and stick with it.

Thorlabs PDA36A), which is then fed into one of the inputs of a JILA-built PID intensity servo. The laser intensity is stabilized through the use of a mixer (MiniCircuits ZX05-1LHW) that allows us to rapidly modulate the amplitude of the RF tone fed into the AOM. The remaining 96% of the light is then coupled into a fiber and sent to the main experiment; care is taken to ensure the polarization stability of the output light, as polarization fluctuations after the fiber output are converted into intensity fluctuations in the main experiment (see Fig. 5.9). These intensity fluctuations are not servoed in our current setup.

5.1.2 The physics package and vacuum cell

The details of the basic experimental system can be found in Ref. [53] and are described briefly here. The double-chambered vacuum cell shown in Fig. 5.2 is the “standard” vacuum + atom chip system used in many of the experiments done in our lab. Variations on this system have been used widely in our group to create BEC. The cell is attached to a larger vacuum chamber via a copper pinch-off tube and baked at 300°C to achieve the low pressures needed for BEC experiments. Once baking is complete, the cell is cleanly separated from the system using the steady hands of a few good graduate students and a set of hydraulic jaws that closes on the copper tube. After a few days of “soft baking” at 120°C, the system is ready to be installed.

In the vacuum cell the bottom chamber (the “2D chamber”) contains a non-evaporable getter (NEG) and a rubidium dispenser. Our rubidium source is a commercial SAES dispenser through which we run about 2.6 – 2.8 A of current.³ A pinhole between the 2D and 3D chamber (top chamber) provides the differential pumping needed to isolate the higher pressure 2D section from the 3D. A picture of the pinhole is shown in Fig. 5.3. A 0.5 L/s ion pump maintains the high vacuum quality ($P \approx 10^{-9}$ torr) in the 3D chamber. A second NEG installed in the spherical cube separating the 2D and 3D chambers also helps maintain vacuum quality. The 3D chamber is capped off at the top by a silicon atom chip. A silicon disk is epoxied to the ambient side of the

³ The optimum current at which to run the dispenser varies from system to system and is largely based on the vacuum quality in the cell. Due to a leak that occurred in our cell during bakeout, the dispenser in this system is run at a relatively low current. Other experiments in our lab that use similar vacuum systems run the Rb dispenser currents as high as 3.4 A.

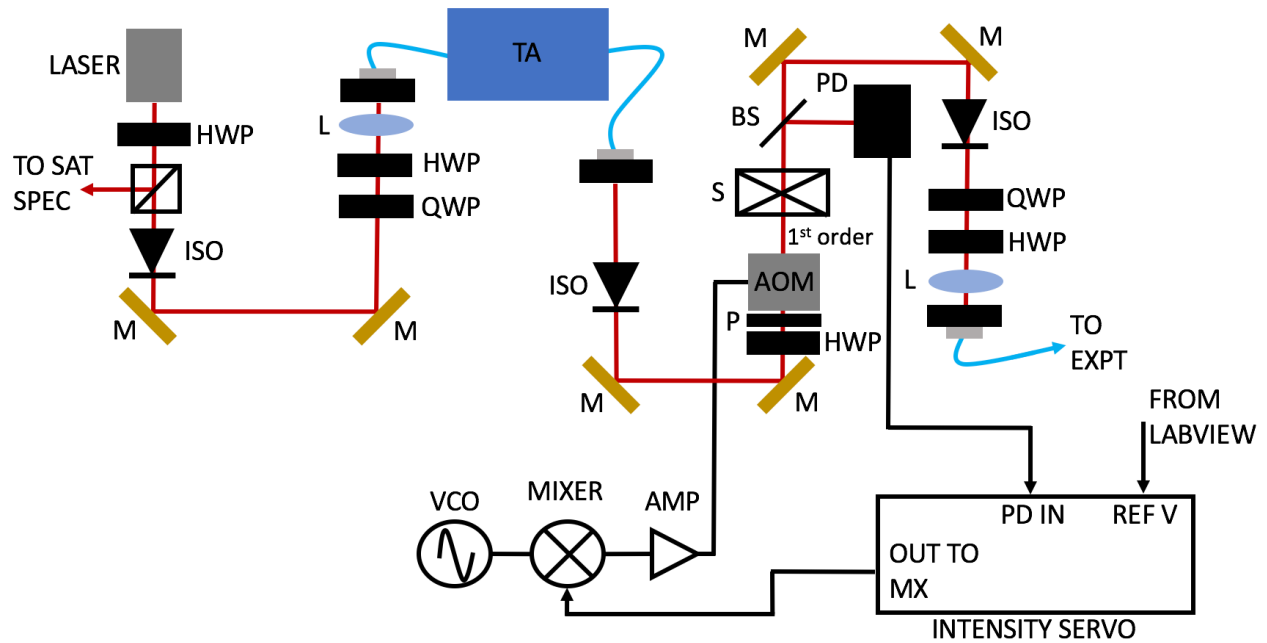


Figure 5.1: Schematic of the lattice laser system with servo. The angle of the first-order AOM output is not shown for ease of drawing. The servo works by comparing a control voltage from LabVIEW to the voltage output of the photodetector. The error signal is sent to a mixer that attenuates the output of a voltage-controlled oscillator (VCO) that controls the AOM frequency. The attenuated signal is then amplified by a 2 W RF amplifier and sent to the AOM. The symbols represent the following: lens (L), mirror (M), shutter (S), optical Faraday isolator (ISO), half-wave plate (HWP), and quarter-wave plate (QWP). Other symbols are defined in the main text.

chip before vacuum is pulled on the chip to minimize chip deflection due to vacuum [105].

Currents run through the wires patterned on the chip produce magnetic trapping potentials in which the atoms are evaporated to BEC. Wires patterned on the ambient side of the chip are gold, while copper wires are patterned on the vacuum side of the chip.⁴ High-conductivity silicon “vias” surrounded by Pyrex allow for currents to be transported from the ambient side of the chip to the vacuum side. The chip used in this work did not use copper wires or vias, but other chips used in the group have [15, 16, 17, 100] and the technology is mentioned here for completeness.

The 3D chamber is surrounded by a coil structure containing the anti-Helmholtz MOT gradient coils and three Helmholtz coils providing uniform bias fields in all three Cartesian directions. Four permanent magnets installed around the 2D cell provide a two-dimensional anti-Helmholtz field for 2D MOT production. All of the optics, cameras, and coils needed to make BEC are contained in a multi-tiered “physics package” made by ColdQuanta, Inc. (Fig. 5.4). The vacuum cell fits in the center of this package, and the magnetic coils are installed by sliding the coil assembly down rails surrounding the cell. The cell is bolted to the second-highest level of the package, and the permanent magnets are fixed to the underside of this level.

The connections to the chip are made by gluing a 14-pin Mill-Max connector (P/N 850-10-050-20-001000) to the gold pads on the edge of the ambient side of the chip using conductive epoxy (CircuitWorks[®] CW2400 Conductive Epoxy). This epoxy is used because it bonds the connector robustly to the chip (but not so robustly that the chip breaks before the connector pops off), does not appear to warp the chip (which could break vacuum), and is easily removed with methanol. Upon removal of the connector and any residual epoxy, no noticeable damage is done to the chip. Thus, mistakes are easily corrected and if damaged, the chip connector can be fixed without pulling the cell completely out of the physics package. In practice the chip connector is cured at room temperature for 24 hours before the external connector is applied.⁵ This connector

⁴ This arises because gold and rubidium do not play well together. The two metals form an alloy, and upon heating this alloy (e.g. by running current through wires on an atom chip), rubidium vapor is released. This can increase the background pressure in the vacuum cell and thus is detrimental to BEC production. See Refs. [99, 105] for more details.

⁵ The author could not find the part number for the female chip connector, but any 14-pin connector that fits in




Figure 5.2: The two-chamber vacuum system used in the experiments described in this thesis. In our system the cell is oriented so that the atom chip is on top. The 2D MOT chamber contains a rubidium dispenser and an NEG. The atoms are pushed up through a pinhole between the 2D cell and the spherical cube (yes, that's actually what it's called) and into the 3D cell. A 0.5 L/s ion pump pumps on the 3D cell, which is topped by the atom chip (the pinhole also serves as a differential pumping stage). Copper wires patterned on the vacuum side of the atom chip are barely visible through the 3D cell. While the 2D cell is not AR-coated, the 3D cell is. The copper pinch-off is also labeled; this separates the vacuum chamber from the larger bake-out station.




Figure 5.3: A view from the underside of the 2D cell, up through the aperture on the cell bottom and focused on the pinhole. A light is being shone through the pinhole to backlight it for clarity. The 0.5 L/s ion pump can be seen off to the right, and the copper pinch-off is to the left. The push beam that moves the atoms from the 2D cell/MOT to the 3D cell/MOT travels through the aperture in the bottom of the cell and through the pinhole. The metallized pads on the underside of the cell (surrounding the aperture) make electrical connections to the rubidium dispenser and NEG. Connections are made to the pads using conductive tape and colloidal silver.




Figure 5.4: (a) CAD and (b) a photograph of the physics package provided by ColdQuanta. The physics package is roughly 28x28x32 cm in size, and a cell similar to that shown in Fig. 5.2 is bolted down in the center of the package. The lower two tiers contain the optics needed to produce a 2D MOT and the push beam that pushes atoms up through the pinhole to the 3D cell. The upper two tiers contain the optics for 3D MOT production, optical pumping, and probing. In the CAD drawing, the cooling beams are shown in red, and the imaging beam is gray. In the photograph, the system is rotated clockwise 90° from the CAD view. The coils surrounding the cell in the CAD drawing are not installed in the photograph, so the photo shows the 3D cell poking above the top tier of the physics package in the center. The cable connecting the 0.5 L/s ion pump to its controller is shown running off to the right.

runs up through a space in the coil assembly and attaches to a simple breakout board above the coil assembly. From this breakout board one can make the appropriate connections to the chip drivers.

5.1.3 BEC production in the experimental system

Once the cell is installed and the chip is connected, we can begin BEC production. Digital and analog control of the experimental elements (e.g. digital signals sent to the optical shutters or analog signals sent to the magnetic coil drivers) is done via a LabVIEW program running on the experimental control computer. This control program generates a timetable for each relevant channel and sends it to an NI PXI 8186 control chassis. The chassis then communicates with a set of NI compact reconfigurable input/output (cRIO) cards that output the relevant analog and digital signals. The temporal resolution of the experimental control is $\approx 20 \mu\text{s}$. More information on the control system is given in Chapter 6.

Atoms are collected in a cigar-shaped two-dimensional MOT then pushed up through the pinhole in the silicon barrier separating the 2D and 3D chambers by a “push beam” directed up through the bottom of the 2D chamber. The push beam consists of a couple mW of cooling and repump light sampled from the 2D MOT fiber.

When the atoms are moved into the 3D chamber, they are captured in a three-dimensional MOT. Once the MOT is loaded, we detune the MOT lasers (from $\approx 2 - 3 \Gamma_{D2}$ to $\approx 4 - 5 \Gamma_{D2}$) and increase the current in the MOT coils by roughly a factor of 2. This compressed MOT (CMOT) increases the density of the atoms and reduces heating due to spontaneous emission. From this compressed MOT we detune the lasers even further ($\approx 8 - 9 \Gamma_{D2}$) and perform polarization-gradient cooling (PGC) with our MOT beams. This cools the atoms below the Doppler limit; practically we reach temperatures of about $20 - 40 \mu\text{K}$, roughly an order of magnitude higher than the theoretical limit of $4 \mu\text{K}$.

the space between the coil mount and the big-Z coil and mates to the male connector will work. The connectors we currently use have flat metal pads on the other side of the connector, and we solder flexible stranded 32-gauge wire (sold by Cooner wire) to the pads and then to the breakout board through-hole terminals.

After sub-Doppler cooling and optically pumping the atoms into the magnetically trappable $|2, 2\rangle$ state, we trap the atoms by running about 20 A of current through a large Z-shaped coil (astutely named the “big-Z” coil) placed above the chip.⁶ This magnetic field, in combination with the bias fields, traps the atoms in a magnetic minimum. By ramping down the currents in the big-Z coil and the bias coils, we can move the atoms to the chip and trap them using the magnetic fields produced by the chip wires and the bias coils. In the current experimental system we employ a “throw and catch” scheme to transfer the atoms from the big-Z trap to the chip trap. Directly after catching the atoms in the initial trap, the bias fields are increased to compress the chip trap and increase the trapping frequency.⁷ This increases the atom rethermalization rate and allows us to more efficiently produce a BEC during forced RF evaporation.

Similar chips to those used in this experiment have been used to make BEC in a compact, transportable system [38, 53]. The basic mathematics underlying atom chip traps is well-summarized in Refs. [53, 103] and we will only give details specific to the experiment described in this work. In this experiment the atom chip trap consists of two wires that cross one another at the chip center, a thicker “main” wire and a thinner “dimple” wire, as shown in Fig. 5.5. By running independent currents through these wires (3.25 and 1.25 A in the main and dimple wires, respectively) and manipulating the bias fields, we trap atoms in a very tight trap roughly 250 μm from the chip surface (see Fig. 5.6). We calculate our mean trapping frequencies to be $\bar{\omega} \approx 2\pi \times 110$ Hz in the initial trap and $\bar{\omega} \approx 2\pi \times 260$ Hz in the compressed trap. Our atom chip current drivers are identical to the ones used in Ref. [15] and each of the two channels runs on two 12 V sealed lead-acid batteries for low-noise operation. Electronic and mechanical fuses protect the chip from being damaged due to excessive currents. Mechanical fuses protect the coils from damage due to undesired overcurrents.⁸

After the atoms are trapped on the chip, we can use the well-established technique of forced

⁶ Do not run 20 A through this coil for longer than a few hundred ms unless you want to fry the coil. This is an unpleasant experience.

⁷ The first stage of RF evaporation takes place during this compression, but this initial cut is not affected appreciably by the changing trap parameters.

⁸ Don’t remove these fuses. You’ll regret your decision.




Figure 5.5: Picture of the atom chip used in the experiments described in this work. The chip is made of high-conductivity silicon coated in an insulating dielectric and measures roughly 2.5 cm by 2.5 cm by $420 \mu\text{m}$ thick. All wires are patterned in gold on the ambient side of the chip. The only two wires used here are the thicker main wire (labeled by the current I_0) and the thinner dimple wire (I_d); these wires intersect in the center of the chip and the atoms are trapped below this intersection.




Figure 5.6: Density plots of the magnetic trap potential used to hold the atoms in (a) the initial chip trapping configuration, (b) the compressed evaporation trap, and (c) the waveguide-like decompressed pre-lattice trap. The color indicates the magnetic field in Gauss. The calculations were made using the Biot-Savart law and assume infinitesimally thin wires. Note the difference in scale in plot (c) relative to (a) and (b).

RF evaporation to evaporate them to degeneracy. A single loop antenna placed just above the chip surface radiates the RF field used to selectively remove the hot atoms from the trap; the cold atoms then rethermalize and the temperature of the atom cloud drops. This process continues until the atoms are Bose-condensed. A typical BEC in our system contains a few 10^4 atoms. Figure 5.7 plots the currents through the coils and chip wires during the experimental sequence from magnetic capture with the big-Z coil to BEC. Note that we define our axes such that the lattice light propagates along the x direction and z is the direction of gravity.

5.2 Building an optical lattice around a BEC system

When the atoms are trapped on the chip and Bose-condensed, we decompress the magnetic trap by lowering the bias fields and main wire current and turning off the dimple wire completely. This produces a waveguide-like trap [103] with little confinement in the longitudinal direction, as shown in Fig. 5.6(c). This is done to better mode-match the magnetic trap to the lattice trap. We then switch off the magnetic trap potential and ramp on the lattice potential over 0.6 ms, effectively “catching” the atoms in the lattice trap. More details on lattice loading will be given in Chapter 6. Due to the Gaussian profile of the lattice beam, the radial confinement along the z -direction holds the atoms against gravity. Figure 5.8 plots how the relevant experimental parameters change during the transfer of the atoms from the BEC trap to the lattice trap.

Our lattice beam has a waist at the focus of $w_L = 40 \mu\text{m}$, and our typical trapping powers are $P_L \approx 10 \text{ mW}$. To account for losses from the electro-optic modulator, we assume that only 30% of the incident lattice light is reflected back onto itself. This gives a lattice depth of $V_0 \approx 20E_R$ and a photon scattering rate of $\Gamma_{\text{sc}} = 0.64 \text{ photons/second}$. We estimate the lattice depth by counting atoms in each momentum state and fitting the relative atom number to the predicted atom number for a Bloch state with a given depth. A more precise determination of the lattice depth is described in Chapter 7.

We build our lattice system around the BEC system by placing a 1” (25.4 mm) thick aluminum breadboard around the physics package. This breadboard, shown in Fig. 5.10, was made




Figure 5.7: The relevant currents run through the big-Z coil (black, scaled down by 50%), X-bias coil (red) Y-bias coil (blue), Z-bias coil (gray), main chip wire (magenta), and dimple wire (cyan) during the following experimental stages: (a) optical pumping, (b) big-Z capture and transport, (c) initial chip loading, (d) chip compression and first-stage evaporation, and (e) the rest of the RF evaporation. The optical pumping stage usually lasts for only 0.5 ms and is extended in time in this plot for clarity.




Figure 5.8: Changes in currents (left axis) and lattice light (magenta, right axis) as we ramp from the BEC trap to lattice. The X-bias and Y-bias coil currents are shown in black and red, respectively, and the main and dimple wire currents are shown in blue and gray, respectively. The amplitude of the lattice beam is given in arbitrary units related to the voltage sent to the intensity servo. Due to the fact that the servo has trouble servoing around zero intensity, the lattice beam is held at an arbitrary high value (with the mechanical shutter closed) as the atoms are evaporated to BEC (end of evaporation shown in a), then ramped down to zero as the atoms are loaded into the waveguide trap (b). The lattice light intensity is nearly zero for about 1 ms as the mechanical shutter is opened (c), then it ramps on after the chip wires are ramped off (d). The X-bias coil is not ramped to zero, but rather a constant value that allows the atoms to maintain their spin polarization in the lattice. The reason for this will be explained in more detail in Chapter 6.

in-house and has a hole cut out so that it slides over the physics package. The breadboard is held up by four brass legs of 3" diameter and length 7.5".⁹ Brass was used for stability at the suggestion of the Ye group here in JILA,¹⁰ the owners of the most stable optical lattice-based clock [18]. Previous iterations of the lattice setup used aluminum legs and 1/2" (12.7 mm) thick breadboards; day-to-day lattice stability increased dramatically once the more robust breadboard setup was installed. On this board we set up all of the optics needed to produce our optical lattice as well as a second axis of imaging that can image the atoms along the lattice propagation direction.

For all stages of experimental optimization, we image the atoms using conventional time-of-flight (TOF) absorption imaging (with background subtraction), but for imaging Bose-condensed atoms, we place a doubling tube in the camera lens system to improve our magnification. Our imaging resolution is roughly $4.6 \mu\text{m}/\text{pixel}$ along the first axis of imaging (orthogonal to the lattice beam propagation), and $3.6 \mu\text{m}/\text{pixel}$ along the second axis. These resolutions are calculated by varying the TOF and fitting the atoms' position in free-fall to a parabola, and the measured resolutions are sufficient for our purposes.

The lattice and imaging optical setup is shown schematically in Fig. 5.9, and a photograph of the whole experiment is shown in Fig. 5.10. The lattice light leaves the fiber and is collimated by a lens L_1 . The light then passes through a half-wave plate (HWP) to adjust its polarization for maximum transmission through the polarizing beamsplitters (PBS) placed on both sides of the cell. The light is then focused with a long focal length lens L_2 onto the atoms. After passing through the cell, a short focal length lens L_3 focuses the atoms through an electro-optic modulator (EOM) and onto a retro-reflecting mirror M_R in a cat's eye scheme.¹¹ The light then reflects back on itself, creating the optical lattice trap. The EOM is used to shake the lattice and will be discussed further in Sec. 5.3.

The second axis of imaging (also shown in Fig. 5.9) has proven critical for rapid realignment

⁹ In this thesis, we will use metric units for most quantities. However, most machined parts are fabricated using English units. To avoid confusion in conversion, when the English units give nice round numbers, they will be used.

¹⁰ Special thanks to G.E. Marti for taking the time to show us the Ye clock setup and for making these brilliant suggestions.

¹¹ We find that a cat's eye scheme is more robust against mirror motion. We have found that with this setup, the mirror motion is negligibly small ($\delta\phi \approx 1 \text{ mrad}$) relative to the lattice shaking.

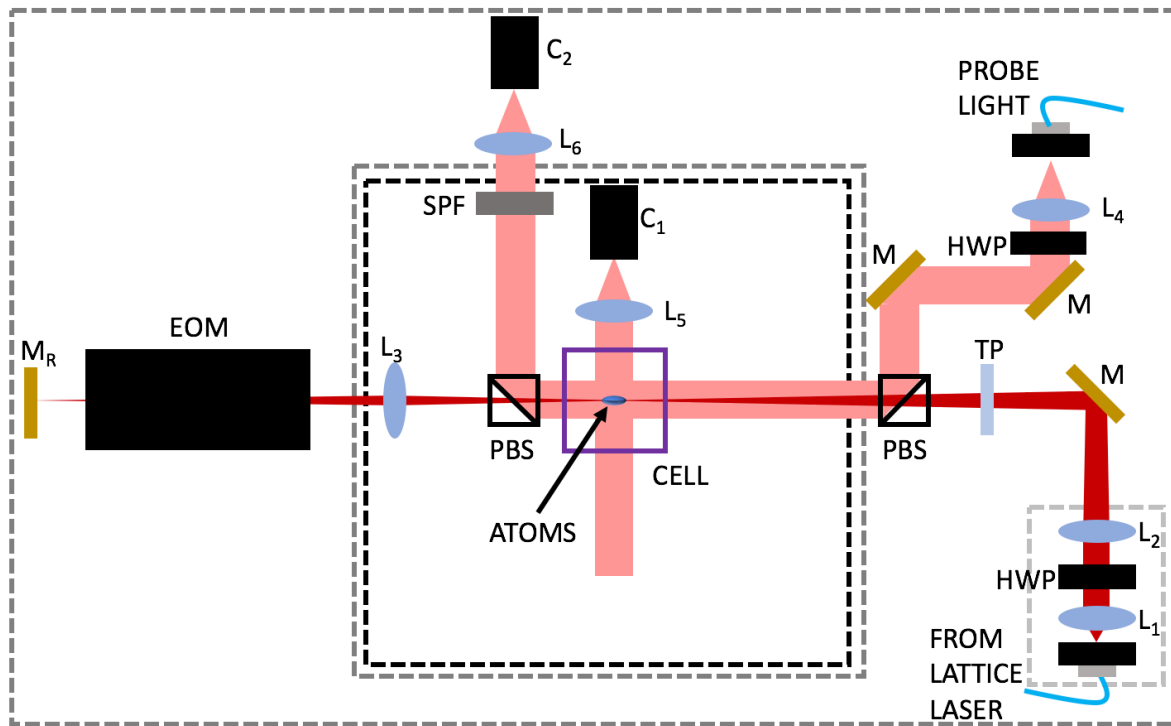


Figure 5.9: Schematic of the optics used in the lattice and imaging systems. The cell is shown outlined in purple, with the atoms (blue) in the center. The probe light for the two axes of imaging are shown in light red, and the lattice light is shown in dark red. Elements surrounded by the black dashed box are contained on the physics package, elements surrounded by the dark gray dashed lines are contained on the secondary lattice board, and the lattice optics surrounded by the light gray box are on a three-axis translation stage. All abbreviations are explained in the text.




Figure 5.10: The entire experiment, including the physics package (black, center) and the auxiliary lattice board surrounding it. (a) View of the experiment showing the lattice launcher (right) and the EOM and lattice retro-reflecting mirror (left) on the auxiliary lattice board. (b) View of the experiment along the direction of lattice propagation showing the lattice launcher (left) and the second axis of imaging launcher (right) on the auxiliary lattice board. The coils used to apply the magnetic field gradient to simulate a signal are shown surrounding the coil assembly (red coils surrounding black coil structure). (inset) A close-up view of the cell surrounded by the coil assembly (coils in orange). The red gradient coil can also be seen to the right of the cell.

of the lattice. Probe light at 780 nm is collimated by a lens L_3 and sent through a half-wave plate to orient its polarization orthogonal to that of the lattice light. Thus, this light is combined with the lattice light on the PBS placed before the cell and picked off by the PBS placed after the cell. The probe light then passes through a short-pass filter (SPF) to filter out most of the residual 852 nm light picked off by the second PBS and is focused onto a CCD camera (C_2 , Basler Scout) identical to that used for the first axis of imaging (C_1).

To align the lattice beam we remove the EOM and retro-reflecting mirror and trap the atoms in a dipole trap. Rough alignment of the lattice beam is done via a three-axis translation stage on which the lattice beam launcher, collimating lens, and focusing lens are mounted. We have constructed an alignment jig consisting of two pinholes that can be attached to the coil mount surrounding the cell; this jig allows us to obtain a rough lattice alignment and ensure that the beam is not angled up or down. We align the lattice to the pinholes on the jig using the translation stage and the angular mirror. The angular mirror is then tilted slightly in the horizontal direction to mitigate the effects of reflections off of the vacuum cell.¹²

The alignment jig ensures that the beam is straight but does not align the beam with the atoms. The beam is then aligned with the atoms using the second axis of imaging and the three-axis translation stage. Small amounts of 852 nm light that leak through the SPF are aligned with the atoms using low TOF values while also monitoring the trapping efficiency along the first axis of imaging. Using this method we quickly and efficiently align the optics and trap the atoms in the dipole trap. The alignment is optimized by catching the atoms in a high-intensity dipole trap, then iteratively lowering the trap intensity and optimizing the beam alignment for longer and longer hold times. For the experiments presented here, we find that hold times of $\approx 10 - 100$ ms are sufficient for alignment purposes.

After the lattice is roughly aligned the second axis imaging beam is blocked, as we find its

¹² Any spurious reflections, if they are constant from shot-to-shot, will simply interfere with the main lattice, changing the lattice depth. The effects of these reflections are difficult to measure, but it is straightforward to measure the lattice depth taking these reflections into account. Regardless, they can change the effect of the shaking and thus should be mitigated.

presence affects the position of the atoms after imaging along the orthogonal direction due to the momentum kick given to the atoms by the presence of the resonant imaging light. Furthermore, after the lattice beam is aligned, we remove the second beamsplitter that was placed after the cell. While the optics used for imaging are reasonably robust against small bumps, care should be taken to avoid bumping the lattice optics, especially the angular mirror directing the lattice beam from the fiber launcher into the cell (see Fig. 5.11).

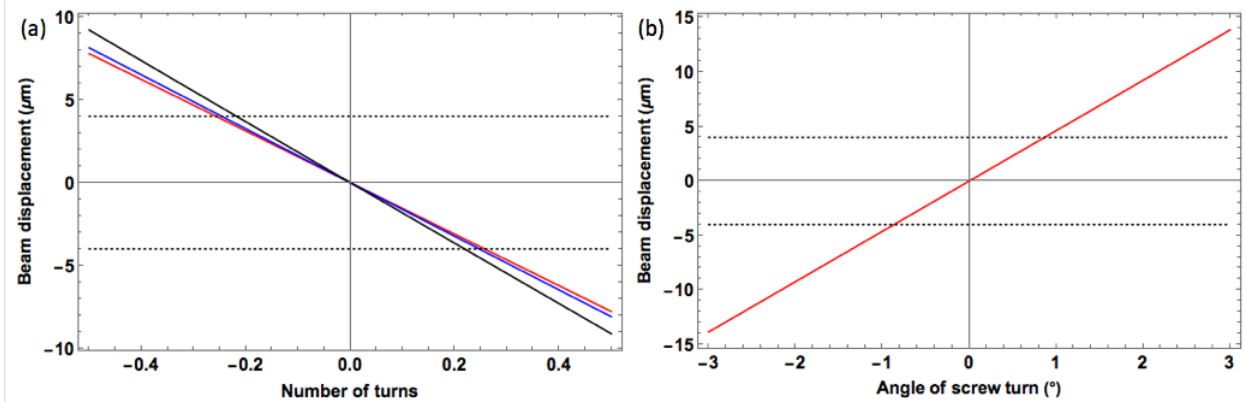


Figure 5.11: Stability of (a) the tweaker plate (TP) and (b) the angular mirror directing the lattice beam into the cell (see Fig. 5.9). The TP displacement is shown for plate angles of 0° (red), 10° (blue), and 20° (black) relative to the plane orthogonal to the incoming laser beam. For maximum resolution (in terms of screw turns per displacement), the TP is mounted near 0° . The number of screw turns is calculated using the different mounts for these two optics: (a) a Newfocus 9873-K mount and (b) a Thorlabs Polaris K1T mount. The angular mirror is extremely sensitive and thus a high-quality mirror mount is used to house this optic. In both plots the experimental image resolution is shown with black dotted lines.

Once the alignment is optimized in the dipole trap, we re-install the retro mirror to make our lattice beam reflect back on itself. A hallmark of a properly-aligned lattice is that the atoms remain trapped with much less power than required for a dipole trap due to the fact that the intensity “seen” by the atoms increases by a factor of 4 when retro-reflecting the lattice beam. We then remove the retro-reflecting mirror¹³ and align the EOM to maximize the amount of light propagating through the EOM. The retro-reflecting mirror is then reinstalled and the retro-reflection is optimized by overlapping the retro-reflecting light with itself when passed through a

¹³ This step is necessary because the EOM crystal changes the position of the focus of the lattice beam after the cat’s eye lens L_3 .

pinhole poked into a card that fluoresces when the infrared lattice light is incident on it (an “IR card”).

Note that the secondary step of aligning the lattice without the EOM is not technically necessary, but experience shows that it is helpful when performing fine alignment of the lattice. The EOM crystal is not anti-reflection coated at 852 nm, and thus it is easy to fool oneself into thinking that the main lattice is created by the retro-reflecting mirror, when really the front facet of the EOM crystal is really providing the lattice. The main symptom of this is that shaking will not have any effect, and this is a common mistake made when first aligning the lattice. Once a shaking protocol (e.g. splitting) has been optimized, fine alignment of the lattice can also be done by shaking the lattice with this protocol and adjusting the lattice beam position and power until optimal splitting is found. It is entirely possible to optimize splitting, completely dismantle and rebuild the lattice system, and then be able to obtain splitting with the same optimized shaking function.¹⁴

For fine alignment of the lattice beam we installed a 2.5 mm thick AR-coated glass “tweaker” plate (TP, Thorlabs XYT-B) after the lattice focusing lens and before the vacuum cell. The tweaker plate was removed from its commercial mount and epoxied to a modified miniature mirror mount (Newfocus 9873-K) for ease of alignment. Plots showing how the beam position changes as a function of the angular mirror mount and tweaker plate angle are shown in Fig. 5.11. We find that practically, the day-to-day positioning of the angular mount does not need to be changed, but the tweaker plate does require some adjustment. Specifically, we must translate the lattice beam up and down to compensate for temperature drifts. These temperature drifts mainly affect our magnetic coil drivers (KEPCO BOP 20-10DM) and resulting drifts in the coil driver output current cause the pre-lattice waveguide magnetic trap position to change over time as the ambient temperature near the drivers changes. Calculations of the drift in the trap position as a function of coil current drift are shown in Fig. 5.12. This figure shows that the drift in the trap position may be attributed mostly to drifts in the chip wire and y -bias fields; these fields set the trap position.

¹⁴ The author speaks (writes?) from experience here.

Conversely, the x -bias field largely sets the trap “bottom” not in position but in energy. That is, this field sets an overall bias so that atoms are trapped in a non-zero magnetic minimum in order to minimize losses due to Majorana spin flips. Thus drifts in this field has a smaller effect on the trap position. In the future additional temperature stabilization will be required, but for the work presented here we simply monitored the atoms’ position along the first imaging axis over time and adjusted the tweaker plate so that the atoms’ position was roughly constant.

5.3 Shaking the lattice and measuring the results

We used two different methods to shake the lattice. For the interferometry results shown in Chapter 6 we used an EOM to shake the lattice. However, a second method for shaking was also explored; this method shakes the lattice by detuning the lattice laser. Both methods will be detailed below, with advantages and disadvantages given for each.¹⁵ Shaking can also be done by modulating the retro-reflecting mirror, but we did not pursue this option in the work presented here. This was largely due to the difficulties of getting a piezo-driven mirror with a high enough bandwidth to drive a mirror an appreciable distance $d \approx \lambda_L$ with decent linearity. Interested readers should refer to Ref. [13] for a good design idea. Our brief excursions exploring this design were not promising, likely due to the large displacements required to obtain the desired shaking amplitudes.

In practice we load our desired shaking waveform onto the arbitrary waveform generator (Agilent 33250A, AWG) before running the experiment. In the experiment the atoms are loaded into a stationary lattice. After loading the AWG is triggered and outputs the desired waveform to the electronics that control the shaking (either the EOM or laser driver, depending on the experiment). After shaking the lattice is turned off and the atoms undergo ballistic expansion in time-of-flight. After 20 ms TOF, the atoms are imaged in the far-field, revealing their momentum state population. In the lattice the momentum of the atoms is quantized in units of $2\hbar k_L$ (see Sec. 2.3.1). As in Chapter 4 it is this momentum state population that we seek to control. Details of

¹⁵ Lattice shaking by frequency shifting one of two independent laser beams is also possible, but we only used one lattice beam reflected back onto itself in this experiment.




Figure 5.12: Drift δz of the z -position of the pre-lattice waveguide magnetic trap as a function of drift δI of the current in the x -bias field (red), y -bias field (blue), main wire (black), and dimple wire current (blue). Calculations show that the drift in the x - and y -directions is negligible, as confirmed by the fact that only the z -position of the lattice beam needs adjustment over the course of a day.

the closed-loop learning system are given in Chapter 6.

5.3.1 Electro-optic shaking

Phase modulation of a retro-reflected lattice may be done by placing an EOM in the lattice beampath after the lattice beam has traveled through the 3D chamber [71], as shown in Fig. 5.9. The phase shift of the returning beam is thus proportional to the voltage applied to the EOM crystal, and this phase shift of the returning beam relative to the incident beam causes the nodes of the lattice to move. One advantage of this scheme is that for a given phase shift ϕ imparted by the EOM, the phase shift in the returning beam is 2ϕ due to the fact that the lattice beam passes twice through the EOM. Thus, lower voltages are needed to impart a given phase shift on the lattice. Our particular EOM is a Coherent/Quantum Technologies Model 28 four-crystal modulator. The EOM has two inputs, and the differential voltage between these two inputs is applied to the crystals and thus determines the phase shift imparted on the lattice beam.

A block diagram of the electro-optic shaking scheme is shown in Fig. 5.13. The desired shaking waveform is generated by the LabVIEW program running on the experiment control computer and loaded into the AWG. When the AWG is triggered by a LabVIEW-controlled digital signal, the output waveform is sent to two amplifiers (FLC Electronics A400DI). One amplifier increases the signal (and noise) by $20\times$, and the other inverts the signal then amplifies it by $20\times$. Thus, the total signal gain is $40\times$. The amplifier outputs (limited to ± 400 V by the amplifier electronics and the AWG) are then sent to the differential inputs on the EOM.

We can measure the linearity and the double-pass phase shift scaling of the EOM by placing it in one arm of an optical Michelson interferometer. A photodetector measures the intensity of the light at the interferometer output. Using the AWG and amplifiers, we apply a triangle-wave modulation at a given frequency and fit the resulting photodetector output to a sinusoid and determine the number of fringes that pass by in one cycle of the triangle wave. This gives us a measurement of the phase shift imparted by the EOM. We can then plot how this phase changes as we change the amplitude of the triangle wave output by the AWG (and thus the amplified voltage




Figure 5.13: Schematic of lattice shaking with an EOM. The lattice is shown in red, while the atoms are blue and the cell is gray. Amplifier gains are given by the numbers on the amplifier elements. The details of shaking are explained in the text.

applied to the EOM). The data and fit are shown in Fig. 5.14. We see that the EOM response is roughly linear, and from the linear fit we can measure the phase shift scaling to be $0.746(6)$ rad/V, where the parentheses give the error in the last digit.

5.3.2 Shaking via frequency modulation

Another method of shaking a lattice can be implemented by controlling the frequency of the lattice laser [56]. The effects of frequency-based shaking can be derived in the following way. Consider the one-dimensional lattice potential in Eq. (2.6), $V(x) \propto \cos(2k_L x)$. We know that in vacuum the light forming the lattice moves at speed c , and we can write the wavenumber $k_L = 2\pi\nu_L/c$, where ν_L is the Cesium-transition-locked frequency of the lattice light. Taking the atoms to be centered at a position d_0 from the retro-reflecting mirror, the potential at this point is

$$V(d_0) = \frac{V_0}{2} \cos(2k_L d_0) = \frac{V_0}{2} \cos(4\pi\nu_L d_0/c). \quad (5.1)$$

If we then detune this frequency an amount $\delta\nu$ from ν_L , we obtain the shifted potential

$$V_{\text{mod}}(d_0) = \frac{V_0}{2} \cos(4\pi\nu_L d_0/c + 4\pi\delta\nu d_0/c) = \frac{V_0}{2} \cos(4\pi\nu_L d_0/c + \phi(\delta\nu)). \quad (5.2)$$

That is, in the vicinity of d_0 , the potential $V_{\text{mod}}(x)$ looks like the regular potential $V(x)$ shifted by some phase $\phi(\delta\nu)$ that is a linear function of the detuning and the distance to the retro-reflecting mirror. Thus, by changing the frequency of the lattice we can approximate shaking. From Eq. (5.2) the frequency shift needed for a π phase shift is given by

$$\delta\nu_\pi = \frac{c}{4d_0}. \quad (5.3)$$

This is illustrated in Fig. 5.15 for parameters that roughly correspond to those used in the experiment. That is we set the distance $d_0 = 25$ cm, giving $\delta\nu_\pi = 300$ MHz.

This method is designed to work perfectly only at the designed position d_0 . As we consider the lattice light at positions away from this optimum position in the lattice, the difference in frequency causes errors in phase. The phase error (as a fraction of π) is plotted in Figure 5.16. The error is




Figure 5.14: Double-pass phase change of the EOM as a function of the voltage output by the AWG (the applied voltage to the crystals is this voltage amplified by $40\times$). The error bars are high for the low-voltage points due to the difficulty in measuring the phase at low applied voltage. A linear fit $\phi_{\text{lin}}(V) = aV$ to the data is shown in red, and a quadratic fit $\phi_{\text{quad}}(V) = aV^2 + bV$ is shown in blue. The linear fit gives $a = 0.746(6)$ rad/V and the quadratic fit gives $a = 0.12(2)$ rad/V² and $b = 0.64(1)$ rad/V. The linear fit is sufficient for our purposes.




Figure 5.15: The unshifted lattice potential (gray) and the lattice potential when a frequency shift of $\delta\nu = 300$ MHz is applied (black), showing the predicted π phase shift. The plot extends over $x = d_0 \pm 3\lambda_L$.

less than 4% over 10^4 wavelengths (for the experimental case considered here). For reference the Rayleigh range for our lattice beam (waist $w_0 = 40 \mu\text{m}$) is 5.9 mm, or about $6900\lambda_L$. Thus, phase errors are likely negligible relative to amplitude errors that arise due to the Gaussian nature of the lattice beam.




Figure 5.16: Phase error (as a fraction of π phase shift) as a function of distance (in wavelengths λ_L) for the experimentally relevant case where $d_0 = 0.25$ m and $\lambda_L = 85$ nm. The error is less than 4% over a range of 10^4 wavelengths.

The laser controller used for the lattice laser has a back panel input that changes the laser frequency proportionally to a voltage applied to the input. The laser shifts by 2 GHz per volt applied to this input, and the shift is applied within $400 \mu\text{s}$. In our case, we would require a voltage $V_\pi = 150$ mV to obtain a π phase shift. Naïvely scaling our bandwidth based on the lower voltages we require, our bandwidth then becomes roughly 167 kHz, which is enough for the experiments detailed here.

In practice, to go from EOM-based shaking to frequency-based shaking we simply change where the AWG output goes. For EOM-based shaking (Fig. 5.13) this output goes into the amplifier that then drove the EOM crystal. When using frequency-based shaking the output goes directly into the frequency jump input on the laser back panel. A TTL signal on another back-panel input signals to the laser that it is to be shifted by the voltage on the frequency shifting input. The

output of the AWG can then be controlled by our LabVIEW program and the experiment proceeds as in the EOM-based case.

5.3.3 Advantages and disadvantages of the two shaking methods

The EOM-based shaking method was the first to be implemented in the experiment. It is reasonably straightforward to set up, and the bandwidth of the shaking is limited to ≈ 300 kHz by the slew rate of the amplifiers. However, our EOM is bulky (8 cm long, and 5 cm in diameter) and requires a relatively high voltage to shift the phase of the lattice beam by π radians. Furthermore, the EOM crystal faces are not anti-reflection (AR) coated. We observe losses of about 75% through the EOM, the majority of which are from spurious reflections from the crystal faces. Assuming a crystal index $n \approx 1.5$ and two deleterious reflections (off of the crystal input and output faces), we assume that there is at most a residual 15% amplitude modulation of the lattice that occurs when shaking. While we will see in Chapter 6 that the optimization algorithm is able to account for these systematics, this residual AM can be deleterious when trying to make experiment and theory agree (e.g. for the work described in Chapter 7).

Frequency-based shaking does not have the high power requirements of EOM-based shaking, as relatively low voltages are required to shake the lattice. Furthermore, there are no additional optical elements required for this method, so there are no spurious reflections. Residual phase errors are also minimal relative to those caused by the spurious reflections. However, the effectiveness of this method is intrinsically tied to the device size. Thus, if we wanted a (relatively) compact system where the lattice beam focus is 5 cm from the retro-reflecting mirror, we would need to detune the laser frequency by 1.5 GHz, in which case the phase error (Fig. 5.16) increases from 4% to almost 20%. This restriction makes this setup unfavorable in most compact and practical setups.¹⁶

In a compact setup, one would almost certainly want to use a low-loss miniature EOM crystal anti-reflection coated at the lattice wavelength. The small crystal size would allow for relatively low

¹⁶ There is speculation that a cavity-based system (i.e. placing a cavity after the atoms, with the retro-reflecting mirror being one of the cavity mirrors) could make this system more compact.

power phase shifting (due to the lower voltage V_π required for a π phase shift) and the AR coating would eliminate reflections to less than 2% (assuming the AR coating is such that $R < 0.025$ at λ_L). As the technology advances, this scheme will become more and more feasible.

5.4 What's next

The next chapter will discuss the experimental results. First, the closed-loop optimization scheme will be discussed. Results of closed-loop optimization of the interferometer will be presented. Then, we will describe the calibration of the interferometer and show results that demonstrate the T^2 scaling of the interferometer sensitivity. It will be demonstrated that interferometer output is sensitive to the direction of the applied signal. We will also show how the interferometer can be optimized in the presence of a DC bias signal and that perturbations on this bias can be measured. Finally, Chapter 7 will extend the theory presented in Chapter 2 and explore shaking (and specifically splitting) from an analytical perspective.

Chapter 6

Optimization and characterization of a shaken lattice interferometer

In this chapter we will report on results of building a shaken lattice interferometer using the experimental setup described in Chapter 5. First, we will report on our efforts to build a closed-loop experimental system in which the experimental optimization and calibration can proceed with minimal human input¹. Section 6.2 will describe learning in the closed-loop system. We will then move on to the results of optimizing and calibrating a shaken lattice interferometer in Sec. 6.3. In this section we will describe how we add a fictitious acceleration signal to the system and determine the interferometer sensitivity. We obtain a sensitivity of T_1^2 , consistent with the simulation results presented in Sec. 4.4. Finally, some brief next steps will be described in Sec. 6.4; these next steps will be elaborated on in Chapter 8.

6.1 Closed-loop optimization in the experimental system

Our experimental control system is very similar to that used in Refs. [15, 53, 99]. The LabVIEW control code serves as the “brain” of the system, creating the timing sequence that is fed to the FPGA in the NI PXI-8186 chassis that controls the analog outputs (NI 9263 and 9269) and digital inputs and outputs (NI 9403) that drive the experiment.² In our system, we use NI cRIO hardware with 16 bit analog output and a temporal resolution of 20 μ s for both analog and digital output. The control code that runs the experiment is adapted from code written for our group by

¹ Someone generally has to be around to make sure the lasers stay locked and the trap bottom does not drift. This can be mitigated by stabilization of the lasers and magnetic fields. See Chapter 8.)

² Our only digital input is a Schmitt trigger circuit that synchronizes our experiment start time with the 60 Hz wall power cycle.

Farhad Majdeteimouri.³

In the past closed-loop operation of cold-atom experiments in our group was not possible. Every experimental run had to be initiated by a human being. In our system, this became tedious when attempting to optimize the shaken lattice interferometry sequence, as we were using LabVIEW to initiate and control the experimental sequence and MATLAB to analyze the experimental images, run the optimization, and communicate with the AWG (see Chapter 5). Therefore, we modified the control code so that experimental control, optimization, and imaging were based on a state-machine architecture in LabVIEW and ran continuously. This scheme requires human input only when the total atom number dips below a certain value; this generally indicates a laser unlock or trap bottom drift.⁴

There are two separate control codes used in the experiments presented in this work. The first implements the CRAB algorithm and optimizes the lattice. This code also has built-in functionality to scan shaking frequencies and amplitudes (without optimization) in order to map out the parameter space or take the single-frequency shaking experimental data shown in Chapter 7. The second calibrates the interferometer by adding a signal to the atoms (see Sec. 6.3.2.1) and recording the final state after shaking. We have added a third control code that optimizes the experiment using the CR bound as the error measure, as discussed in Sec. 8.3. However, these experiments are just getting started and have not yet produced results.

In Fig. 6.1 we show a cartoon block diagram of the closed-loop experiment as it is used during optimization. We set up the experimental parameters using the LabVIEW front end, then LabVIEW loads the proper timing sequence onto the FPGA and the proper shaking function onto the AWG. The experiment then runs and subsequent time-of-flight images of the atoms are recorded. If the atom number is above a certain level set by the user, the image is analyzed and

³ So much thanks to Farhad for sitting down and explaining some of the relevant bits of the code. This made the implementation of a closed-loop system so much easier.

⁴ Currently, a laser unlock is easy to fix without having to break out of the experimental loop. A trap bottom drift requires one to stop and restart the experiment, but functionality to change the trap bottom without breaking the experimental loop could be added without much hassle. In practice we simply allow the experiment to stabilize over a few hours. We then run late into the night when the absence of people and stable temperatures cause the experiment to run smoothly. For example, the best data presented in this chapter was taken at about one a.m. during Fall 2017.

the momentum state population of the atoms is measured via atom counting. If the atom number is below this threshold, the control code stops and waits for human input (e.g. relocking the lasers and sending the command for the experiment to continue). After three separate runs with the same experimental setup, the mean momentum state population is calculated, giving us the vector \vec{P}_f corresponding to this final state (see 3.3). The error of this state relative to the desired state \vec{P}_d is then calculated via Eq. (3.4) and the algorithm proceeds as described in 3.5. All of the optimization is coded into LabVIEW so that the entire system runs smoothly.

6.2 Learning in the closed-loop system

6.2.1 Narrowing the frequency band used in the optimization

In general learning is much faster when one narrows down the search space to only include the most “influential” parameters; this is explored in more detail in Chapter 7. While our simulations are relatively cheap in terms of computation and graduate student time, the experiment is much more time-consuming. Each experimental run takes roughly one minute, although this time could be lessened significantly in a practical implementation of the system [39]. Thus, we want to limit the frequency band in which we split.

In order to understand the best frequency band in which to limit our optimization, we shook the lattice for 1 ms with a shaking amplitude of 1 V (2 Vpp) and scanned the shaking frequency from 10 to 40 kHz.⁵ The typical envelope function Eq. (3.5) multiplies each shaking function to ensure smooth turn-on and turn-off. After shaking, we measured the error of the final state relative to the split state. Results of this are shown in Fig. 6.2.

We see from Fig. 6.2 that the region from 18 – 30 kHz gives the lowest error relative to the split state. Outside of this regime, the error hovered around 50 – 60% and changed very little as the frequency was scanned. Thus, this is the frequency band used in all of the optimizations presented in this chapter. In Chapter 7 we analyze the effect of shaking frequency on splitting in more detail.

⁵ The analysis done in Ch. 7 was just getting started around the time that this data was taken. Therefore, a more “brute force” approach was taken to obtain the results presented in this section.

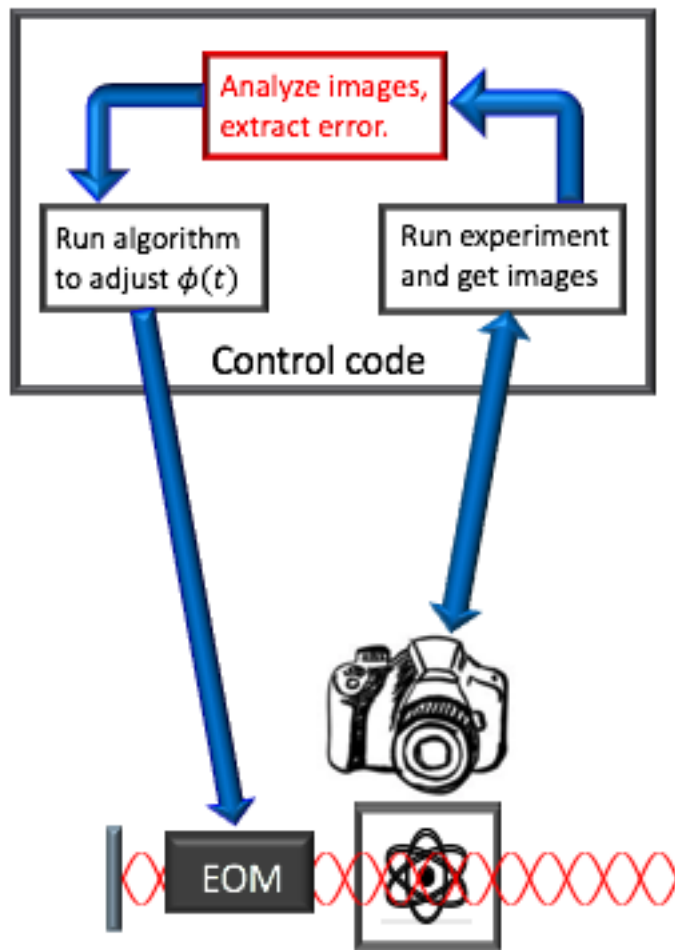


Figure 6.1: A block diagram of the closed-loop optimization protocol used in the experiments shown in this chapter. The control code runs the experiment, analyzes the images produced by the camera, and updates the lattice shaking function based on the optimization algorithm.




Figure 6.2: Results of scanning the shaking frequency from (a) 10 – 18.25 kHz, (b) 18 – 28 kHz and (c) 26 – 42 kHz and measuring the percent error of the final state with respect to the split state. The different colored curves simply denote different experimental runs. The points (error bars) denote the mean (standard deviation) of the error. The data was taken over the course of three days (each separate plot corresponds to one day of data), and this explains small “jumps” in the error as one moves from plot to plot.

From the analyses presented there, we see (in retrospect) that this band is most “influential” due to the fact that the transitions between the ground Bloch band $r = 0$ and the excited bands $r' = 1$ and 2 lie in this band.

6.2.2 An experimental learning curve

A sample learning curve for the experiment is shown in Fig. 6.3. To obtain Fig. 6.3 we optimized splitting using the CRAB algorithm discussed in Sec. 3.5. For each iteration the percent error is calculated as in Eq. (3.4) for each of the 5 simplex vertices used here, and error corresponding to the best of these 5 vertices is shown in Fig. 6.3. The curve in (a) shows the initial optimization, and (b) shows further optimization using the optimum found in curve (a). This shows that we were able to stop optimization and pick it up again at (roughly) the same place after e.g. improving the evaporation scheme (which would change due to the trap bottom drift discussed in Chapter 5).

Such curves are typical for the experimental learning process,⁶ but it should be noted that we found that in some cases we could optimize within one or two iterations. The physics behind such lucky optimization runs is discussed in more detail in Chapter 7 and in such fortuitous cases we were able to quickly learn how to split the atom wavefunction.

Experimentally, we were also able to implement the dCRAB algorithm discussed in Sec. 3.5.3. This allowed us to continue optimizing even if the optimization got stuck in a local minimum. In practice, we “re-rolled” our optimization frequencies after the percent error remained stagnant for N_I iterations. In practice, we set N_I to 2 or 3. As described in Sec. 3.5.3, re-rolling takes the individual vertex in the simplex that corresponds to the best shaking function and moves it forward unchanged to the next iteration (a form of elitism). The other vertices in the simplex were identical to the best individual, but a second shaking function was added to it as in Eq. (3.33). In this way, we were able to keep the optimization from stagnating.⁷

⁶ Although we did see some improvement between the curves shown in (a) and (b), usually the small “jump” expected due to experimental drift was slightly deleterious.

⁷ We generally used 5 vertices to generate our experimental simplex.




Figure 6.3: Sample learning curves for the experiment. The points (error bars) show the average (standard deviation) of three points. The optimization resulting in the learning curve in (a) was done first, then the results were improved upon in a second optimization run that took the best results of optimization (a) and produced the learning curve in (b).

Due to experimental limitations such as imaging noise, shot-to-shot drift, finite cloud momentum, and heating, we were limited to a splitting error of $E_{\min} \approx 10\%$. The effect of finite cloud momentum is discussed in Sec. 6.2.3, and the limits of optimization are discussed in more detail in Sec. 6.3.1. As a result an error of about 10% was roughly the level at which our optimization was converged sufficiently. In Sec. 6.3.2, we show that we can obtain T_1^2 sensitivity in the experimental shaken lattice interferometer (consistent with the simulation results presented in Sec. 4.4), even with such errors.

6.2.3 The effect of a finite momentum width

All of the simulations done in Chapter 4 assumed that the atom cloud had an infinitesimally small momentum width. However, in the experiment, the atoms have a finite momentum width. This width is set by the repulsive atom-atom interactions and the trap in which the atom cloud is evaporated to BEC (see Sec. 5.1.3)⁸. After time-of-flight, we measure an atom momentum width of $\delta p \approx 0.6\hbar k_L$; this value changes based on the lattice depth and the atom number, but this width is typical for the experiments presented in this chapter and Chapter 7. Simulations with these cloud widths do not generally obtain errors below a few percent, so this is likely one of the hard limits to our optimization protocol.

6.3 Experimental shaken lattice interferometry

This section will present results on experimental shaken lattice interferometry and is based on the results presented in Ref. [118]. Most figures in this section come from Ref. [118]. First, we will describe how we build the interferometer sequence. In order to quantify the interferometer sensitivity, we built five separate atom interferometers with different interrogation times. This is detailed in Sec. 6.3.1. Then, in Sec. 6.3.2 we will describe how the interferometer sensitivity was measured. The method of applying a simulated acceleration signal is described in Sec. 6.3.2.1 and

⁸ While we do decompress the trap when going from BEC to the optical lattice, the atoms still have an appreciable momentum width when loading into the lattice

the resulting sensitivity measurements are described in Sec. 6.3.2.2. From there we will describe some next experimental steps that will improve on the results presented here.

6.3.1 Optimizing the interferometers

As stated above we perform gradient-free, closed-loop optimization based on the dCRAB algorithm described in Sec. 3.5. As stated in Chapter 3 experiments have used this algorithm to optimize the state inversion of a BEC [14, 114], in Ramsey interferometry schemes [115], or to calibrate qubit operations in diamond NV centers [41]. Optimization protocols have also been used in cold atom [98, 121] and quantum optics experiments [64] as well as in efforts to find efficient pulse schemes in light-pulse atom interferometry [109].

Our experiment is based on the compact BEC setup described in Chapter 5. For the work presented here we load our Bose-condensed ^{87}Rb atoms into the ground Bloch state [50] of a red-detuned optical lattice with a depth $V_0 \approx 14E_R$. We shake the lattice using the EOM-based technique described in Sec. 5.3.1.

To build the interferometer sequence, the shaking function is optimized to cause the atoms in the lattice to undergo the desired state-to-state transformations. In particular we “stitch” together shaking functions corresponding to different interferometer operations, just as in Chapter 4 [119]. As before, to implement an atom beamsplitter we begin by loading atoms into the ground Bloch state of the lattice. The lattice is subsequently shaken to split the atom wavefunction so that roughly half of the atoms occupy each of the $\pm 2\hbar k_L$ momentum states. We can then optimize separate propagation protocols that maintain the atoms’ momentum-space splitting. To recombine the atoms back into the ground state (in the absence of an applied signal), the optimized splitting shaking protocol is run in reverse. Each separate protocol is $T = 0.2$ ms in duration and is multiplied by the envelope function in Eq. (3.5) to ensure smooth turn-on and turn-off of the shaking. The smooth turn-on and turn-off allows the shaking functions to be stitched together without discontinuity. In this way we optimize five separate interferometers with interrogation times of $T_n = 0.4n$ ms, where $n = 1, \dots, 5$.

To optimize our interferometer we define the target state to be the split state. The dCRAB algorithm picks five frequencies at random within our chosen frequency band of 18 – 30 kHz (see Sec. 6.2) and assigns each frequency five separate Fourier sine and cosine amplitudes. The five waveforms described by these Fourier coefficients become the five vertices of a simplex in frequency space. Using the Nelder-Mead algorithm, the simplex is modified and iteratively converges upon the target state. Error is determined in the usual way by building a vector \vec{P} with components P_n containing the relative population of atoms in the $2n\hbar k_L$ momentum states. In practice there is negligible population in the $\pm 6\hbar k_L$ states, so n is truncated to $\pm N$ where $N = 2$. The percent error E is then defined as in Eq. (3.4), where the desired momentum population vector is $\vec{P}_{\text{des, sp}} = (0, 0.5, 0, 0.5, 0)$.

Two examples of optimized shaking functions are shown in Fig. 6.4, corresponding to interferometry sequences with interrogation times of 0.4 ms and 2 ms in length. While splitting requires relatively high shaking amplitudes, smaller amplitudes are required to maintain this state during propagation. This is likely due to the fact that the split state is similar in momentum state to the first excited Bloch state of the lattice, so less modulation is required to maintain a state close to a lattice eigenstate than to transform from one state to another nearly orthogonal state. Note that the splitting and recombination times are included in the definition of the total interrogation time because they are not negligibly small relative to the propagation time.

The error in splitting begins at roughly 10% and increases as propagation protocols are added (see Fig. 6.5(a)). This error arises due to spurious atoms detected in the $0\hbar k_L$ momentum state due to atom localization in the deep lattice potential [101] and atom-atom interactions that cause heating and loss of visibility during the experiment [23, 24, 96, 101] or during time-of-flight [43]. We are also limited by asymmetry in the two split clouds and the finite momentum spread of the condensed atoms as they are loaded into the lattice. Upon recombination our errors are $< 10\%$ as shown in Fig. 6.5(b). Errors in recombination manifest largely in population of higher-order momentum states as a result of accumulated errors in the splitting and propagation protocols. The error in recombination is lower than that of splitting because spurious atoms detected in the $0\hbar k_L$




Figure 6.4: Two example shaking functions. (a) Splitting and recombination shaking protocols, totaling $T_I = 0.4$ ms. (b) Splitting and recombination protocol from (a), but with 4 propagation steps added in between, resulting in an interrogation time $T_I = 2$ ms. We optimize the interferometer so that the atoms remain split at the end of each propagation step. In both cases the second half of the shaking protocol is simply the reflection of the first half.

momentum state are no longer deleterious.

6.3.2 Determining the interferometer sensitivity

Once we have the interferometer sequence optimized, it can be used to measure something practical. However, it is interesting to know just how sensitive the interferometer is. In particular the sensitivity scaling of the sensor is important, as this tells us how the interferometer response ties to the interrogation time. In order to do this we must add a signal to the system and record the interferometer response. To back out the sensitivity, we measure the classical Fisher information as in Sec. 4.4 and use the Cramer-Rao bound to obtain the lower bound on the experimentally measurable sensitivity.

6.3.2.1 Applying a signal via gradient coils

Currently, accelerating our interferometer is not practical, so we must get creative. To calibrate the atoms' response to an applied signal, a pair of magnetic coils placed outside the cell provides a magnetic field gradient $G = \partial B / \partial x$ along the lattice direction.⁹ These coils can be seen in Fig. 5.10(b) and the accompanying inset. A small bias magnetic field remains on while the atoms are trapped in the lattice to maintain the atoms' spin polarization and thus their sensitivity to magnetic fields. This sensitivity to a magnetic field gradient also shows that our system can be configured as a magnetic field gradiometer. Note that if we were measuring pure accelerations, we would not need to ensure that the atoms are spin-polarized in the lattice.

The gradient G gives rise to an effective acceleration

$$a_{\text{eff}} = G g_F m_F \mu_B / m \quad (6.1)$$

where $g_F = 1/2$ is the Landé g-factor [106], and μ_B is the Bohr magneton. The data to produce this calibration factor is shown in Fig. 6.6. In practice we calibrate the acceleration due to the

⁹ Because the coils are offset from the cell center (due to the need for optical access at that point), a gradient is also applied in the y -direction orthogonal to the lattice direction (and to gravity), but the applied gradient is so small that the only effect is a minor and negligible shift in the lattice trap position along this direction.




Figure 6.5: Percent error in the (a) splitting and (b) recombination protocols as a function of (a) the splitting time $T_1/2$ and (b) the total interrogation time T_1 . The insets show example experimental images from each protocol, where the color represents optical density (OD). (a, inset) An image from optimized splitting of the atoms into equal population in the $\pm 2\hbar k_L$ momentum states and (b, inset) recombining the atoms into the ground Bloch state.

gradient field by loading the atoms into a dipole trap and measuring their velocity as a function of hold time in the trap while varying G . The applied acceleration a_{eff} increases linearly with current through the gradient coils ($a_{\text{eff}} = 0.71 \pm 0.16 \text{ m/s}^2/\text{A}$), as expected from calculations using the Biot-Savart law.




Figure 6.6: Data showing the atoms' acceleration as a function of magnetic field gradient G applied via two coils placed on the side of the cell. The coils (visible in Fig. 5.10(b)) are oriented so that the current through them goes in opposite directions, resulting in a magnetic field gradient felt by the atoms.

6.3.2.2 Measuring the interferometer sensitivity

To quantify the performance of the interferometer, we measure how the momentum state population vector \vec{P}_a changes as an acceleration signal a is applied to the atoms. As determined in our earlier work [119], because there are more than two momentum states considered in this interferometry scheme, we cannot assign a phase difference based on the relative path length traveled

by two arms of the interferometer. Thus, we use the classical Fisher information (CFI) to define the lowest detectable acceleration δa based on the Cramer-Rao (CR) bound [48]. As before, we use the expression for the CFI given in Eq. (4.10) [119]. We numerically evaluate our derivative using a two-point forward-difference scheme. The CR bound allows us to find the minimum detectable acceleration $\delta a = 1/\sqrt{F_{C,P}}$ as in Eq. (4.11).

Results of performing this analysis are shown in Fig. 6.7. The data is fit in Mathematica using the Levenberg-Marquardt scheme [66, 72] to a function $f(T_I) = aT_I^{-b} + c$ where b is the sensitivity scaling and c is a noise-limited offset that we can quantify. Therefore we fit only the values of a and b . In order to mitigate the effects of imaging noise we set a threshold OD below which we do not count atoms. We find that the optimum value of this threshold is $OD_{\text{thresh}} \approx 0.05 - 0.06$ depending on the imaging noise. The largest contributor to this noise is imbalance in the exposure time between the absorption and background images due to our minimum experimental resolution of $20 \mu\text{s}$. Small, slow drifts in the probe beam power contribute to the drift in the noise from day-to-day. The residual noise in imaging contributes to the offset c in the fit.

This offset is measured by measuring the CR bound with no atoms present, then dividing this number by the ratio of the detected atom number with and without atoms actually being present. Future work will focus on the reduction of this offset by improving the exposure balance in the imaging system. Furthermore, the atom signal-to-noise can be improved by minimizing the heating of the atoms in the lattice [12], which will also allow for longer interrogation times. However, longer interrogation times will increase the likelihood of decoherence due to phase diffusion [55], which can be mitigated by lowering the atom density before loading into the lattice.

Our fit (Fig. 6.7) gives a scaling $b = 1.96 \pm 0.13$ using our measured value of $c = 0.014(3)$, consistent with the expected T_I^2 scaling. The measured value for c and the fit for b are consistent with our results when we leave both b and c to be free parameters; in this case, we measure $b = 2.20 \pm 0.34$ and $c = 0.015(2)$. Data taken on different days gives scaling that is consistent with the expected T_I^2 scaling, and the data presented here is a typical example. Furthermore, data taken where a signal is applied to unshaken atoms is indistinguishable from noise and shows

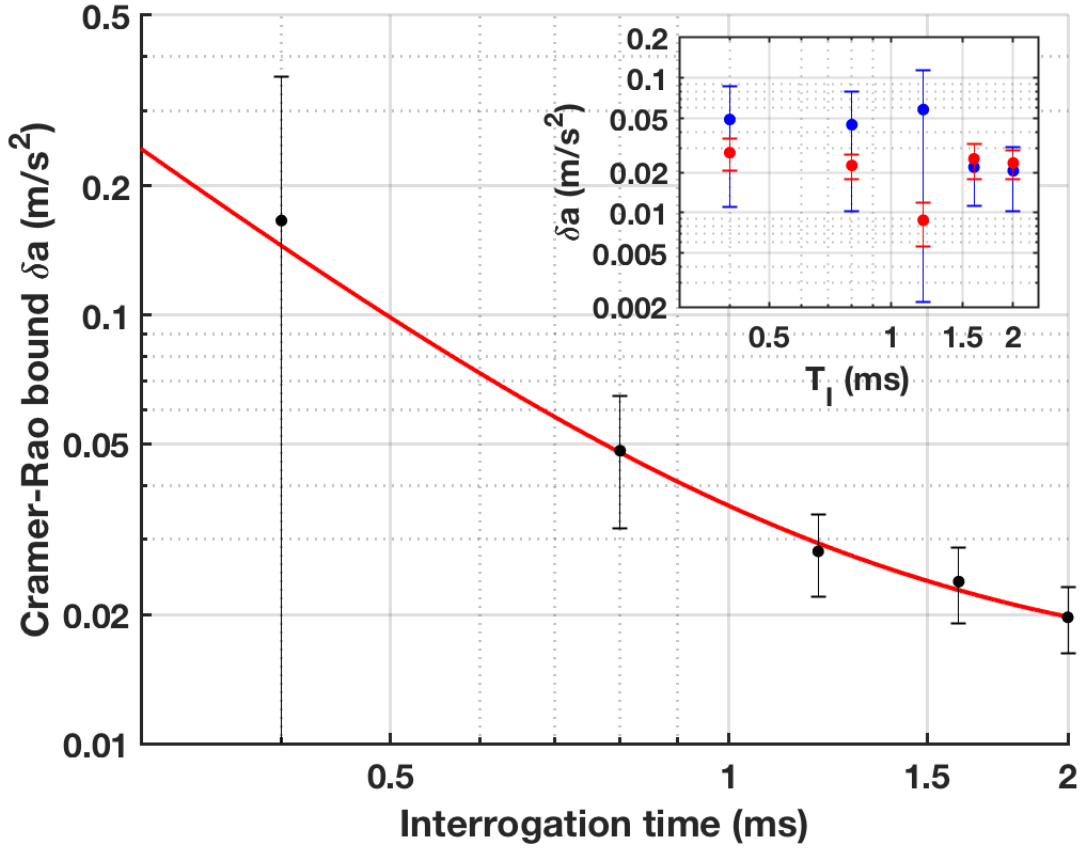


Figure 6.7: Minimum detectable effective acceleration δa plotted as a function of interrogation time for each of the five interferometers optimized for this work (black) and fit (red) to $f(T_I) = aT_I^b + c$. The scaling value b is consistent with the expected T_I^2 scaling, and the offset c arises due to imaging noise and is measured experimentally. (inset) Data taken with no atoms present (blue) and no shaking applied to the atoms (red) showing no signal other than imaging noise. Blue data is scaled by the ratio of the relative atom number as explained in the text.

no discernable scaling law, showing that the shaking is a coherent process. For the accelerations presented here, the interrogation time is much less than the Bloch oscillation time $\tau_B \propto 1/a$. Thus we do not expect the atoms' population to change dramatically when accelerated without shaking. Due to the sensitivity scaling with T_I^2 , as we increase T_I , the lower values of a that we are sensitive to strengthen this assumption. Thus, we can neglect the effects of Bloch oscillations in our analyses.

6.3.3 Experimental confirmation of other interesting aspects of shaken lattice interferometry

We were able to show in Chapter 4 two unique and interesting aspects of shaken lattice interferometry. First, we showed that one can measure both the magnitude and direction of the applied signal, then we demonstrated that one can tune the interferometer transfer function. This work showed that we can tune the sensitivity of the interferometer to the class of signals that we want to measure.

Here, we demonstrate experimentally our ability to measure the magnitude and direction of an applied signal. Furthermore we take first steps in the direction of transfer function tunability by optimizing the interferometer as in Sec. 6.3.1 but in the presence of a bias signal.

6.3.3.1 Directionality measurements in the shaken lattice interferometer

We can calibrate the interferometer response to a signal by recording how the final state of the interferometer after shaking changes with the applied signal. Because the lattice shaking breaks the symmetry of the system [119], we can determine the sign of an applied signal. We measure the variation of the atoms' final momentum state after the interferometry sequence, as shown in Fig. 6.8. The data show that the final state after an acceleration a is applied is distinct from the final state after an acceleration $-a$ is applied. This ability to distinguish the signal direction differentiates our interferometer from the typical light-pulse atom interferometer where the atom population varies sinusoidally between two states.




Figure 6.8: Momentum population of the atoms after the $T_1 = 2$ ms interferometer sequence as a function of the applied acceleration signal. Atoms in the $2n\hbar k_L$ state are denoted by open blue circles ($n = -2$), blue crosses ($n = -1$), black dots ($n = 0$), red pluses ($n = 1$) and red asterisks ($n = 2$). As the applied signal is varied away from zero, we can distinguish positive and a negative signals. The dotted lines are cubic spline fits to guide the eye.

6.3.3.2 Optimizing in the presence of a bias acceleration

Finally we show steps towards the tunability of the interferometer transfer function [119]. We optimize the interferometer in the standard Michelson configuration but add a bias signal $a_{\text{bias}} = -0.71 \text{ m/s}^2$ during optimization. We then measure the atoms' final recombined momentum state after the addition of signals $a_{\text{bias}} \pm \Delta a$, as shown in Fig. 6.9. From this data, we see that we can distinguish the sign of Δa by observing the final state of the atoms. Further extensions of this work include increasing the magnitude of a_{bias} and optimization of the interferometer to an AC-varying signal, as predicted in Ref. [119]. This will allow the interferometer to be optimized for sensitivity to any signal of interest.

6.4 Next steps

In this chapter we have presented experimental results of interferometry using atoms trapped in an optical lattice, showing that shaken lattice interferometry scales as T_I^2 . Notably, the sign of the applied signal may be measured, and the interferometer may be optimized in the presence of a bias signal. We show that the limitation on our interferometer sensitivity is set by imaging noise, which may be mitigated with some straightforward experimental improvements. These improvements will be detailed in Chapter 8, but we outline some brief goals below.

If we were to improve the imaging signal-to-noise ratio, we could then work to increase the interferometer sensitivity. We could accomplish this in one of two ways. First, we could increase the interferometer interrogation time. We could also increase the magnitude of interferometer splitting (e.g. from $\pm 2\hbar k_L$ to $\pm 6\hbar k_L$). This is interesting because the sensitivity scales directly with the splitting order in free-space interferometers, and we expect the same to be true in shaken lattice interferometry. Furthermore, if we were able to split the atoms with high enough momentum, we could potentially accelerate them in separate moving lattices. The feasibility of higher-order splitting as it relates to this goal is described in Chapter 7. The methods described there could also lower the heating due to atom-atom interactions [23, 24, 96].

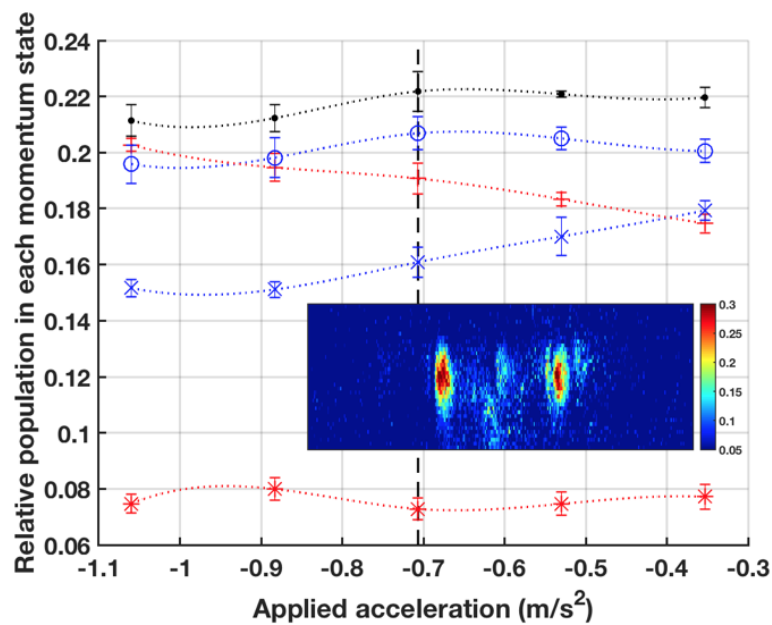


Figure 6.9: As in Fig. 6.8, plot of the momentum state variation as a function of applied acceleration with the biased interferometer, showing variation of the final state as the acceleration is varied from the bias value of $a_{\text{bias}} = -0.71 \text{ m/s}^2$ (black dashed line). Data points and splines colored as in Fig. 6.8. (inset) An experimental image of the optimized split state in the biased interferometer. OD is indicated by the colorbar on the right.

It would also be of interest to build a three-axis accelerometer using a three-dimensional lattice. Ideally, each axis could be interrogated simultaneously but without crosstalk. This also opens up the possibility of operating the shaken lattice interferometer in a two-dimensional Sagnac gyroscope configuration. This configuration is detailed in Chapter 8 and will allow us to build a rotation sensor. In this way, we could make progress towards the ultimate goal of this project: building a compact six-axis inertial sensor [119].

Chapter 7

An analysis of splitting in the shaken lattice interferometer

7.1 Introduction

Previous work presented in this thesis and in Refs. [92,93,119] has demonstrated that one can cause atoms to undergo state-to-state transitions in a phase-modulated (or “shaken”) lattice system. This is the bread and butter of the work presented here. However, in the previous chapters rigorous analysis has occasionally been eschewed in favor of simple proof-of-principle demonstrations. In order to truly improve shaken lattice interferometry we must analyze shaking in more detail. This chapter performs this analysis, and it is hoped that the results presented here will be of use as the system is improved by future students.

As such, we explore here in detail how one can shake an optical lattice to transform the wavefunction of atoms trapped in the lattice. Specifically, we seek to reduce the dimensionality of the shaking control landscape. The motivation for this is twofold: first, by simplifying the optimization landscape we can improve the efficiency of learning [81]. This is particularly important in experiments like ours that are limited by drift or finite cycle times. Second, we wish to limit heating and decoherence in the shaken lattice system. Recent work has shown that atoms in a shaken lattice are susceptible to decoherence [4, 54, 68, 122] when shaken at certain amplitudes and frequencies, both in the presence and absence of a signal. Furthermore, shaking of a BEC trapped in an optical lattice has been shown to cause heating in the condensate due to atom-atom interactions [34, 96]. Atom scattering into transverse modes has also been shown to be deleterious [23, 24]. Therefore, it is of interest to analyze the lattice shaking protocols that drive

these state-to-state transitions and reduce the subspace to eliminate deleterious shaking frequencies.

The desired transformation considered in this chapter is the first step of lattice-based interferometry, the splitting of the atom wavefunction in momentum space. That is, we wish to start with atoms in the ground state of the lattice and transform them to a “split” state with an error of less than 1%. In particular, we consider in this chapter the simplest cases of splitting the atom population equally into the $\pm 2n\hbar k_L$ states for $n = 1, 2$, and 3. Note that in previous chapters of this work only splitting with $n = 1$ has been considered; this chapter attempts to better understand both this first-order splitting and higher-order splitting with $n > 1$. Furthermore, we will introduce the relative phase θ between the two counterpropagating momentum components. That is, $|\psi_n(2n\hbar k_L, \theta)\rangle = e^{i\theta}|\psi_n(-2n\hbar k_L, \theta)\rangle$ for the n th order split state $|\psi_n(p, \theta)\rangle$. Previously we have ignored this relative phase, as it was inaccessible in experiment. In simulation we can access θ , and we use this to draw conclusions about the wavefunction splitting later in the chapter.

For the simplest case of $2\hbar k_L$ splitting we show that if the lattice is shaken at frequencies near the Bloch band 0 to band 1 transition, we can split the atom wavefunction to within the desired error. This transition gives a relative phase difference of π between the two momentum states in the resulting split state. If we shake the lattice at half of the band 0 to band 2 transition frequency, we can split the atoms to within 1% error with zero relative phase between the two momentum states. In each case the simulation results are backed up by analytics. This simple shaking scheme is not suitable for higher-order splitting because the transition rate between bands drops precipitously as we transition from band 0 to higher-lying bands. However, we find that if we optimize splitting via a genetic algorithm using only the band-to-band transition frequencies, we can achieve 4 and $6\hbar k_L$ splitting within 1%.

We show that after splitting to third-order ($\pm 6\hbar k_L$) we can load the atoms into counterpropagating moving lattices and accelerate them to improve the splitting, potentially achieving acceleration sensitivities that scale as T_I^3 [74]. Thus, we consider splitting only up to third order in this chapter.

The chapter is structured as follows: In Sec. 7.2, we motivate the description of the lattice

dynamics in terms of the Bloch states and describe the split state in terms of these states. Section 7.3 will describe an analytic treatment of the problem. Section 7.4 discusses the efficacy of $2\hbar k_L$ and higher-order splitting with single-frequency shaking. Section 7.5 shows results of optimization with a genetic algorithm where we restrict ourselves to shaking at the band-to-band transition frequencies. Section 7.6 concludes.

7.2 Band decomposition of the lattice wavefunction

As a reminder, for the simulation results presented in this work, we have made the following assumptions: First, we assume that the atoms are delocalized in the lattice, i.e. in a superfluid state [44]. We will assume that the atoms are initialized with quasimomentum $q = 0$ and this quasimomentum does not change. Finally, we assume the atoms are non-interacting and that the lattice is infinite.

Because we are working in the regime where the atoms are delocalized in an infinite lattice, the Bloch states $|\Psi_r^{(q)}\rangle$ discussed in Sec. 2.3.1 are a convenient basis for calculations, where r denotes the band number and q is the quasimomentum. As in Chapter 4 the lattice depth was chosen to be $V_0 = 10E_R$, where the recoil energy $E_R = \hbar^2 k_L^2 / 2m$ for atoms with mass m . The band energy E is plotted against the quasimomentum q in Fig. 2.1. The atoms begin in the state corresponding to the ground Bloch band $r = 0$ with $q = 0$. Since we assume that the quasimomentum is always zero we will suppress the index q in what follows and denote the Bloch states $|\Psi_r^{(0)}\rangle$ as simply $|r\rangle$.

We know from Sec. 2.4 that the Bloch states at zero quasimomentum populate only the $2n\hbar k_L$ momentum states (for integer n). Therefore a complete description of the atom wavefunction can be given by the amplitudes and (relative) phases of the wavefunction in these quantized momentum states. Of particular interest is the relative phase θ between the two momentum components of the split state, as defined in section 7.1.

Experimentally one cannot determine these relative phases in time-of-flight images, as we only have access to the amplitudes in experimental measurements. As such, in simulation and experiment we define the vector \vec{P} as in Sec. 3.3 where the components P_n contain the probability

of finding an atom in a given momentum state $2n\hbar k_L$. If we consider an ensemble of atoms, this vector gives the relative population of atoms in each momentum state. We can then define an “error” E_{ab} between two momentum states described by vectors \vec{P}_a and \vec{P}_b as in Eq. (3.4). Note that if we are comparing any state to the split state, E will be independent of θ and is thus a more useful quantity to look at when considering the results in the context of what is experimentally observable. Thus, we use this expression as an error measure to quantify how well our optimization algorithm is doing.

For bands $r > 0$ there is considerable similarity between the Bloch states and split states of various orders. This is most easily seen when one looks at the momentum state population of the Bloch states corresponding to different bands, as shown in Fig. 7.1 [50]. Interestingly, there are two separate Bloch states at different band energies that roughly correspond to each split state. To glean further insight, we calculate the inner product D_{nr} between the n th order split state $\psi_n(p, \theta)$ and the state $|r\rangle$ as

$$D_{nr} = |\langle r | \psi_n(p, \theta) \rangle|^2. \quad (7.1)$$

From Fig. 7.2, we see that the difference between two Bloch states corresponding to bands $l > 0$ and $l' = l + 1$ is that the lower-energy state $|l\rangle$ has a relative phase difference $\theta = \pi$ between the $\pm(l + 1)\hbar k_L$ states, and the higher-energy band $|l'\rangle$ is almost identical, except $\theta = 0$ (and thus the two states are orthogonal). This is commensurate with the fact that states corresponding to adjacent bands have opposite parity.

In the rest of this chapter, we refer to state-to-state transitions between different bands. The transition frequencies for transitions from the ground band to the first 5 excited bands are shown in Fig. 7.3. For an example of the energy scales at play here, we tabulate the transition frequencies $f_{r,r'}$ between two bands r and r' in Table 7.1. The frequencies given in Table 7.1 assume that we are working with ^{87}Rb atoms (as we do in all of this work) at a lattice depth of $V_0 = 10E_R$. We see that the band transitions lie between 0 and 121 kHz, and this is the regime in which our driving is simulated.

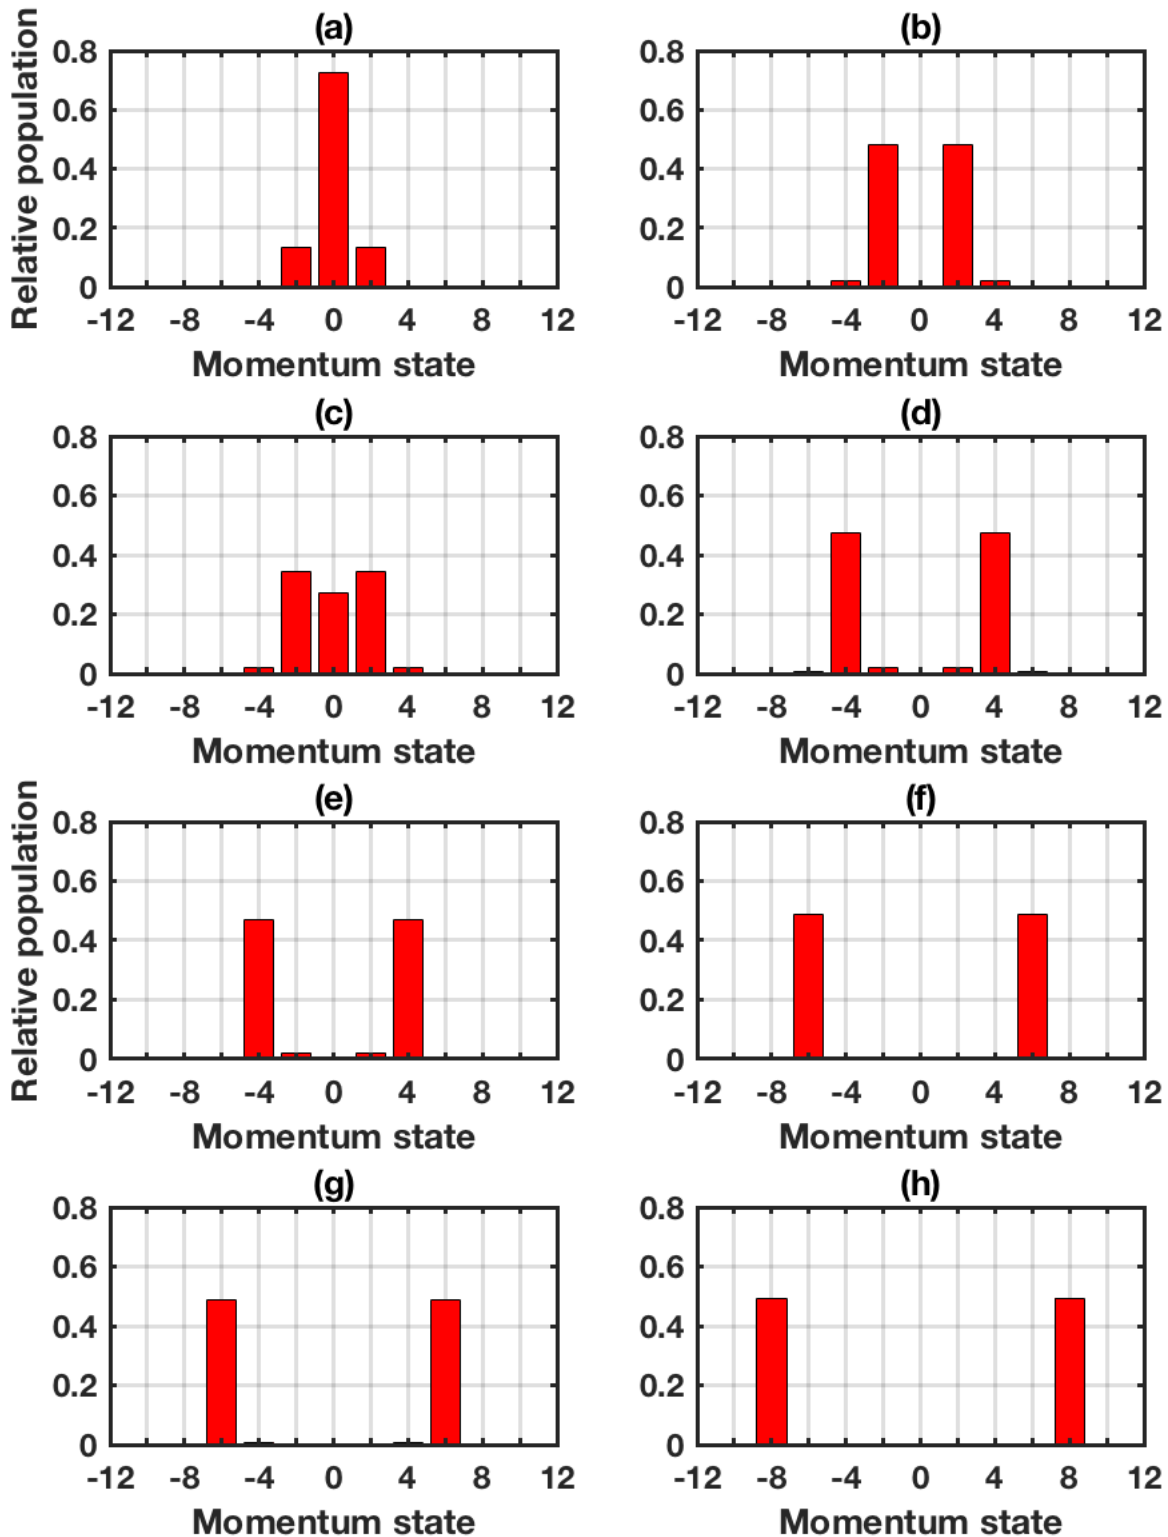


Figure 7.1: (a-h, in order of increasing band number from $r = 0$ to 7) Momentum state populations and energies for the first 8 Bloch states $|r\rangle$ for atoms trapped in a lattice with $V_0 = 10E_R$. Note that for $r > 0$ the states begin to resemble split states of higher and higher orders n .




Figure 7.2: Value of the projection integral in Eq. (7.1) versus the state $|r\rangle$ corresponding to the band r for splitting order $n = 1$ (red circles), $n = 2$ (black diamonds), and $n = 3$ (blue squares) for (a) $\theta = 0$ and (b) $\theta = \pi$. As the band number increases, the Bloch wavefunctions look more and more like the split states, with the relative phase θ between the $\pm 2n\hbar k_L$ momentum states equal to $\theta = 0$ (π) for even (odd) band numbers. Alternating bands have relative phases θ of 0 and π , depending on the band parity.




Figure 7.3: Band-to-band transition frequencies (in units of $\omega_R = E_R/\hbar$) versus lattice depth V_0 (in units of E_R) for the transition from band $r = 0$ to $r' = 1$ (red, solid), 2 (blue, dashed), 3 (gray, dotted), 4 (magenta, dash-dotted), and 5 (black). The band 6 transitions are almost degenerate with the band 5 transitions for the entire range of depths considered here, so the $r' = 6$ curve would completely overlap the $r' = 5$ curve.

Table 7.1: Band transition frequencies $f_{r,r'}$, $V_0 = 10E_R$. The single (double) asterisk marks frequencies with matrix element overlaps $M_c(M_s) > 0.1$ (see Fig. 7.6 and Sec. 7.3).

Band n	Band m	$f_{r,r'}$ (kHz)
0	1	17.89**
0	2	24.61*
0	3	58.14
0	4	58.25
0	5	121.19
1	2	6.72**
1	3	40.25*
1	4	40.36**
1	5	103.30
2	3	33.53**
2	4	33.64*
2	5	96.58
3	4	0.10
3	5	63.0*
4	5	62.9**

The next section will analytically explore the dynamics of atoms trapped in a shaken optical lattice. This case is somewhat difficult to solve analytically, but some basic results can be applied in the simple case of single-frequency shaking at low amplitudes. We verify these results numerically, showing that we can split the atom wavefunction with a phase difference of 0 or π , depending on our driving frequency. Experimental results verify the numerics. From this we gain some insight into how to restrict our optimization parameters and present the results of rapid optimization of higher-order splitting in section 7.5.

7.3 The analytics of phase modulation of an optical lattice

The Hamiltonian for a phase-modulated (shaken) lattice with general shaking function $\phi(t)$ is given in Eq. (2.16) and repeated here for clarity

$$H = \frac{p^2}{2m} + \frac{V_0}{2} \cos(2k_L x + \phi(t)). \quad (7.2)$$

For the specific case where $\phi(t) = \alpha \sin(\omega t)$, we can apply the Jacobi-Anger expansion to Eq. (7.2). Using this we can write the potential term $V(x, t)$ in Eq. (7.2) as

$$V(x, t) = V_0 \left\{ \cos(2k_L x) \left[J_0(\alpha)/2 + \sum_{k=1}^{\infty} J_{2k}(\alpha) \cos(2k\omega t) \right] - \sin(2k_L x) \sum_{k=1}^{\infty} J_{2k-1}(\alpha) \sin[(2k-1)\omega t] \right\}. \quad (7.3)$$

Equation (7.3) shows that we can decompose the phase modulation into a term representing the carrier (first term) and a set of amplitude modulation terms containing both sine and cosine terms. The strength of these amplitude modulation terms are given by the Bessel functions $J_k(\alpha)$ where α is the amplitude of the phase modulation. Because the sine terms are odd, they will drive transitions between opposite parity states while the even cosine terms will drive transitions between states with the same parity [50].

By taking the J_0 term in Eq. (7.3) as the bare Hamiltonian $H_0(x)$, we can write the rest of the terms as a perturbation $H'(x, t)$. Using Fermi's golden rule, we can then write down the

transition rate $\Gamma_{r,r'}$ from state $|r\rangle$ to $|r'\rangle$ resulting from shaking at a frequency ω as

$$\begin{aligned} \Gamma_{r,r'} &= \frac{2\pi}{\hbar} V_0^2 \sum_{k=1} \left[J_{2k}^2(\alpha) |\langle r' | \cos(2k_L x) | r \rangle|^2 \times \delta(E_{r,r'} - 2k\hbar\omega) \right. \\ &\quad \left. + J_{2k-1}^2(\alpha) |\langle r' | \sin(2k_L x) | r \rangle|^2 \times \delta(E_{r,r'} - (2k-1)\hbar\omega) \right] \end{aligned} \quad (7.4)$$

where $E_{r,r'} = \hbar\omega_{r,r'} = E_r - E_{r'}$ is the energy difference between states $|r\rangle$ and $|r'\rangle$. The transition rate $\Gamma_{r,r'}$ is governed by the squares of the Bessel functions $J_k^2(\alpha)$ (where α is the amplitude of the phase modulation) and magnitudes of the transition matrix elements $|M_{r,r'}^{(s)}|^2 = |\langle r' | \sin(2k_L x) | r \rangle|^2$ and $|M_{r,r'}^{(c)}|^2 = |\langle r' | \cos(2k_L x) | r \rangle|^2$.

7.4 Single-frequency splitting

This section consists of three parts. The first subsection will show first-order single-frequency shaking results via simulation. Next, we will show experimental evidence that supports the simulation results. Finally, we will discuss just how far we need to split to implement an accelerating lattice scheme.

7.4.1 Theory and simulation of single-frequency splitting

For a given value of α the amplitude of the Bessel functions $J_k(\alpha)$ dies off as k increases. For $\alpha \leq 0.3$ we can keep two terms, simplifying the potential in Eq. (7.3) to

$$\begin{aligned} V(x, t) &= V_0 \left[J_0(\alpha) \cos(2k_L x) / 2 \right. \\ &\quad \left. - J_1(\alpha) \sin(\omega t) \sin(2k_L x) + J_2(\alpha) \cos(2\omega t) \cos(2k_L x) \right]. \end{aligned} \quad (7.5)$$

As with Eq. (7.5) if we keep two terms in Eq. (7.4), we obtain

$$\begin{aligned} \Gamma_{r,r'} &= \frac{2\pi}{\hbar} V_0^2 \left[J_2^2(\alpha) |M_{r,r'}^{(c)}|^2 \delta(E_{r,r'} - 2n\hbar\omega) \right. \\ &\quad \left. + J_1^2(\alpha) |M_{r,r'}^{(s)}|^2 \delta(E_{r,r'} - (2n-1)\hbar\omega) \right]. \end{aligned} \quad (7.6)$$

From this we see that weak shaking of the lattice at $\omega_{r,r'} = 2\pi \times f_{r,r'}$ will drive transitions between Bloch states $|r\rangle$ and $|r'\rangle$ if they have opposite parity and shaking at $\omega_{r,r'}/2$ will drive transitions between $|r\rangle$ and $|r'\rangle$ if they have the same parity. In general, shaking at $\omega_{r,r'}/N$ will drive transitions between $|r\rangle$ and $|r'\rangle$ with parity determined by whether N is odd or even. This is in keeping with the results in Ref. [108] for the case of the amplitude-modulated lattice (where only like-parity transitions are allowed) and the phase-modulation results in Ref. [120]. The difference in our work is that we approach this problem from a standpoint of inducing band-to-band transitions to perform atom beamsplitting for interferometry.

The above analysis shows that if we begin in the ground state $|r = 0\rangle$ and shake at $\omega = \omega_{01}$ ($\omega_{02}/2$), we will drive odd (even) parity transitions between bands $r = 0$ and $r' = 1$ ($r' = 2$). We simulate both cases, solving the TDSE with the Hamiltonian in Eq. (7.2) with a single frequency ω and amplitude $\alpha = 0.3$ for $T \approx 1$ ms.¹

The band transition frequencies ω_{01} and $\omega_{02}/2$ are plotted in Fig. 7.4 versus the lattice depth. Given $V_0 = 10E_R$, results for odd parity shaking at $\omega = \omega_{01} = 2\pi \times 17.88$ kHz are shown in Fig. 7.5(a), and results for even parity shaking at $\omega_{02}/2 = 2\pi \times 12.3$ kHz are shown in 7.5(b). For the simulation results presented in Fig. 7.5, at each timestep we calculate the percent error relative to the split state as in Eq. (3.4) and the inner product between the simulated state $|\Phi(t)\rangle$ at time t and the first order split state $|\psi_1(p, \theta)\rangle$ as in Eq. (7.1).

We see that when the percent error is lowest in Fig. 7.5(a), the projection of the state $|\psi\rangle$ onto the split states is highest for the split state $|\psi_1(p, \theta = \pi)\rangle$. This shows that we are in fact splitting with relative phase $\theta = \pi$ between the two split arms. Conversely, in Fig. 7.5(b) we achieve splitting with $\theta = 0$. Thus, by controlling the shaking frequency we can control the relative phase between the two split arms.

For higher amplitudes first-order perturbation theory becomes less and less applicable, and we can no longer use Fermi's golden rule to accurately describe the physics of the situation. In

¹ Because we scale our time and energy scales to avoid excessively large and small numbers, we actually end up shaking for $T = 1.04$ ms.




Figure 7.4: The band transition frequencies ω_{01} (red, solid) and $\omega_{02}/2$ (blue, dashed) as a function of the lattice depth.




Figure 7.5: Results of shaking simulations for a lattice depth of $V_0 = 10E_R$. The percent error relative to the split state (red, see Eq. (3.4)), projection of the current state $|\psi\rangle$ onto the split state with $\theta = 0$ (blue, dashed, see Eq. (7.1)) and $\theta = \pi$ (black, dot-dashed), plotted versus shaking time for a shaking frequency of (a) ω_{01} and (b) $\omega_{02}/2$ and a shaking amplitude of $\alpha = 0.3$ rad.

this case we must keep more terms in the Jacobi-Anger sums of Eq. (7.5) and go to higher orders in perturbation theory. Furthermore, we cannot use this simple picture to obtain higher-order splitting. This is due to the fact that the matrix elements $|M_{r,r'}^{(c)}|^2$ and $|M_{r,r'}^{(s)}|^2$ become relatively small as we consider transitions from the state $|r = 0\rangle$ to higher-lying states with $|r > 2\rangle$. This is shown in Fig. 7.6(a) where for higher-band transitions the relevant matrix element is at least one order of magnitude below the lower-band transitions. As such, the transition strength is much lower and the transitions become less favorable.

However, we can make transitions from the ground state $|r = 0\rangle$ to an intermediate state in band $r' = 1$ or 2 and then to higher-lying states in bands $r'' > 2$. If we make these intermediate state transitions, the matrix elements become more favorable. This is shown in Fig. 7.6(b-c). As expected, the cosine transition matrix elements are strongest when making transitions between states in next-to-adjacent bands, but the sine matrix elements are strongest when making transitions between states in adjacent bands. Interestingly, when considering the sine matrix elements we see that it is also favorable to make transitions between states in bands $r = 1$ to $r' = 0, 2$ or 4 . We also observe that transitions between bands $r = 3$ and $r' = 2$ and 6 are favorable, but transitions between bands $r = 3$ and $r' = 4$ are not. This is possibly due to the avoided crossing between bands 3 and 4 at $q = 0$ (see Fig. 2.1). We will find in Sec. 7.5 that the strongest transitions in Fig. 7.6 can be used to simplify the optimization landscape for higher-order splitting in Sec. 7.5.

7.4.2 Experimental results

In this section we demonstrate that the splitting schemes described in the previous subsection and shown in Fig. 7.5 are viable experimentally. To do this we start with atoms loaded into the ground state of a lattice of (intentionally) unknown depth. We then shake the lattice for a time $T = 0.2$ ms with varying amplitude A at different frequencies. Unlike the results presented in Chapter 6 we use frequency modulation of the lattice laser to shake the lattice, as presented in Sec. 5.3.2. This is done to eliminate the effects of losses and spurious reflections due to the EOM.

We find that we can split the atom wavefunction to within $E \approx 10\%$ at frequencies cor-




Figure 7.6: (a) The matrix elements $|M_c|^2$ (red) and $|M_s|^2$ (black) for band-to-band transitions from bands $r = 0$ to r' plotted versus final band r' . (b-c) Plots of the matrix elements (b) $|M_c|^2$ and (c) $|M_s|^2$ for band-to-band transitions from bands r (x-axis) to r' (y axis). The colorbar on the right gives the magnitude of the transition matrix element.

responding to ω_{01} and $\omega_{02}/2$. This is shown in Fig. 7.7. As in Chapter 6 we do not generally obtain splitting to better than 10% due to spurious atoms detected in the $0\hbar k_L$ momentum state (e.g. due to heating and imaging noise), the finite momentum spread of the condensed atoms in the lattice [43, 101], and the finite signal-to-noise ratio in imaging. Simulations show that for our experimental momentum width of $0.6\hbar k_L$ we are limited to errors of about 4 – 8%, depending on our shaking frequency.

From the results shown in Fig. 7.7, we estimate the lattice depth to be $V_0 \approx 15.3E_R$. In this case the relevant band-to-band transition frequencies are $\omega_{01} = 21.7$ kHz and $\omega_{02} = 2 \times 17$ kHz. This not only confirms the simulation results from the previous subsection but provides us with a reliable way to approximate the lattice depth. In our current interferometry experiments the absolute lattice depth is less important than the day-to-day and shot-to-shot stability of the lattice depth. Thus, an approximate depth combined with the stability provided by a lattice laser intensity servo is sufficient for our purposes.

7.4.3 How much must we split?

Due to the fact that the single-frequency shaking does not work as well for higher-order splitting (see Fig. 7.6), we must work harder to obtain higher-order splitting. This is the subject of Sec. 7.5. However, before we dive into the next section it is instructive to demonstrate how much we must split the atoms to be able to accelerate them in a moving lattice.

If we truncate the Bessel function expansion of Eq. (7.5) to two terms and do some trigonometry, we obtain

$$V(x, t) = 2V_0(J_0(\alpha) \cos(2k_L x)/4 + J_1(\alpha)(\cos(2k_L x - \omega t) + \cos(2k_L x + \omega t))). \quad (7.7)$$

Equation (7.7) describes a carrier lattice and two counterpropagating moving lattices with velocity $v = \pm\omega/2k_L$. If we can split the atoms to a high enough order, we can trap the split atoms in one of the two moving lattices. The atoms will then move with the lattice if we accelerate and decelerate




Figure 7.7: Experimental results showing the percent overlap with the split state versus shaking amplitude A for frequencies of $17 \text{ kHz} \approx \omega_{02}/2$ (red), 21.5 kHz (blue), and 22 kHz (black), where $\omega_{01} \approx 21.7 \text{ kHz}$ for a lattice depth $V_0 \approx 15.3E_R$ (see Fig 7.4). The points show the mean of three separate experimental runs with the same shaking function and the error bars show the standard deviation of these points.

the lattice. This will allow us to obtain interferometry with T_1^3 sensitivity to an applied signal [74]. In this case the moving atoms will not be able to “see” the counterpropagating lattices and will thus not be affected by them.² In this case, the atoms in the positive (negative) momentum state will be trapped in the lattice moving with positive (negative) velocity. Then, if the lattices are accelerated by changing the shaking frequency such that the magnitude of the counterpropagating lattice velocity changes, the atoms should follow the lattices in which they are trapped. The atoms will thus accelerate as the lattices are accelerated, given that this is done slowly enough [30, 102].

We find that if we begin with atoms split to third order (that is, $\pm 6\hbar k_L$), we can shake the lattice at $\omega = 12\hbar k_L^2/m = 12\omega_R$ with $\alpha = 1$ such that the lattice is moving with $v = \pm 6\hbar k_L/m$ without disturbing the atom wavefunction appreciably. Here, $\omega_R = E_R/\hbar$ is the recoil frequency of the atoms in the lattice. In this case the atoms maintain their split state to within $\approx 1\%$, as shown in Fig. 7.8(a). Furthermore, simulations show that if the atoms begin in the $\pm 8\hbar k_L$ split state and are trapped in a lattice shaken at $\omega = 16\hbar k_L^2/m$, the atoms will continue to maintain their splitting to within 1%, as shown in Fig. 7.8(b). This trend continues for even higher splitting orders.

This analysis shows that if we can achieve third-order splitting we can then accelerate the atoms in the lattice with negligible perturbation. In the next section we will show how to optimize such splitting by shaking at frequencies corresponding to band-to-band transitions.

7.5 Optimization of higher-order splitting using band-to-band transitions

This section presents results of simulations optimizing splitting protocols up to $n = 3$. The optimization is done via the genetic algorithm described in Chapter 3. As in Chapter 4, the simulation solves the TDSE using the Hamiltonian in Eq. (7.2) as in Sec. 7.3 and the GA uses the error as given in Eq. (3.4) as a fitness function (and tries to minimize this error). We will compare results using the full frequency bandwidth up to the $r = 0 \rightarrow r' = 5'$ transition to optimize the lattice shaking to results where only single-photon band-to-band and two-photon

² The atoms moving with one of the lattices must be in an eigenstate of the lattice, but shaking can always be modified to prepare the split atoms so that they resemble the ground state of the moving lattice with depth $V_0 J_1(\alpha)$.




Figure 7.8: The percent error relative to the (a) third- ($n = 3$) and (b) fourth-order ($n = 4$) split state (red, see Eq. (3.4)), projection of the current state $|\Phi(t)\rangle$ onto the split state with $\theta = 0$ (blue, dashed, see Eq. (7.1)) and $\theta = \pi$ (black, dot-dashed), plotted versus shaking time for a shaking amplitude $\alpha = 1$ rad. In each case, there are two counterpropagating lattices moving at velocities $v = \pm 2n\hbar k_L/m$.

half-band transitions are used. By the term “half-band transitions,” we mean that we shake the lattice at a frequency $\omega_{r,r'}/2$ to cause even parity transitions. We know from Ref. [96] that off-resonant shaking in the presence of atom-atom interactions causes heating. Furthermore, to avoid the transverse scattering described in Refs. [23, 24] we want to shake at single- and two-photon band-to-band resonances so that no excess energy is available for transverse scattering. Therefore we wish to restrict our optimization to resonant transitions in order to limit the heating due to these factors. Note that as in the rest of this thesis our simulations do not take interactions into account.

In the simulations presented here we shook the lattice for $T \approx 0.5$ ms and attempted to obtain $\pm 2n\hbar k_L$ splitting for $n = 1, 2$, and 3. As in the rest of this work, to ensure smooth turn-on and turn-off of the shaking function, we multiplied each shaking function by the envelope function in Eq. 3.5. Due to the inherent randomness in the GA, we ran each class of simulations 10 separate times and took the best result for our analysis. We plot the best results (in terms of percent error) in Fig. 7.9.

We ran five different classes of optimization simulations. One class included every frequency in the band from DC up to $r = 0 \rightarrow r' = 5'$, another included only the 10 band transition frequencies in this region, and a third included the 10 half-band transition frequencies. These frequencies are tabulated in Table 7.1. A fourth simulation class included all 20 band and half-band transition frequencies, and a fifth chose only the 9 frequencies corresponding to appreciable (> 0.1) matrix element overlap in Fig. 7.6 (marked with asterisks in Table 7.1). We plot the lowest error achieved after 1000 iterations in Fig. 7.9. Note that convergence below $10^{-3}\%$ is limited by phase errors in the split step method.

From the results presented in Fig. 7.10, we see that in all cases we can split the atom wavefunction to better than 0.3%. We see in Fig. 7.9(a) that the error is lowest if we include all frequencies, but in this regime the interaction-induced heating (which is not present in our current simulation model) will be highest. For simulations restricted to the select strongest band transitions we can obtain splitting to better than 0.05%. While the use of more complex fitness

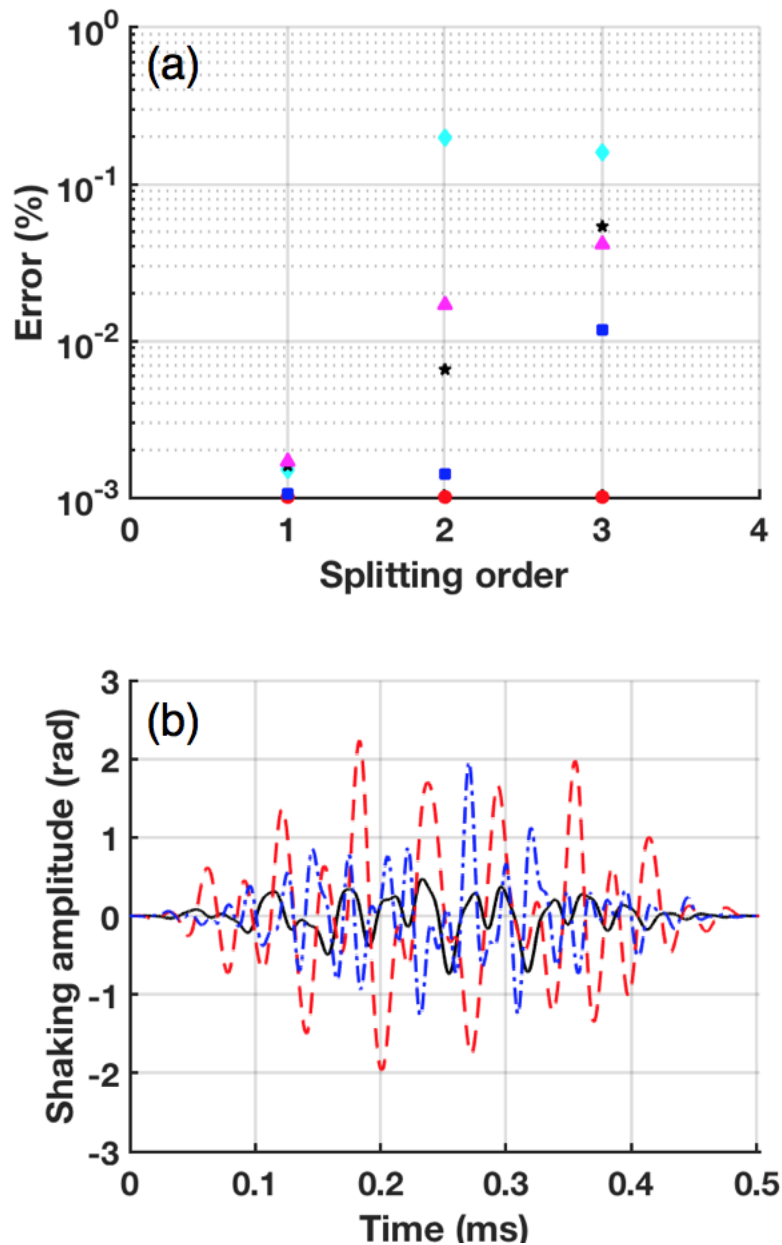


Figure 7.9: (a) Percent error vs. splitting order for $n = 1, 2$ and 3 for the five simulations considered in this work. The frequency ranges considered here are: band transition frequencies (cyan diamonds), half-band transitions (black stars), both band and half-band transitions (blue squares), frequencies with non-negligible matrix element overlap (magenta triangles), and all frequencies in the band (red dots). (b) The best optimized shaking functions for the select frequency case corresponding to the magenta triangles in (a) for $n = 1$ (black, solid), $n = 2$ (blue, dot-dashed), and $n = 3$ (red, dashed).

functions (e.g. the fitness function used for first-order splitting in Chapter 4) may further improve this splitting efficiency, we obtain good results by simply minimizing the error in Eq. (3.4). In summary, by truncating our search space from ≈ 50 frequencies (limited by the discrete temporal sampling inherent in the numerics) to ≈ 10 frequencies, we can still split with sufficiently low error.

Even though we cannot access the relative phase θ of the two split arms of these optimized split states experimentally, it is of interest to examine them in simulation in order to better understand the shaking dynamics. Thus, we plot the final phase θ of the optimized split state for the best results of each of the five simulation classes and three splitting orders in Fig. 7.10. We include the error from Fig. 7.9 for easy reference and comparison.

The results show that the phase dynamics of multi-frequency splitting are more complex than the simple model presented in Sec. 7.3 predicts. For example, when we split using the half-band transition frequencies we would expect that the phase θ be near zero. However, we find that this phase is closer to $\theta = -\pi$ for first-order splitting and $-\pi/2$ for higher-order splitting. Therefore, our simple model derived in the case of single-frequency shaking has broken down. Unfortunately it is analytically difficult to consider multiple frequency shaking due to the fact that neither the Jacobi-Anger expansion nor the results of Floquet analysis applies. Furthermore, as shown in Fig. 7.9(b) the shaking function amplitude is about an order of magnitude greater than that used to justify the truncation of the sum in Eq. (7.3) and apply first-order perturbation theory.

We can, however, make some general inferences from our simple model. The fact that $|\theta| \approx \pi$ for the first-order split state likely comes as a result of the fact that the two-photon matrix element $|M_{0,2}^{(c)}|^2$ is about a factor of two lower than the single-photon element $|M_{0,1}^{(s)}|^2$ connecting $r = 0$ and $r' = 1$. Thus, transitions between the odd-parity states are more favorable, as can be seen from Fig. 7.5 where the dynamics of shaking at ω_{01} are far less complex than those of shaking at $\omega_{02}/2$.

For higher-lying states, the multiple pathways for an atom to get from $|r = 0\rangle$ to the higher-lying states means that the even and odd parity transitions tend to interfere, and the split state will in general be a mixture of the two bands. From Fig. 2.1 we see that these higher-lying states $|r' > 2\rangle$ corresponding to splitting with $n > 1$ are nearly degenerate. This interference causes the

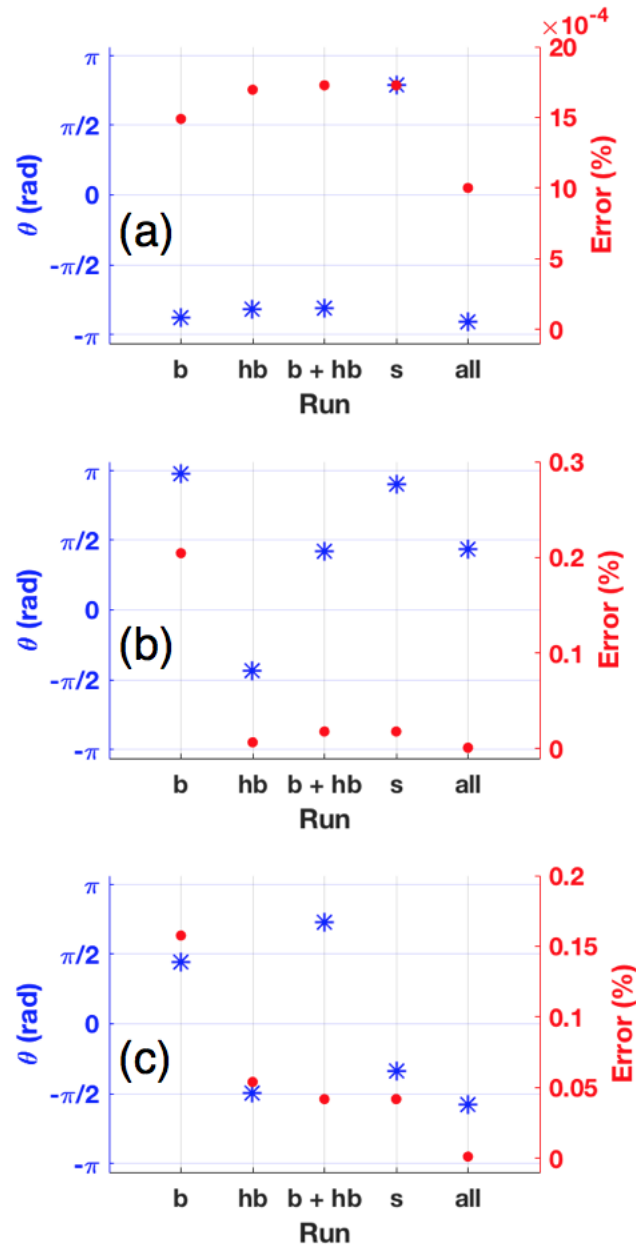


Figure 7.10: Phase θ between the split arms (blue asterisks, left axis) and error (red dots, right axis) for (a) $n = 1$, (b) $n = 2$, and (c) $n = 3$. The best of ten runs for each of the five simulation classes is shown. The classes are labeled on the x-axis as follows: band transition frequencies (b), half-band transitions (hb), both band and half-band transitions (b+hb), select frequencies with non-negligible matrix element overlap (s), and all frequencies in the band (all)

phase difference between the two split arms to average to $|\theta| = \pi/2$. In these cases the optimized split state is not dominated by population transfer into a single higher-lying band but rather this state is composed of a mixture of states corresponding to two nearly degenerate bands.

From a purely experimental standpoint our results dramatically simplify the optimization landscape that we must explore. This allows for sufficient error minimization within a low number of iterations. That is, even though simulations with more frequencies tend to converge to lower errors, they take longer to do so. For example, if we run 10 simulations to optimize splitting of the atom wavefunction with the select transition frequencies indicated in Table 7.1, for $n = 1, 2$, and 3, we can achieve convergence to better than 1% error in (on average) 1, 10, and 28 iterations, respectively. In each case, the error for the simulations with all frequencies in the band is $> 70\%$, as shown in Fig. 7.11. Figure 7.11 also shows that if we start with atoms in the $n = 2$ split state, we can optimize transfer into the $n = 3$ split state within 1% within 10 iterations. In this case the total splitting time will double, but optimization of $6\hbar k_L$ splitting is possible with fewer than 20 total iterations. This learning speedup is extremely important for practical implementations of shaken lattice interferometry in that optimization happens more quickly and effectively, allowing for fast optimization of the interferometer sequence.




Figure 7.11: Results of genetic optimization showing the mean (open markers) and best (closed markers) error for splitting with (a) $n = 1$, (b) $n = 2$ and (c) $n = 3$ versus number of iterations of the optimization algorithm. The black points indicate optimizations with all frequencies in the band from DC to 121 kHz, and the red points indicate optimizations using only the truncated frequency space corresponding to the select band-to-band transitions indicated in Table 7.1. The error improvement is much faster with the truncated subspace.

In the experiment, if the lattice depth is known (e.g. via the measurement scheme described in Sec. 7.3 or in Ref. [50, 101]) we can restrict our shaking to the transitions with appreciable transition matrix elements. Then a closed-loop algorithm (e.g. the dCRAB algorithm used to obtain the results in Chapter 6) can be used to efficiently optimize the splitting protocol. Thus we have found a reduced subspace that allows for faster optimization of the system and reduces heating due to off-resonant shaking [96].

7.6 Conclusion

In conclusion we demonstrate a simple means of using the band-to-band transitions to implement an atom beamsplitter in an optical lattice. We develop a theoretical model that allows us to use a single shaking frequency to implement low-order splitting. However, at higher frequencies our simple model breaks down and we must incorporate multiple frequencies in order to obtain good splitting. Due to heating caused by atom-atom interactions it is of interest to restrict our shaking frequencies to those resonant with single- and two-photon transitions between bands. We show that we can obtain higher-order splitting up to $n = 3$ with an error $< 0.1\%$ by optimizing shaking with a learning algorithm using a reduced subspace of frequencies corresponding to the strongest band and half-band transition resonances. This simplification of the optimization landscape allows for faster optimization with less deleterious heating effects due to atom-atom interactions. Finally, we show that higher-order splitting can be implemented by accelerating the atoms in the optical lattice and can potentially allow for interferometry with sensitivity that scales with the cube of the interrogation time. This opens up potential new pathways for improving and expanding interferometry using atoms trapped in a shaken optical lattice. Future work that will advance this goal will be described in Chapter 8.

Chapter 8

Shaken lattice interferometry in the future

At the end of nearly four years working on shaken lattice interferometry, it's quite clear that the experiment that has been built has produced some interesting results. This thesis has by and large covered a lot of the initial proof-of-principle tests. We have done simulations that showed what this system could theoretically be capable of. Furthermore, we have produced experimental results showing that the predictions made by our simulation work are possible (at least at the level at which we have explored it).

The point of this chapter is simply to speculate on the future. Section 8.1 speculates on the experimental improvements that could be made to the experimental system and how these improvements would subsequently improve the results. Section 8.2 will discuss briefly the improvements that could be made to the simulation. This section will also outline the basics of a semiclassical simulation of a shaken lattice gyroscope, the next step in shaken-lattice-based inertial sensing. Section 8.3 present some final thoughts on the future of SLI.

The idea here is not to dictate what another student would do but rather to offer some suggestions. What actually ends up happening in the future is a unique combination of student and advisor interests and personalities as well as funding opportunities. That, in a lot of ways, is what's so fun about doing science, especially in JILA.

8.1 Future experimental improvements

The most obvious experimental improvements that could be made are those that improve the interferometer interrogation time. Higher interrogation times correspond to higher sensitivity. Furthermore, when implementing an AC-signal-sensitive reciprocal interferometer, longer interrogation times mean that we are sensitive to lower-frequency signals.

Interrogation time improvements can proceed in a number of ways. Atom stability in the lattice can be improved by the use of common lattice intensity stabilization techniques. This would allow for longer interrogation times limited only by photon scattering rates and collisions with background particles. Markus Greiner’s group at Harvard has been able to hold lithium atoms in an optical lattice on the order of a minute [12]. The most remarkable difference between the Harvard experiment and ours is the use of a high-powered 1064 nm laser and two-stage lattice beam intensity servo in that experiment. The first stage of the servo system is AOM-based and high bandwidth. The second is lower-bandwidth and based on a polarization rotator. The two systems act independently.

Our single-stage servo system is reasonably high bandwidth but has issues with ramping the lattice on and off.¹ Furthermore, as we ramp off the magnetic pre-lattice trap, the trap minimum tends to move towards the chip as the currents are ramped off, as shown in Fig. 8.1. Thus, we are loading the atoms into the lattice in a suboptimal manner. This in combination with the issues involved with ramping the chip trap off, results in somewhat of a “throw and catch” scheme being used for atom transfer from the magnetic trap into the lattice. This scheme imparts a non-zero quasimomentum to the atoms, and this can change the relative phase θ of the two split states.² Therefore, it is worthwhile to spend some time studying the magnetic trap to lattice trap transfer scheme and optimizing it.

One could even abandon the chip altogether and move towards an all-optical approach based

¹ It’s hard to servo around zero intensity.

² If the lattice is optimally aligned, this quasimomentum is very small (perhaps 10% of $\hbar k_L$). This presents experimentally as asymmetry in the population in the $\pm 2\hbar k_L$ side clouds visible when the atoms are in the ground band of the lattice. See Fig. 2.3(b).




Figure 8.1: Location below the chip of the atom chip trap as the pre-lattice waveguide trap is ramped off over an arbitrary ramp time (on the x -axis). As the chip trap is ramped off, the trap location moves rapidly towards the chip.

on BEC production using a crossed dipole trap [7,9,25]. In this way one could design optical access in all three lattice dimensions using a simple vacuum cell geometry with a clear aperture where the chip once was. There is, as always, the possibility of this being simply a “grass is greener on the other side” situation, as making a BEC in a dipole trap offers its own set of challenges. If one wishes to stay with the chip-based geometry, then they should carefully analyze the transfer from the chip into the lattice to mitigate excitations as much as possible.

From the results shown in Chapters 6 and 7 we can see that one of the main factors limiting convergence below 10% error is the finite momentum width of the condensate. Thus we must work to narrow this width in future experiments. One solution to this problem is to perform most of the evaporation in the tight atom chip trap, then move to a looser decompressed trap (like the waveguide trap used before we ramp on the lattice beams) for the final stage of evaporation. If this is done adiabatically, the condensate wavefunction will follow the trap deformation and the looser trap will result in a narrower momentum distribution. A scheme like delta-kick cooling [5,84] could also be an alternative means of narrowing the momentum width.

In order to make a six-axis inertial sensor, we need to be able to interrogate the atoms along each of the three Cartesian dimensions. One of the advantages of an all-optical scheme would be the relative ease of construction of a three-dimensional lattice scheme. However, even in the absence of such a scheme, a three-dimensional optical lattice could be constructed through the use of on-chip optics [107] and windows in the chip [15,17,99,100]. One issue that the next student would need to overcome is the issue of optimizing the z -axis accelerometer in the presence of gravity to make meaningful measurements of applied signals along this axis. This should be possible using the DC bias optimization methods discussed in Chapters 3 and 4.

We note that it would be reasonably straightforward (if somewhat difficult due to limited space) to implement a second lattice dimension along the direction orthogonal to the current lattice axis and gravity in the current experimental scheme. This would allow for studies of how the atom-atom interactions affect the shaken lattice interferometer, as one could prepare in one dimension a deep lattice state with a low occupation number on each site [42,43,44,101] and a superfluid

state (characterized by the distinct momentum peaks we rely on for interferometry) along the other dimension. This would result in a 1-D array of one-dimensional, low atom number interferometers operating in parallel.

Finally, there are some other obvious improvements. Chip and coil drivers with lower noise and improved temperature stability would result in a more stable and repeatable experimental system. The most convenient result of such improvements would be the mitigation or elimination of trap bottom drift due to drift in the magnetic fields, as discussed in Sec. 5.2. As another matter of convenience, one could even implement a set of computer-controlled piezo mirror mounts that optimize the lattice alignment every morning (and tweak said alignment as the system drifts over the course of the day). A more stable laser system would allow for the experiment to be run continuously without the need for graduate student input. In a day and age when a number of atomic physics experiments run continuously and can be controlled even remotely, we should work towards similar levels of control and operation.³ Achieving this would allow one to get better data faster.

8.2 Future simulation improvements

A number of refinements could be made to the simulation to better understand the effects of the most obvious thing that we ignore: atom-atom interactions. As discussed in Sec. 4.5 the simplest refinement would be to include mean-field effects through the use of the Gross-Pitaevskii equation [90], which simply adds a mean-field term to the Schrödinger equation. This can be easily solved using our split-step method. More complex refinements (e.g. due to finite temperature effects [45]) would require the use of more complex formulations and will not be considered here.

Multi-dimensional simulations with the split-step method are straightforward as long as there is no crosstalk between the different lattice axes, rendering the equations that govern the different axes separable. When considering shaken lattice accelerometry this is the case. However, when

³ Thanks to M. Pigneur for showing us how her experiment could be controlled remotely, inspiring waves of jealousy.

considering a rotation sensor we will see that the two axes are no longer separable. Thus, we will have to resort to different methods. One such method is described in Sec. 8.2.1.

8.2.1 A shaken lattice gyroscope

As initially discussed in Sec. 6.4, in order to build a six-axis inertial sensor we need to configure our shaken lattice interferometer as a gyroscope. Taking a clue from the Sagnac effect discussed in Chapter 1 we can enclose an area by using a two-dimensional optical lattice. To do this we perform the conventional Michelson interferometry sequence in one dimension while simultaneously shaking along the second dimension to move the atoms out and back along this dimension. A cartoon of this is shown in Fig. 8.2.

8.2.2 A semiclassical Hamiltonian

In order to simulate the effects of a constant rotation Ω on the system, we can write down the semiclassical Hamiltonian for a non-interacting atom moving in a rotating frame. This gives

$$H = \frac{p_x^2}{2m} + \frac{p_y^2}{2m} - \frac{V_0}{2} \cos(2k_L x + \phi_x(t)) - \frac{V_0}{2} \cos(2k_L y + \phi_y(t)) + \frac{1}{2} m \Omega^2 (x^2 + y^2) + \Omega (x p_y + y p_x) \quad (8.1)$$

Here, we see the typical shaking lattice term in both the x and y dimensions. The next terms are the centrifugal terms proportional to x^2 and y^2 .⁴ Finally, we have the Coriolis terms proportional to $x p_y$ and $y p_x$. It is these terms that lead to the inseparability of Eq. (8.1) and thus render our split-step method useless. Thus, we need another simulation method. A possible alternative is described in Sec. 8.2.3.

8.2.3 Towards simulations of a shaken lattice gyroscope

While the split-step method will not work for our simulations of a shaken lattice gyroscope, we can use the Chebyshev method [62, 65, 86, 111] to solve the TDSE with the Hamiltonian in

⁴ One of the problems with these terms is that they are quadratic in position. Thus for large enough rotation rates, this background harmonic potential will overwhelm the lattice potential and the system will no longer be quantized in momentum space.

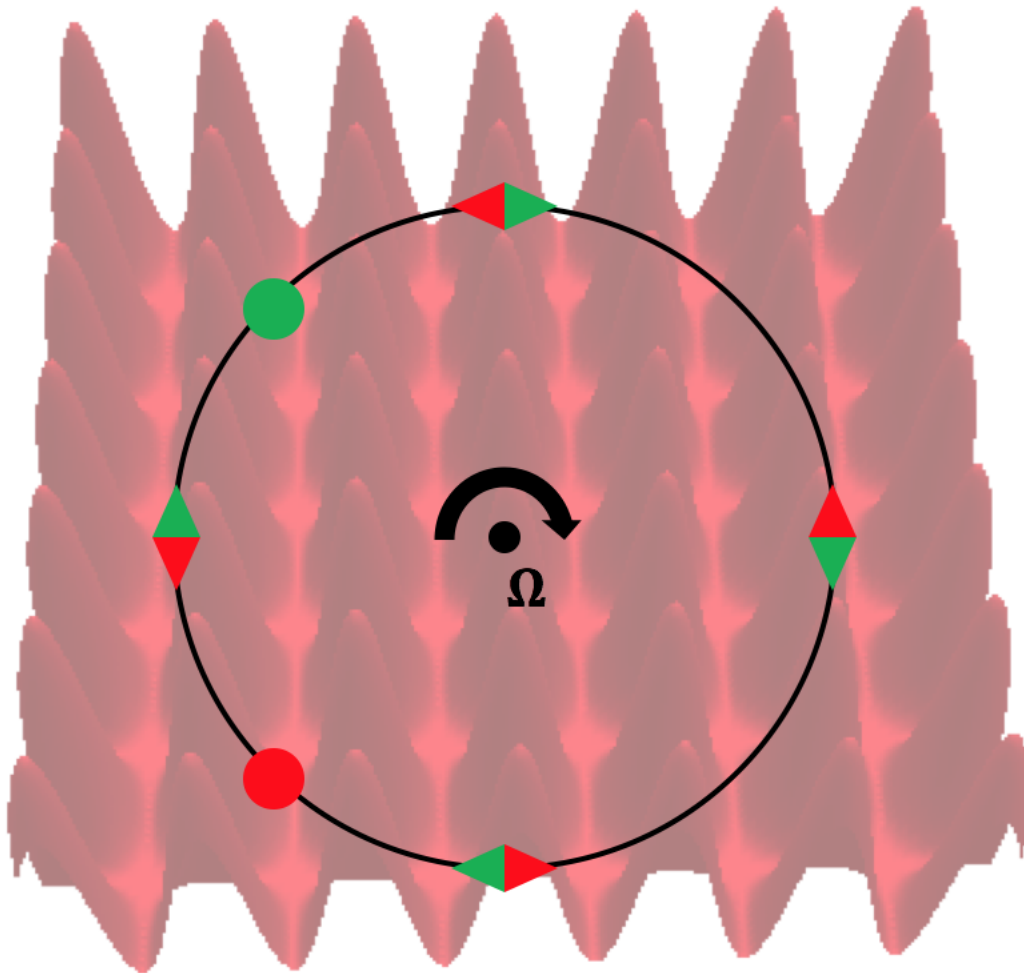


Figure 8.2: A cartoon of an atom-based Sagnac gyroscope in a two-dimensional lattice. The atoms (depicted as red and green circles) enclose an area as they undergo the conventional Michelson atom interferometry sequence in one dimension while moving out and back in the orthogonal dimension.

Eq. (8.1). This method uses the Chebyshev polynomials to construct an approximation to the Hamiltonian propagator $\exp(iHt/\hbar)$ and thus solve the TDSE, which can be written as

$$\psi(t) = \exp(iHt/\hbar)\psi(0) \quad (8.2)$$

The Chebyshev polynomials⁵ are useful in that they exist in the domain $[-1, 1]$ and range from $[-1, 1]$. Thus, the maximum error that they can obtain is lower than any other set of polynomials with higher ranges. Most usefully, there is no restriction on the separability of the Hamiltonian H , so this method can be used to solve more general problems than the split operator method we used in this work.

We are currently working to implement this method, but the citations provided above should allow the interested reader to write their own implementation of the algorithm.

8.2.4 Experimental rotation sensing

The shaken lattice gyroscope is one of those rare experiments that is (currently) easier to simply carry out the experiment than it is to simulate. We have a Precitech model RT200 air bearing rotation table in the lab that can rotate the system while it is being interrogated. For past experiments, we rotated the table in a (reasonably well) controlled manner by attaching a neodymium magnet to a corner. We oriented a coil of wire so that the magnetic field lines from the Nd magnet ran through the area enclosed by the coil. We could force the table to move simply by pulsing a current through the coil.

Once the experiment is mounted on the rotation table and a two-dimensional lattice is set up, one could optimize the Sagnac scheme described in the previous sections. After optimization the rotation table can be used to calibrate and test the inertial sensor just as the Michelson accelerometer scheme was tested in Chapter 6. Sensitivity scaling can be measured using the classical Fisher information just as in the accelerometer scheme. As with the accelerometer work presented

⁵ Again, Wikipedia provides a wonderful general reference for the properties of these polynomials, but most of what is relevant can be found in the references cited in the previous sentence.

here, we expect quadratic scaling of the gyroscope sensitivity with respect to interrogation time, based on our simple model described in Chapter 1.

8.3 Final thoughts

Experimentally, there are some interesting next steps that could be taken once the necessary improvements are made. The optimization of higher-order splitting is an interesting one that could lead towards the continuous acceleration lattice described in Sec. 7.4.3; this scheme would give another factor of the interrogation time T_I in the sensitivity scaling [74]. A longer interrogation time would allow us to improve on the results presented in Chapter 6 and perhaps even optimize a reciprocal interferometer that could then be used to measure an AC acceleration signal.

Furthermore, we could abandon entirely the idea of modeling the interferometer on the typical Michelson-style scheme that has been the focus of this work. Since our sensitivity measure is the classical Fisher information (CFI) we could use this as our optimization metric. As stated in Chapter 6 we have already begun preliminary studies of this method in both simulation in experiment. This is reasonably straightforward to implement in both cases (although experimentally we need at least two different data points to be able to calculate the required derivatives in Eq. (4.10)). We based our interferometer scaling initial assumptions on analogous light-pulse atom interferometry systems [74], and building a multi-path interferometer optimized with the CFI could allow for higher sensitivity scalings. Experimental and theoretical analyses of this scheme could produce some very interesting results.

I have thoroughly enjoyed working on shaken lattice interferometry for the last few years. Even if it does not completely revolutionize the field of atom interferometry and we do not all become terribly rich and famous,⁶ I can only hope that some of the concepts and ideas that have been presented in this work are useful in other atom interferometry, atomic physics, and scientific research in the future. And to you, reader, for reading this far, I thank you.⁷

⁶ and we probably won't

⁷ Now go get some research done. What are you doing reading the concluding paragraphs of a PhD thesis anyway?!

Bibliography

- [1] B. P. Abbott et al. Observation of gravitational waves from a binary black hole merger. Phys. Rev. Lett., 116:061102, 2016.
- [2] B. P. Abbott et al. GW170814: A three-detector observation of gravitational waves from a binary black hole coalescence. Phys. Rev. Lett., 119:141101, 2017.
- [3] B. P. Abbott et al. GW170817: Observation of gravitational waves from a binary neutron star inspiral. Phys. Rev. Lett., 119:161101, 2017.
- [4] A. Alberti, V. V. Ivanov, G. M. Tino, and G. Ferrari. Engineering the quantum transport of atomic wavefunctions over macroscopic distances. Nat. Phys., 5:547, 2009.
- [5] H. Ammann and N. Christensen. Delta kick cooling: A new method for cooling atoms. Phys. Rev. Lett., 78:2088, 1997.
- [6] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of Bose-Einstein condensation in a dilute atomic vapor. Science, 269:198, 1995.
- [7] K. J. Arnold and M. D. Barrett. All-optical Bose-Einstein condensation in a 1.06 μm dipole trap. Opt. Comm., 284:3288, 2011.
- [8] B. Barrett, A. Bertoldi, and P. Bouyer. Inertial quantum sensors using light and matter. Phys. Scr., 91:053006, 2016.
- [9] M. D. Barrett, J. A. Sauer, and M. S. Chapman. All-optical formation of an atomic Bose-Einstein condensate. Phys. Rev. Lett., 87, 2001.
- [10] P. Berg, S. Abend, G. Tackmann, C. Schubert, E. Giese, W. P. Schleich, F. A. Narducci, W. Ertmer, and E. M. Rasel. Composite-light-pulse technique for high-precision atom interferometry. Phys. Rev. Lett., 114:063002, 2015.
- [11] N. Beverini, A. Di Virgilio, J. Belfi, A. Ortolan, K. U. Schreiber, A. Gebauer, and T. Klügel. High-accuracy ring laser gyroscopes: Earth rotation rate and relativistic effects. J. Phys.: Conf. Ser., 723:012601, 2016.
- [12] S. Blatt, A. Mazurenko, M. F. Parsons, C. S. Chiu, F. Huber, and M. Greiner. Low-noise optical lattices for ultracold ${}^6\text{Li}$. Phys. Rev. A, 92:021402(R), 2015.
- [13] T. C. Briles, D. C. Yost, A. Cingöz, J. Ye, and T. R. Schibli. Simple piezoelectric-actuated mirror with 180 kHz servo bandwidth. Opt. Express, 18:10, 2010.

- [14] R. Bücker, T. Berrada, S. van Frank, J-F. Schaff, T. Schumm, J. Schmiedmayer, G. Jäger, J. Grond, and U. Hohenester. Vibrational state inversion of a Bose-Einstein condensate: optimal control and state tomography. J. Phys. B: At. Mol. Opt. Phys., 46:104012, 2013.
- [15] S. C. Caliga. Experimental realization of atomtronic circuit elements in non-equilibrium ultracold atomic systems. PhD thesis, University of Colorado at Boulder, 2016.
- [16] S. C. Caliga, C. J. E. Straatsma, and D. Z. Anderson. Transport dynamics of ultracold atoms in a triple-well transistor-like potential. New J. Physics, 18:025010, 2016.
- [17] S. C. Caliga, C. J. E. Straatsma, and D. Z. Anderson. Experimental demonstration of an atomtronic battery. New J. Physics, 19:013036, 2017.
- [18] S. L. Campbell, R. B. Hutson, G. E. Marti, A. Goban, N. Darkwah Oppong, R. L. McNally, L. Sonderhouse, J. M. Robinson, W. Zhang, B. J. Bloom, and J. Ye. A Fermi-degenerate three-dimensional optical lattice clock. Science, 358:90, 2017.
- [19] T. Caneva, T. Calarco, and S. Montangero. Chopped random-basis quantum optimization. Phys. Rev. A, 84:022326, 2011.
- [20] B. Canuel, F. Leduc, D. Holleville, A. Gauguier, J. Fils, A. Viridis, A. Clairon, N. Dimarcq, Ch. J. Borde, A. Landragin, and P. Bouyer. Six-axis inertial sensor using cold-atom interferometry. Phys. Rev. Lett., 97:010402, 2006.
- [21] O. Carnal and J. Mlynek. Young’s double-slit experiment with atoms: A simple atom interferometer. Phys. Rev. Lett., 66(21):2689, 1991.
- [22] R. Charrière, M. Cadoret, N. Zahzam, Y. Bidel, and A. Bresson. Local gravity measurement with the combination of atom interferometry and Bloch oscillations. Phys. Rev. A, 85:013639, 2012.
- [23] S. Choudhury and E. J. Mueller. Stability of a Bose-Einstein condensate in a driven optical lattice: Crossover between weak and tight transverse confinement. Phys. Rev. A, 92:063639, 2015.
- [24] S. Choudhury and E. J. Mueller. Transverse collisional instabilities of a Bose-Einstein condensate in a driven one-dimensional lattice. Phys. Rev. A, 91:023624, 2015.
- [25] J.-F. Clément, J.-P. Brantut, M. Robert de Saint-Vincent, R. A. Nyman, A. Aspect, T. Bourdel, and P. Bouyer. All-optical runaway evaporation to Bose-Einstein condensation. Phys. Rev. A, 79:061406(R), 2009.
- [26] A. D. Cronin, J. Schmiedmayer, and D. E. Pritchard. Optics and interferometry with atoms and molecules. Rev. Mod. Phys., 81:1051, 2009.
- [27] J. Dalibard and C. Cohen-Tannoudji. Laser cooling below the Doppler limit by polarization gradients: simple theoretical models. J. Opt. Soc. Am. B, 6:11, 1989.
- [28] K. B. Davis, M.-O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-Einstein condensation in a gas of sodium atoms. Phys. Rev. Lett., 75:3969, 1995.

- [29] S. M. Dickerson, J. M. Hogan, A. Sugarbaker, D. M. S. Johnson, and M. A. Kasevich. Multi-axis inertial sensing with long-time point source atom interferometry. Phys. Rev. Lett., 111:083001, 2013.
- [30] B. A. Dinardo and D. Z. Anderson. A technique for individual atom delivery into a crossed vortex bottle beam trap using a dynamic 1d optical lattice. Rev. Sci. Instr., 87:123108, 2016.
- [31] P. Doria, T. Calarco, and S. Montangero. Optimal control technique for many-body quantum dynamics. Phys. Rev. Lett., 106:190501, 2011.
- [32] S. Du. Atom-chip Bose-Einstein condensation in a portable vacuum cell. PhD thesis, University of Colorado at Boulder, 2005.
- [33] I. Dutta, D. Savoie, B. Fang, B. Venon, C. L. Garrido Alzar, R. Geiger, and A. Landragin. Continuous cold-atom inertial sensor with 1 nrad/sec rotation stability. Phys. Rev. Lett., 116:183003, 2016.
- [34] A. Eckardt. Colloquium: Atomic quantum gases in periodically driven optical lattices. Rev. Mod. Phys., 89:011004, 2017.
- [35] K. Eckert, P. Hyllus, D. Bruß, U. V. Poulsen, M. Lewenstein, C. Jentsch, T. Müller, E. M. Rasel, and W. Ertmer. Differential atom interferometry beyond the standard quantum limit. Phys. Rev. A, 73:013814, 2006.
- [36] M. Egorov, B. Opanchuk, P. Drummond, B. V. Hall, P. Hannaford, and A. I. Sidorov. Measurement of s-wave scattering lengths in a two-component Bose-Einstein condensate. Phys. Rev. A, 87:053614, 2013.
- [37] S. Ezekiel. Optical gyroscope options: Principles and challenges. Optical Fiber sensors 2006, 2006.
- [38] D. M. Farkas, K. M. Hudek, E. A. Salim, S. R. Segal, M. B. Squires, and D. Z. Anderson. A compact, transportable, microchip-based system for high repetition rate production of Bose-Einstein condensates. Appl. Phys. Lett., 96:093102, 2010.
- [39] D. M. Farkas, E. A. Salim, and J. Ramirez-Serrano. Production of rubidium Bose-Einstein condensates at a 1 Hz rate. 2014.
- [40] M. D. Feit, J. A. Fleck, and A. Steiger. Solution of the Schrödinger equation by a spectral method. J. Comp. Phys., 47:412, 1982.
- [41] F. Frank, T. Uden, J. Zoller, R. S. Said, T. Calarco, S. Montangero, B. Naydenov, and F. Jelezko. Autonomous calibration of single spin qubit operations. 2017.
- [42] M. Greiner. Ultracold quantum gases in three-dimensional optical lattice potentials. PhD thesis, Ludwig-Maximilians Universität München, 2003.
- [43] M. Greiner, I. Bloch, O. Mandel, T. W. Hänsch, and T. Esslinger. Bose-Einstein condensates in 1D- and 2D optical lattices. Appl. Phys. B, 73:769, 2001.
- [44] M. Greiner, O. Mandel, Tilman Esslinger, T. W. Hänsch, and I. Bloch. Quantum phase transition from a superfluid to a Mott insulator in a gas of ultracold atoms. Nature, 413:40, 2002.

- [45] A. Griffin, T. Nikuni, and E. Zaremba. Bose-Condensed Gases at Finite Temperatures. Cambridge University Press, 2009.
- [46] R. Grimm, M. Weidemüller, and Y.B. Ovchinnikov. Optical dipole traps for neutral atoms. Adv. in Atom., Mol., and Opt. Phys., 42:95, 2000.
- [47] T. L. Gustavson, P. Bouyer, and M. A. Kasevich. Precision rotation measurements with an atom interferometer gyroscope. Phys. Rev. Lett., 78:2046, 1997.
- [48] S.A. Haine. Mean-field dynamics and Fisher information in matter wave interferometry. Phys. Rev. Lett., 116:230404, 2016.
- [49] P. Hamilton, M. Jaffe, J. M. Brown, L. Maisenbacher, B. Estey, and H. Müller. Atom interferometry in an optical cavity. Phys. Rev. Lett., 114:100405, 2015.
- [50] J. Hecker-Denschlag, J. E. Simisarian H. Häffner, C. McKenzie, A. Browaeys, D. Cho, K. Helmerson, S. L. Rolston, and W. D. Phillips. A Bose-Einstein condensate in an optical lattice. J. Phys. B: At. Mol. Opt. Phys., 35:3095, 2002.
- [51] M. J. Holland and K. Burnett. Interferometric detection of optical phase shifts at the Heisenberg limit. Phys. Rev. Lett., 71(9):1355, 1993.
- [52] M. Holthaus. Floquet engineering with quasienergy bands of periodically driven optical lattices. J. Phys. B: At. Mol. Opt. Phys., 49:013001, 2015.
- [53] K. M. Hudek. A Compact and Transportable Ultracold Matter System and Progress Towards a Continuously Operating Neutral Rydberg Atom Quantum Computer. PhD thesis, University of Colorado at Boulder, 2012.
- [54] V. V. Ivanov, A. Alberti, M. Schioppo, G. Ferrari, M. Artoni, M. L. Chiofalo, and G. M. Tino. Coherent delocalization of atomic wave packets in driven lattice potentials. Phys. Rev. Lett., 100:043602, 2008.
- [55] J. Javanainen and M. Wilkens. Phase and phase diffusion of a split Bose-Einstein condensate. Phys. Rev. Lett., 78:4675, 1997.
- [56] N. B. Jørgenson, M. G. Bason, and J. F. Sherson. One- and two-qubit quantum gates using superimposed optical-lattice potentials. Phys. Rev. A., 89:032306, 2014.
- [57] R. S. Judson and H. Rabitz. Teaching lasers to control molecules. Phys. Rev. Lett., 68:1500, 1992.
- [58] M. Kasevich and S. Chu. Atomic interferometry using stimulated Raman transitions. Phys. Rev. Lett., 67:181, 1991.
- [59] D. W. Keith, C. R. Eckstrom, Q. A. Turchette, and D. E. Pritchard. An interferometer for atoms. Phys. Rev. Lett., 66(21):2693, 1991.
- [60] W. Ketterle, D. S. Durfee, and D. M. Stamper-Kurn. Making, probing and understanding Bose-Einstein condensates. In M. Inguscio, S. Stringari, and C. E. Wieman, editors, Proceedings of the International School of Physics "Enrico Fermi", Course CXL, page 67. IOS Press, 1999.

- [61] J. Kitching, S. Knappe, and E. A. Donley. Atomic sensors—a review. IEEE Sensors Journal, 11(9):1749, 2011.
- [62] R. Kosloff. Propagation methods for quantum molecular dynamics. Annu. Rev. Phys. Chem., 45:145, 1994.
- [63] 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, 528:530, 2015.
- [64] M. Krenn, M. Malik, R. Fickler, R. Lapkiewicz, and A. Zeilinger. Automated search for new quantum experiments. Phys. Rev. Lett., 116:090405, 2016.
- [65] C. Leforestier, R.H. Bisseling, C. Cerjan, M.D. Feit, R. Friesner, A. Guldberg, A. Hammerich, G. Joucard, W. Karrlein, H.-D. Meyer, N. Lipkin, O. Roncero, and R. Kosloff. A comparison of different propagation schemes for the time dependent Schrödinger equation. Journal of Computational Physics, 94:59–80, 1991.
- [66] K. Levenberg. A method for the solution of certain problems in least squares. Quart. Appl. Math, 2:164, 1944.
- [67] H. J. Lewandowski, D. M. Harber, D. L. Whitaker, and E. A. Cornell. Simplified system for creating a Bose-Einstein condensate. J. Low Temp. Phys., 132:309, 2003.
- [68] H. Lignier, C. Sias, D. Ciampini, Y. Singh, A. Zenesini, O. Morsch, and E. Arimondo. Dynamical control of matter-wave tunneling in periodic potentials. Phys. Rev. Lett., 99:220403, 2007.
- [69] C. Lovecchio, F. Schäfer, M. Ali Khan, J. Herrera, F. S. Cataliotti, T. Calarco, S. Montangero, and F. Caruso. Optimal preparation of quantum states on an atom-chip device. Phys. Rev. A, 93:010304(R), 2016.
- [70] K. B. MacAdam, A. Steinbach, and C. Wieman. A narrow-band tunable diode laser system with grating feedback and a saturated absorption spectrometer for Cs and Rb. Am. J. Phys., 60:12, 1992.
- [71] O. Mandel, M. Greiner, A. Widera, T. Rom, T. W. Hänsch, and I. Bloch. Coherent transport of neutral atoms in spin-dependent optical lattice potentials. Phys. Rev. Lett., 91:1, 2003.
- [72] D. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. SIAM J. Appl. Math, 11:431, 1963.
- [73] G. D. McDonald and C. C. N. Kuhn. Space-time area in atom interferometry.
- [74] G. D. McDonald, C. C. N. Kuhn, S. Bennetts, J. E. Debs, J. D. Close, and N. P. Robins. A faster scaling in acceleration-sensitive atom interferometers. Eur. Phys. Lett., 105:63001, 2014.
- [75] J. M. McGuirk, M. J. Snadden, and M. A. Kasevich. Large area light-pulse atom interferometry. Phys. Rev. Lett., 85(21):4498, 2000.

- [76] K. M. Mertes, J. W. Merrill, R. Carretero-González, D. J. Frantzeskakis, P. G. Kevrekidis, and D. S. Hall. Nonequilibrium dynamics and superfluid ring excitations in binary Bose-Einstein condensates. Phys. Rev. A, 99:190402, 2007.
- [77] H. J. Metcalf and P. van der Straten. Laser Cooling and Trapping. Springer-Verlag, 1999.
- [78] A. Miffre, M. Jacquy, M. Büchner, G. Tréneç, and J. Viguè. Atom interferometry. Phys. Scr., 74:C15, 2006.
- [79] J. Monnier. Optical interferometry in astronomy. Rep. Prog. Phys., 66:789, 2003.
- [80] J. Monnier. Optical distortion evaluation in large area windows using interferometry. 14th Int. Symp. on Nondestructive Characterization of Materials, 2015.
- [81] K. W. Moore, R. Chakrabarti, G. Riviello, and H. Rabitz. Search complexity and resource scaling for the quantum optimal control of unitary transformations. Phys. Rev. A, 83:012326, 2011.
- [82] O. Morsch, J. H. Müller, M. Cristiani, D. Ciampini, and E. Arimondo. Bloch oscillations and mean-field effects of Bose-Einstein condensates in 1D optical lattices. Phys. Rev. Lett., 87:140402, 2001.
- [83] T. Müller, M. Gilowski, M. Zaiser, P. Berg, C. Schubert, T. Wendrich, W. Ertmer, and E. M. Rasel. A compact dual atom interferometer gyroscope based on laser-cooled rubidium. Eur. Phys. J. D, 53:273, 2009.
- [84] H. Müttinga et al. Interferometry with Bose-Einstein condensates in microgravity. Phys. Rev. Lett., 110:093602, 2013.
- [85] J. Nayak. Optical distortion evaluation in large area windows using interferometry. 14th Int. Symp. on Nondestructive Characterization of Materials, 2015.
- [86] M. Ndong, H. Tal-Ezer, R. Kosloff, and C. P. Koch. A Chebyshev propagator for inhomogeneous Schrödinger equations. Annu. Rev. Phys. Chem., 45:145, 1994.
- [87] J. A. Nelder and R. Mead. A simplex method for function minimization. The Computer Journal, 7:4, 1965.
- [88] J.P. Palao and R. Kosloff. Optimal control theory for unitary transformations. Phys. Rev. A, 68:062308, 2003.
- [89] J.P. Palao, R. Kosloff, and C. P. Koch. Protecting coherence in optimal control theory: State-dependent constraint approach. Phys. Rev. A, 77:063412, 2008.
- [90] C. J. Pethick and H. Smith. Bose-Einstein Condensation in Dilute Gases. Cambridge University Press, 2008.
- [91] N. Poli, F.-Y. Yang, M. G. Tarillo, A. Alberti, M. Prevedelli, and G. M. Tino. Precision measurement of gravity with cold atoms in an optical lattice and comparison with a classical gravimeter. Phys. Rev. Lett., 106:038501, 2011.
- [92] S. Pötting, M. Cramer, and P. Meystre. Momentum-state engineering and control in Bose-Einstein condensates. Phys. Rev. A, 64:063613, 2001.

- [93] S. Pötting, M. Cramer, C. H. Schwalb, H. Pu, and P. Meystre. Coherent acceleration of Bose-Einstein condensates. Phys. Rev. A, 64:023604, 2001.
- [94] E. L. Raab, M. Prentiss, A. Cable, S. Chu, and D. E. Pritchard. Trapping of neutral sodium atoms with radiation pressure. Phys. Rev. Lett., 59:2631, 1987.
- [95] N. Rach, M. M. Müller, T. Calarco, and S. Montangero. Dressing the chopped-random-basis optimization: A bandwidth-limited access to the trap-free landscape. Phys. Rev. A, 92:062343, 2015.
- [96] M. Reitter, Jakob Näger, K. Wintersperger, C. Sträter, I. Bloch, A. Eckardt, and U. Schneider. Interaction dependent heating and atom loss in a periodically driven optical lattice. Phys. Rev. Lett., 119:200402, 2017.
- [97] F. Riehle, Th. Kisters, A. Witte., J. Helmcke, and Ch. J. Bordè. Optical Ramsey spectroscopy in a rotating frame: Sagnac effect in a matter-wave interferometer. Phys. Rev. Lett., 67(2):177, 1991.
- [98] W. Rohringer, R. Bücker, S. Manz, T. Betz, Ch. Koller, M. Göbel, A. Perrin, J. Schmiedmayer, and T. Schumm. Stochastic optimization of a cold atom experiment using a genetic algorithm. Appl. Phys. Lett., 93:264101, 2008.
- [99] E. A. Salim. Ultracold matter systems and atomtronics instrumentation. PhD thesis, University of Colorado at Boulder, 2011.
- [100] E. A. Salim, S. C. Caliga, J. B. Pfeiffer, and D. Z. Anderson. High resolution imaging and optical control of Bose-Einstein condensates in an atom chip magnetic trap. Appl. Phys. Lett., 102:084104, 2013.
- [101] R. E. Sapiro, R. Zhang, and G. Raithel. Reversible loss of superfluidity of a Bose-Einstein condensate in a 1D optical lattice. New J. Phys., 11:013013, 2009.
- [102] D. Schrader, S. Kuhr, W. Alt, M. Müller, V. Gomer, and D. Meschede. An optical conveyor belt for single neutral atoms. Appl. Phys. B, 73:819, 2014.
- [103] S. R. Segal. Progress towards an ultracold atomic Sagnac gyroscope. PhD thesis, University of Colorado at Boulder, 2010.
- [104] S. E. Sklarz and D. J. Tannor. Loading a Bose-Einstein condensate into an optical lattice: An application of optimal control theory to the nonlinear Schrödinger equation. Phys. Rev. A, 66:053619, 2002.
- [105] M. B. Squires. High repetition rate Bose-Einstein condensate production in a compact, transportable vacuum system. PhD thesis, University of Colorado at Boulder, 2008.
- [106] D.A. Steck. Rubidium 87 D line data. available online, 2015. revision 2.1.5.
- [107] C. J. E. Straatsma, M. K. Ivory, J. Duggan, J. Ramirez-Serrano, D. Z. Anderson, and E. A. Salim. On-chip optical lattice for cold atom experiments. Opt. Lett., 40:3368, 2015.
- [108] C. Sträter and A. Eckardt. Interband heating processes in a periodically driven optical lattice. Z. Naturforsch., 71:10, 2016.

- [109] S. S. Szigeti, J. E. Debs, J. J. Hope, N. P. Robins, and J. D. Close. Why momentum width matters for atom interferometry with Bragg pulses. *New J. Phys.*, 14:023009, 2012.
- [110] G. Tackmann, P. Berg, C. Schubert, S. Abend, M. Gilowski, W. Ertmer, and E. M. Rasel. Self-alignment of a compact large-area atomic Sagnac interferometer. *New J. Phys.*, 14:0152002, 2012.
- [111] H. Tal-Ezer and R. Kosloff. An accurate and efficient scheme for propagating the time dependent Schrödinger equation. *J. Chem. Phys.*, 81:3967, 1984.
- [112] Q. Thommen, J. C. Garreau, and V. Zehnlé. Theoretical analysis of quantum dynamics in one-dimensional lattices: Wannier-Stark description. *Phys. Rev. A*, 65:053406, 2002.
- [113] R. Trubko, J. Greenberg, M. T. St. Germaine, M. D. Gregoire, W.F. Holmgren, I. Hromada, and A. D. Cronin. Atom interferometer gyroscope with spin-dependent phase shifts induced by light near a tune-out wavelength. *Phys. Rev. Lett.*, 114:140404, 2015.
- [114] S. van Frank, M. Bonneau, J. Schmiedmayer, S. Hild, C. Gross, M. Cheneau, I. Bloch, T. Pichler, A. Negretti, T. Calarco, and S. Montangero. Optimal control of complex atomic quantum systems. *Nat. Sci. Rep.*, 6(34187):34187, 2016.
- [115] S. van Frank, A. Negretti, T. Berrada, R. Bücker, S. Montangero, J.-F. Schaff, T. Schumm, T. Calarco, and J. Schmiedmayer. Interferometry with non-classical motional states of a Bose-Einstein condensate. *Nat. Comm.*, 5:4009, 2014.
- [116] S. w. Chiow, T. Kovachy, H.-C. Chien, and M. A. Kasevich. $102\hbar k$ large area atom interferometers. *Phys. Scr.*, 91:053006, 2016.
- [117] Y.-J. Wang, D. Z. Anderson, V. M. Bright, E. A. Cornell, Q. Diot, T. Kishimoto, M. Prentiss, R. A. Saravanan, S .R. Segal, and S. Wu. Atom michelson interferometer on a chip using a Bose-Einstein condensate. *Phys. Rev. Lett.*, 94:090405, 2005.
- [118] C. A. Weidner and D. Z. Anderson. Experimental demonstration of shaken lattice interferometry. Submitted for review., 2018.
- [119] C. A. Weidner, H. Yu, R. Kosloff, and D. Z. Anderson. Atom interferometry using a shaken optical lattice. *Phys. Rev. A*, 95:043624, 2017.
- [120] M. Weinberg, C. Öschl ager, C. Str ater, S. Prella, A. Eckardt, K. Sengstock, and J. Simonet. Multiphoton interband excitations of quantum gases in driven optical lattices. *Phys. Rev. A*, 92:043621, 2015.
- [121] P. B. Wigley, P .J. Everitt, A. van dan Hengel, J. W. Bastian, M. A. Sooriyabandara, G. D. McDonald, K. S. Hardman, C. D. Quinlivan, P. Manju, C. C. N. Kuhn, I. R. Petersen, A. N. Luiten, J. J. Hope, N. P. Robins, and M. R. Hush. Fast machine-learning online optimization of ultra-cold-atom experiments. *Sci. Rep.*, 6:25890, 2016.
- [122] A. Zenesini, H. Lignier, D. Ciampini, O. Morsch, and E. Arimondo. Coherent control of dressed matter waves. *Phys. Rev. Lett.*, 102:100403, 2009.
- [123] X. Zhang, R. P. del Aguila, T. Mazzoni, N. Poli, and G. M. Tino. Trapped-atom interferometer with ultracold Sr atoms. *Phys. Rev. A*, 94:043608, 2016.